

Theoretical Aspects of Stochastic Signal Quantisation and Suprathreshold Stochastic Resonance

by

Mark Damian McDonnell

B.Sc. (Mathematical & Computer Sciences),
The University of Adelaide, Australia, 1997

B.E. (Electrical & Electronic Engineering, Honours),
The University of Adelaide, Australia, 1998

B.Sc. (Applied Mathematics, First Class Honours),
The University of Adelaide, Australia, 2001


Thesis submitted for the degree of

Doctor of Philosophy

in

The School of Electrical and Electronic Engineering,
Faculty of Engineering, Computer and Mathematical Sciences
The University of Adelaide, Australia

February, 2006



© 2006
Mark Damian McDonnell
All Rights Reserved



For Juliet

Contents

Heading	Page
Contents	v
Abstract	xv
Statement of Originality	xvii
Acknowledgments	xix
Thesis Conventions	xxiii
Publications	xxv
List of Figures	xxvii
List of Tables	xxxiii
Chapter 1. Introduction and Motivation	1
1.1 Background and Motivation	2
1.1.1 Stochastic Resonance and Sensory Neural Coding	3
1.1.2 Low Signal-to-Noise Ratio Systems and Sensor Networks	4
1.2 Specific Research Questions	5
1.2.1 From Stochastic Resonance to Stochastic Quantisation	5
1.3 Thesis Overview and Original Contributions	5
Chapter 2. Historical Landscape	9
2.1 Introduction	10
2.1.1 Chapter Structure	10
2.2 Stochastic Resonance	10
2.2.1 A Brief History of Stochastic Resonance	14
2.2.2 Stochastic Resonance in Threshold Systems	22

2.2.3	Stochastic Resonance in Neural Systems	34
2.3	Information and Quantisation Theory	35
2.3.1	Definition of Mutual Information	35
2.3.2	The Basics of Quantisation Theory	37
2.3.3	Quantisation Literature Review	39
2.4	Chapter Summary	42
2.4.1	Original Contributions for Chapter 2	42
2.4.2	Further Work	42
Chapter 3. How Should Stochastic Resonance be Measured?		45
3.1	Introduction	46
3.1.1	How Do I Measure Thee? Let Me Count The Ways	46
3.1.2	Chapter Structure	49
3.2	The SNR Gain Debate	49
3.2.1	Can SNR Gains Occur At All?	49
3.2.2	Are SNR Gains Meaningful?	51
3.3	The Data Processing Inequality and SR	52
3.3.1	The Data Processing Inequality	53
3.3.2	Example 1: Asymmetric Binary Channel	56
3.3.3	Example 2: Binary Erasure Channel	58
3.3.4	Discussion	63
3.4	Cross-Spectral Measures	64
3.4.1	Power and Cross-Power Spectral Densities	65
3.4.2	New Interpretation	71
3.4.3	Applying the Frequency Dependant SNR Formula	72
3.4.4	Discussion	75
3.5	Chapter Summary	78
3.5.1	Original Contributions for Chapter 3	78
3.5.2	Further Work	79
Chapter 4. Suprathreshold Stochastic Resonance: Encoding		81

4.1	Introduction	82
4.1.1	Chapter Structure	82
4.2	Literature Review	83
4.2.1	Information Theoretic Studies of SR in Threshold Systems	83
4.2.2	Original Work on SSR	83
4.2.3	SSR in Neurons and Cochlear Implant Encoding	84
4.2.4	Work on SSR by Other Authors	86
4.2.5	Similarity to SSR in Other Work	87
4.2.6	Discussion	88
4.3	Suprathreshold Stochastic Resonance	90
4.3.1	The SSR Model	90
4.3.2	Measuring SSR with Mutual Information	93
4.3.3	SSR for Specific Signal and Noise Distributions	99
4.3.4	Exact Result for Uniform Signal and Noise, and $\sigma \leq 1$	108
4.3.5	Numerical Results	112
4.4	Channel Capacity for SSR	126
4.4.1	Matched and Fixed Signal and Noise PDFs	127
4.4.2	Channel Capacity in General	129
4.4.3	Coding Efficiency	132
4.5	SSR as Stochastic Quantisation	133
4.5.1	Encoder Transfer Function	137
4.5.2	Decoding the SSR Output	138
4.6	Chapter Summary	139
4.6.1	Original Contributions for Chapter 4	140
4.6.2	Further Work	141
 Chapter 5. Suprathreshold Stochastic Resonance: Large N Encoding		 143
5.1	Introduction	144
5.1.1	Mutual Information	144
5.1.2	Literature Review	146
5.1.3	Chapter Structure	147

- 5.2 SSR for Large N and $\sigma = 1$ 148
 - 5.2.1 Average Conditional Entropy for Large N and $\sigma = 1$ 148
 - 5.2.2 Mutual Information for Large N and $\sigma = 1$ 150
- 5.3 $I(x, y)$ for Large N , Uniform Signal & Noise & $\sigma \leq 1$ 155
- 5.4 Mutual Information for Large N and Arbitrary σ 159
 - 5.4.1 The Gaussian Approximation to the Binomial Distribution 159
 - 5.4.2 Conditional Entropy for Large N 160
 - 5.4.3 Output Distribution and Entropy for Large N 165
 - 5.4.4 Mutual Information for Large N 172
 - 5.4.5 Relationship to Fisher Information 174
- 5.5 Channel Capacity for Large N 176
 - 5.5.1 General Channel Capacity 176
 - 5.5.2 Capacity for Matched Signal and Noise 177
 - 5.5.3 Gaussian Signal and Noise 178
 - 5.5.4 Numerical Verification and Other Distributions 180
- 5.6 Chapter Summary 181
 - 5.6.1 Original Contributions for Chapter 5 183
 - 5.6.2 Further Work 184

Chapter 6. Suprathreshold Stochastic Resonance: Decoding 187

- 6.1 Introduction 188
 - 6.1.1 Summary of Relevant Results from Chapters 4 and 5 188
 - 6.1.2 Decoding a Quantised Signal 189
 - 6.1.3 Chapter Structure 192
- 6.2 Averaging Without Thresholding 193
- 6.3 Linear Decoding Theory 195
 - 6.3.1 Constant Decoding 197
 - 6.3.2 Matched Moments Linear Decoding 198
 - 6.3.3 Wiener Optimal Linear Decoding 199
- 6.4 Linear decoding for SSR 200
 - 6.4.1 Decoding Measures for SSR 200

6.4.2	Specific Signal and Noise Distributions	202
6.4.3	Constant Decoding	210
6.4.4	Matched Moments Linear Decoding	213
6.4.5	Wiener Optimal Linear Decoding	214
6.5	Nonlinear Decoding Schemes	218
6.5.1	MAP Decoding	218
6.5.2	Maximum Likelihood Decoding	220
6.5.3	Minimum Mean Square Error Distortion Decoding	222
6.6	Decoding Analysis	228
6.6.1	Distortion Comparison	228
6.6.2	Transfer Functions	232
6.7	An Estimation Perspective	234
6.7.1	Definitions	234
6.7.2	Application to the SSR model	236
6.7.3	SSR Encoding	238
6.7.4	SSR Linear Decoding	239
6.7.5	SSR Nonlinear Decoding	240
6.8	Output Signal-to-Noise Ratio	246
6.8.1	SNR in the DIMUS model	247
6.8.2	Stocks' SNR Measure for SSR	248
6.8.3	Gammaitoni's Dithering Formula	251
6.9	Chapter Summary	251
6.9.1	Original Contributions for Chapter 6	252
6.9.2	Further Work	253
 Chapter 7. Suprathreshold Stochastic Resonance: Large N Decoding		 255
7.1	Introduction	256
7.1.1	Chapter Structure	256
7.2	Mean Square Error Distortion for Large N	256
7.2.1	Wiener Linear Decoding	257
7.2.2	MMSE Optimal Decoding	260

- 7.3 Large N Estimation Perspective 264
 - 7.3.1 Asymptotic Large N Theory 265
 - 7.3.2 Application to SSR 266
- 7.4 Discussion on Stochastic Resonance Without Tuning 268
- 7.5 Chapter Summary 270
 - 7.5.1 Original Contributions for Chapter 7 270
 - 7.5.2 Further Work 271

Chapter 8. Optimal Stochastic Quantisation 273

- 8.1 Introduction 274
 - 8.1.1 Optimal Quantisation Literature Review 275
 - 8.1.2 Chapter Structure 278
- 8.2 Optimal Quantisation Model 278
 - 8.2.1 Moment Generating Function 280
 - 8.2.2 Conditional Output Moments 281
 - 8.2.3 Transition Probabilities 282
 - 8.2.4 Optimisation Problem Formulation 284
- 8.3 Optimisation Solution Algorithms 284
 - 8.3.1 Unconstrained Multidimensional Optimisation 284
 - 8.3.2 Dealing with Local Optima 285
- 8.4 Optimal Quantisation for Mutual Information 287
 - 8.4.1 Absence of Noise 287
 - 8.4.2 Results in the Presence of Noise 289
- 8.5 Optimal Quantisation for MSE Distortion 296
 - 8.5.1 Linear Decoding 296
 - 8.5.2 Nonlinear Decoding 296
 - 8.5.3 Absence of Noise 297
 - 8.5.4 Results in the Presence of Noise 298
- 8.6 Discussion of Results 303
 - 8.6.1 Observations 303
 - 8.6.2 Mathematical Description of Optimal Thresholds 306

8.6.3	Local Maxima	308
8.6.4	An Estimation Perspective	313
8.7	Locating the Final Bifurcation	317
8.8	Chapter Summary	319
8.8.1	Original Contributions for Chapter 8	319
8.8.2	Further Work	320
 Chapter 9. SSR, Neural Coding, and Performance Tradeoffs		 323
9.1	Introduction	324
9.1.1	Review of Relevant Material	325
9.1.2	Constrained Optimisation	327
9.1.3	Chapter Structure	328
9.2	Information Theory and Neural Coding	328
9.2.1	Energy Constraints	329
9.2.2	All Thresholds Equal	331
9.2.3	Energy Constrained Optimal Quantisation	337
9.2.4	Fixed Thresholds and a Minimum Information Constraint	340
9.2.5	Discussion and Further Work	340
9.3	Rate-Distortion Tradeoff	342
9.3.1	Rate-Distortion Theory	343
9.3.2	Rate-Distortion Tradeoff for SSR	348
9.3.3	Rate-Distortion Tradeoff for Optimised Thresholds	352
9.4	Chapter Summary	353
9.4.1	Original Contributions for Chapter 9	354
9.4.2	Further Work	354
 Chapter 10. Conclusions and Future Directions		 357
10.1	Thesis Summary	358
10.1.1	Stochastic Resonance	358
10.1.2	Suprathreshold Stochastic Resonance	358
10.1.3	Stochastic Quantisation with Identical Thresholds	359
10.1.4	Optimal Stochastic Quantisation	359
10.1.5	Stochastic Quantiser Performance Tradeoffs	360
10.2	Closing Remarks	360

Appendix A. Binary Channel Calculations	363
A.1 Example 1: Asymmetric Binary Channel	363
A.2 Example 2: Chapeau-Blondeau’s Erasure Channel	365
Appendix B. Derivations for Suprathreshold Stochastic Resonance Encoding	367
B.1 Maximum Values and Modes of $P(n x)$	367
B.2 A Proof of Equation (4.38)	368
B.3 Distributions	369
B.3.1 Gaussian Signal and Noise	369
B.3.2 Uniform Signal and Noise	370
B.3.3 Laplacian Signal and Noise	370
B.3.4 Logistic Signal and Noise	371
B.3.5 Cauchy Signal and Noise	372
B.3.6 Exponential Signal and Noise	373
B.3.7 Rayleigh Signal and Noise	374
B.4 Proofs for Specific Cases that $Q(\tau)$ is a PDF	375
B.4.1 Gaussian Signal and Noise	375
B.4.2 Laplacian Signal and Noise	375
B.5 Numerical Integration of the Mutual Information	376
B.5.1 Integrating Over the Input’s Support	376
B.5.2 Numerical Integration Using $Q(\tau)$	377
B.5.3 Comments on Numerical Integration	377
Appendix C. Derivations for large N Suprathreshold Stochastic Resonance	379
C.1 Proof of Eqn. (5.9)	379
C.2 Derivation of Eqn. (5.13)	380
C.3 Proof that $S(x)$ is a PDF	381
Appendix D. Derivations for Suprathreshold Stochastic Resonance Decoding	383
D.1 Conditional Output Moments	383
D.2 Output Moments	384
D.2.1 Even Signal and Noise PDFs, All Thresholds Zero	385

D.3 Correlation and Correlation Coefficient Expressions	386
D.3.1 Input Correlation Coefficient at Any Two Thresholds	386
D.3.2 Output Correlation Coefficient at Any Two Thresholds	387
D.3.3 Input-Output Correlation	387
D.4 A Proof of Prudnikov’s Integral	388
D.5 Minimum Mean Square Error Distortion Decoding	391
D.5.1 MMSE Reconstruction Points and Distortion	391
D.5.2 MMSE Decoded Output is Uncorrelated with the Error	392
D.5.3 Relationship of MMSE to Backwards Conditional Variance	393
D.6 Fisher Information	394
D.6.1 First Derivation	394
D.6.2 Second Derivation	395
D.7 Proof of the Information and Cramer-Rao Bounds	396
Bibliography	399
List of Acronyms	421
Index	423
Biography	425

Abstract

Quantisation of a signal or data source refers to the division or classification of that source into a discrete number of categories or states. It occurs, for example, when analog electronic signals are converted into digital signals, or when a large amount of data is binned into histograms. By definition, quantisation is a lossy process, which compresses data into a more compact representation, so that the number of states in a quantiser's output are usually far fewer than the number of possible input values. Most existing theory on the performance and design of quantisation schemes specify only deterministic rules governing how data is quantised.

By contrast, *stochastic quantisation* is a term intended to pertain to quantisation where the rules governing the assignment of input values to output states are *stochastic*, rather than *deterministic*. One form of stochastic quantisation that has already been widely studied is *dithering*. However, the stochastic aspect of dithering is usually restricted so that it is equivalent to adding random noise to a signal, prior to quantisation. The term *stochastic quantisation* is intended to be far more general, and apply to the situation where the *rules* of the quantisation process are stochastic.

The inspiration for this study comes from a phenomenon known as *stochastic resonance*, which is said to occur when the presence of noise in a system provides a better performance than the absence of noise. Specifically, this thesis discusses a particular form of stochastic resonance known as suprathreshold stochastic resonance, which occurs in an array of identical, but independently noisy threshold devices, and demonstrates how this effect is essentially a form of stochastic quantisation.

The motivation for this study is two fold. Firstly, stochastic resonance has been observed in many forms of neurons and neural systems, both in models and in real physiological experiments. The model in which suprathreshold stochastic resonance occurs was designed to model a *population of neurons*, rather than a single neuron. Unlike single neurons, the suprathreshold stochastic resonance model supports stochastic resonance for input signals that are not entirely or predominantly subthreshold. Hence, it has been conjectured that the suprathreshold stochastic resonance effect is utilised by populations of neurons to encode noisy sensory information, for example, in the cochlear nerve.

Secondly, although stochastic resonance has been observed in many different systems, in a wide variety of scientific fields, to date very few applications inspired by stochastic resonance have been proposed. One of the reasons for this is that in many circumstances, utilising stochastic resonance to improve a system is sub-optimal when compared to systems optimised to operate without requiring stochastic resonance. However, given that stochastic resonance is so widespread in nature, and that many new technologies have been inspired by natural systems—particularly biological systems—applications incorporating aspects of stochastic resonance may yet prove revolutionary in fields such as distributed sensor networks, nano-electronics and biomedical prosthetics.

Hence, as a necessary step towards confirming the above two hypotheses, this thesis addresses in detail for the first time various theoretical aspects of stochastic quantisation, in the context of the suprathreshold stochastic resonance effect. The original work on suprathreshold stochastic resonance considers the effect from the point of view of an information channel. This thesis comprehensively reviews all such previous work. It then extends such work in several ways; firstly, it considers the suprathreshold stochastic resonance effect as a form of stochastic quantisation; secondly it considers stochastic quantisation in a model where all threshold devices are not necessarily identical, but are still independently noisy; and thirdly, it considers various constraints and tradeoffs in the performance of stochastic quantisers.

Statement of Originality

This work contains no material that has been accepted for the award of any other degree or diploma in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text.

I give consent to this copy of my thesis, when deposited in the University Library, being made available in all forms of media, now or hereafter known.

20 February 2006

Signed

Date

Acknowledgments

Completing a PhD and writing a thesis is a lengthy process, and there are many people and organisations to thank.

Firstly, I would like to express my sincere gratitude to my primary supervisor, Professor Derek Abbott. Thankyou for all your advice, and your never-ending enthusiasm, positivity, encouragement, support and understanding. Completing a PhD is not only about performing quality research and fulfilling the academic requirements; I have found that a good supervisor, as well as facilitating such academic matters, is also one who can always make time to discuss research; one who goes beyond the call of duty to provide support when the twists and turns of life make research difficult; encourages the writing of papers and attendance at conferences; gently nudges towards productivity during times of slackness; and generally helps minimise the frustration of red-tape. Derek, your assistance in all these matters greatly facilitated the research described in this thesis.

Thankyou also to my co-supervisor, Elder Professor Charles Pearce, of the Applied Mathematics discipline, in the School of Mathematical sciences. I have greatly appreciated your insightful comments on aspects of the mathematics presented in this work, your general encouragement, and feedback on my thesis.

My warmest appreciation also goes to Dr Nigel Stocks, of the School of Engineering, at The University of Warwick, UK. I have been fortunate enough to visit Warwick three times through the course of this PhD, where I have thoroughly enjoyed our lengthy discussions. Thankyou very much for your collaboration, and I look forward to continuing post-PhD. While at Warwick university, I have also appreciated the assistance of Nigel's postdoctoral researchers, David Allingham, and Sasha Nikitin.

Thankyou also to Nigel's collaborator, Dr Rob Morse, formerly of The Mackay Institute of Communication and Neuroscience, at Keele University, UK, and now at Aston University, UK. I very much appreciate you giving me a tour of your laboratory, and taking the time for several stimulating discussions.

I would also like to thank several others with whom I have collaborated during this PhD. Firstly, Ferran Martorell, who visited my university throughout 2004. Thankyou

Acknowledgments

for our ongoing discussions. Thankyou also to Antonio Rubio, Lazslo Kish and Swaminathan Sethuraman.

Within the School of Electrical and Electronic Engineering, I would like to express my gratitude to the Head of Department, Mike Liebelt, Laboratory Manager, Stephen Guest, and to the administrators, Yadi Parrot, Colleen Greenwood, Ivana Rebellato, and Rose-Marie Descalzi. You have all always been very happy to drop everything and help whenever I have had a question or request, and just generally made life very easy. Also, thankyou to the computer support staff, David Bowler and his team, for your continuous work in making the computer systems run smoothly.

Thanks also to those with whom I have had technical discussions from time to time, including David Haley, Andrew Allison, Said Al-Sarawi, David O'Carroll, Bruce Davis, Riccardo Mannella and Stuart Nankivell.

It has been very enjoyable having an office in the Centre for Biomedical Engineering Laboratory. Many thanks to all the co-students with whom I have enjoyed coffee and beers, including (in chronological order) Leonard Hall, Greg Harmer, Sam Mickan, Brad Ferguson, Sreeja Rajesh, Leo Lee, Dominic Osborne, Matthew Berryman, Adrian Flitney, Gretel Png, Inke Jones and Frederic Bonnet. Also to those who aren't actually in the lab, but still have been there for the occasional coffee and beer—Brian Ng, Ben Jamali, Nigel Brine, Bobby Yau, and Damath Ranasinghe. Apologies to anyone I've temporarily forgotten!

Special thanks must go to Matthew Berryman, who proof read my thesis and provided many comments that greatly improved its clarity. I owe you, Matt.

Writing this thesis has been greatly assisted by music from The Killers and Coldplay.

Thankyou very much to the following agencies, who have either provided me with living scholarships for this PhD, or provided funding for travel to overseas collaborations or conferences: The Department of Education, Science and Training (DEST); the Cooperative Research Centre for Sensor, Signal and Information Processing (CSSIP); The University of Adelaide; The School of Electrical and Electronic Engineering at The University of Adelaide; The Australian Research Council (ARC); The Australian Academy of Science (AAS); The Australian Federation of University Women (AFUW); The Institute of Electrical and Electronics Engineers (IEEE), SA Section; The Santa Fe Institute; The IEEE International Symposium on Information Theory (ISIT) 2004 and

the United States Air Force Research Laboratory; the Leverhulme Trust; and finally, the DR Stranks Travelling Fellowship.

I have also been employed from time to time during this PhD, both part-time and during a leave-of-absence. I greatly appreciate the support of Neal Young and Ian Allison, both from the Australian Antarctic Division, as well as Chris Coleman, for your support in providing me with this work.

Many thanks also to Stuart Nankivell for some photoshop magic.

Finally, I wish to acknowledge the greatest helpers of all during this PhD—my family. Thankyou to Mum, Dad, Grandma, Andrew, and the rest of my family for providing me with the upbringing and family life that has brought me to this point. With all my love.

And last, but foremost, all my love and gratitude to my wife, Juliet. This thesis is dedicated to you.

– Mark D. McDonnell

NOTE:

This figure/table/image has been removed to comply with copyright regulations. It is included in the print copy of the thesis held by the University of Adelaide Library.

Derek Abbott's 'Stochastic group' at The University of Adelaide. Photo taken in Santa Fe, USA, 19 June 2003. From left: Adrian Flitney, Mark McDonnell, Andrew Allison, Derek Abbott and Matthew Berryman.

NOTE:

This figure/table/image has been removed to comply with copyright regulations. It is included in the print copy of the thesis held by the University of Adelaide Library.

Nigel Stocks' Warwick University 'stochastic resonance' group. Photo taken at Warwick University, UK, August 2005. From left: Rob Morse, Nigel Stocks, Mark McDonnell, David Allingham and Sasha Nikitin.

NOTE:

This figure/table/image has been removed to comply with copyright regulations. It is included in the print copy of the thesis held by the University of Adelaide Library.

Dr Rob Morse, and his lab at The Mackay Institute of Communication and Neuroscience, at Keele University, UK, 17 August 2005.

Thesis Conventions

Typesetting This thesis is typeset using the $\text{\LaTeX}2\text{e}$ software. Processed plots and images were generated using Matlab 6.1 (Mathworks Inc.). WinEdt build 5.3 was used as an effective interface to the Miktex version of \LaTeX .

Spelling Australian English spelling is adopted, as defined by the Macquarie English Dictionary (Delbridge *et al.* 1997).

Referencing The Harvard style is used for referencing and citation in this thesis.

Publications

Book Chapters

MCDONNELL-M. D., STOCKS-N. G., PEARCE-C. E. M., ABBOTT-D. (2006). Information transfer through parallel neurons with stochastic resonance, *Emerging Brain-Inspired Nano-Architectures*, Eds. V. Beiu & U. Rueckert, Imperial College Press (accepted 14 Jun. 2005).

Journal Publications

MCDONNELL-M. D., STOCKS-N. G., PEARCE-C. E. M., AND ABBOTT-D. (2006). Optimal information transmission in nonlinear arrays through suprathreshold stochastic resonance, *Physics Letters A* (accepted 16 Nov. 2005, In Press).

MCDONNELL-M. D., STOCKS-N. G., PEARCE-C. E. M., AND ABBOTT-D. (2005). Quantization in the presence of large amplitude threshold noise, *Fluctuation and Noise Letters*, **5**, pp. L457–L468.

MCDONNELL-M. D., STOCKS-N. G., PEARCE-C. E. M., AND ABBOTT-D. (2003). Stochastic resonance and the data processing inequality, *Electronics Letters (IEE)*, **39**, pp. 1287–1288.

MCDONNELL-M. D., ABBOTT-D., AND PEARCE-C. E. M. (2002). A characterization of suprathreshold stochastic resonance in an array of comparators by correlation coefficient, *Fluctuation and Noise Letters*, **2**, pp. L205–L220.

MCDONNELL-M. D., ABBOTT-D., AND PEARCE-C. E. M. (2002). An analysis of noise enhanced information transmission in an array of comparators, *Microelectronics Journal*, **33**, pp. 1079–1089.

Conference Publications

MCDONNELL-M. D., STOCKS-N. G., PEARCE-C. E. M., AND ABBOTT-D. (2005). How to use noise to reduce complexity in quantization, in A. Bender. (ed.), *Proc. SPIE Complex Systems*, Brisbane, Australia, Vol. 6039, pp. 115–126.

MCDONNELL-M. D., STOCKS-N. G., PEARCE-C. E. M., AND ABBOTT-D. (2005). Optimal quantization and suprathreshold stochastic resonance, in N. G. Stocks., D. Abbott., and R. P. Morse. (eds.), *Proc. SPIE Noise in Biological, Biophysical, and Biomedical Systems III*, Austin, USA, Vol. 5841, pp. 164–173.

MCDONNELL-M. D., STOCKS-N. G., PEARCE-C., AND ABBOTT-D. (2005). Analog to digital conversion using suprathreshold stochastic resonance, in S. F. Al-Sarawi. (ed.), *Proc. SPIE Smart Structures, Devices, and Systems II*, Sydney, Australia, Vol. 5649, pp. 75–84.

MCDONNELL-M. D., AND ABBOTT-D. (2004). Optimal quantization in neural coding, *Proc. IEEE International Symposium on Information Theory*, Chicago, USA, p. 496.

- MCDONNELL-M. D.**, AND **ABBOTT-D.** (2004). Signal reconstruction via noise through a system of parallel threshold nonlinearities, *Proc. 2004 IEEE International Conference on Acoustics, Speech, and Signal Processing*, Montreal, Canada, Vol. 2, pp. 809–812.
- MCDONNELL-M. D.**, **STOCKS-N. G.**, **PEARCE-C. E. M.**, AND **ABBOTT-D.** (2004). Optimal quantization for energy-efficient information transfer in a population of neuron-like devices, in Z. Gingl., J. M. Sancho., L. Schimansky-Geier., and J. Kertesz. (eds.), *Proc. SPIE Noise in Complex Systems and Stochastic Dynamics II*, Maspalomas, Spain, Vol. 5471, pp. 222–232.
- MCDONNELL-M. D.**, **SETHURAMAN-S.**, **KISH-L. B.**, AND **ABBOTT-D.** (2004). Cross-spectral measurement of neural signal transfer, in Z. Gingl., J. M. Sancho., L. Schimansky-Geier., and J. Kertesz. (eds.), *Proc. SPIE Noise in Complex Systems and Stochastic Dynamics II*, Maspalomas, Spain, Vol. 5471, pp. 550–559.
- MCDONNELL-M. D.**, AND **ABBOTT-D.**, AND **PEARCE-C. E. M.** (2003). The data processing inequality and stochastic resonance, in L. Schimansky-Geier., D. Abbott., A. Neiman., and C. Van den Broeck. (eds.), *Proc. SPIE Noise in Complex Systems and Stochastic Dynamics*, Santa Fe, USA, Vol. 5114, pp. 249–260.
- MCDONNELL-M. D.**, **STOCKS-N. G.**, **PEARCE-C. E. M.**, AND **ABBOTT-D.** (2003). Neural mechanisms for analog to digital conversion, in D. V. Nicolau. (ed.), *Proc. SPIE BioMEMS and Nanoengineering*, Perth, Australia, Vol. 5275, pp. 278–286.
- MCDONNELL-M. D.**, AND **ABBOTT-D.** (2002). Open questions for suprathreshold stochastic resonance in sensory neural models for motion detection using artificial insect vision, in S. M. Bezrukov. (ed.), *UPoN 2002: Third International Conference on Unsolved Problems of Noise and Fluctuations in Physics, Biology, and High Technology*, Washington D.C., USA, Vol. 665, American Institute of Physics, pp. 51–58.
- MCDONNELL-M. D.**, **STOCKS-N. G.**, **PEARCE-C. E. M.**, AND **ABBOTT-D.** (2002). Maximising information transfer through nonlinear noisy devices, in D. V. Nicolau. (ed.), *Proc. SPIE Biomedical Applications of Micro and Nanoengineering*, Melbourne, Australia, Vol. 4937, pp. 254–263.
- MCDONNELL-M. D.**, **PEARCE-C. E. M.**, AND **ABBOTT-D.** (2001). Neural information transfer in a noisy environment, in N. W. Bergmann. (ed.), *Proc. SPIE Electronics and Structures for MEMS II*, Adelaide, Australia, Vol. 4591, pp. 59–69.
- MARTORELL-F.**, **MCDONNELL-M. D.**, **RUBIO-A.**, AND **ABBOTT-D.** (2005). Using noise to break the noise barrier in circuits, in S. F. Al-Sarawi. (ed.), *Proc. SPIE Smart Structures, Devices, and Systems II*, Sydney, Australia, Vol. 5649, pp. 53–66.
- MARTORELL-F.**, **MCDONNELL-M. D.**, **ABBOTT-D.**, AND **RUBIO-A.** (2004). Generalized noise resonance: Using noise for signal enhancement, in D. Abbott., S. M. Bezrukov., A. Der., and A. Sánchez. (eds.), *Proc. SPIE Fluctuations and Noise in Biological, Biophysical, and Biomedical Systems II*, Maspalomas, Spain, Vol. 5467, pp. 163–174.

List of Figures

Figure		Page
1.1	Typical stochastic resonance plot	2
<hr/>		
2.1	Frequency of stochastic resonance papers by year	12
2.2	Qualitative illustration of SR occurring in a simple threshold based system	25
2.3	Thresholding a periodic signal at its mean	30
2.4	Threshold SR for a periodic, but not sinusoidal signal	32
2.5	Ensemble averages of thresholded and unthresholded signals	33
<hr/>		
3.1	Generic nonlinear system	54
3.2	Channel capacity for several channels	61
3.3	Ratio of C_2 to C_1	62
3.4	Bandpass signal example	74
3.5	PSD of bandpass signal	74
3.6	PSDs and CSD of input and output signals	76
3.7	Output SNR against frequency	76
3.8	Output SNR against input noise	77
3.9	Coherence magnitude against input noise	77
<hr/>		
4.1	Array of N noisy threshold devices	85
4.2	Mutual information for the case of $P(x) = R(\theta - x)$	98
4.3	Various probability distributions	102
4.4	$Q(\tau)$ for various even PDFs and $\theta = 0$	104
4.5	Relative entropy between $P(x)$ and $R(x)$	105
4.6	PDF of a random phase sine wave	106
4.7	Exact $I(x, y)$ for uniform signal and noise and $\sigma \leq 1$	111

List of Figures

4.8	$I(x, y)$, $H(y x)$ and $H(y)$ for Gaussian signal and noise	114
4.9	$I(x, y)$, $H(y x)$ and $H(y)$ for uniform signal and noise	115
4.10	$I(x, y)$, $H(y x)$ and $H(y)$ for Laplacian signal and noise	115
4.11	$I(x, y)$, $H(y x)$ and $H(y)$ for logistic signal and noise	116
4.12	$I(x, y)$, $H(y x)$ and $H(y)$ for Cauchy signal and noise	116
4.13	$I(x, y)$, $H(y x)$ and $H(y)$ for $N = 127$ and five signal and noise pairs with even PDFs	117
4.14	$I(x, y)$, $H(y x)$ and $H(y)$ for Rayleigh signal and noise	119
4.15	$I(x, y)$, $H(y x)$ and $H(y)$ for exponential signal and noise	119
4.16	$I(x, y)$, $H(y x)$ and $H(y)$ for $N = 127$ and the Rayleigh and exponential cases	120
4.17	$I(x, y)$, $H(y x)$ and $H(y)$ for a uniform signal and Gaussian noise	122
4.18	$I(x, y)$, $H(y x)$ and $H(y)$ for a Laplacian signal and Gaussian noise	122
4.19	$I(x, y)$, $H(y x)$ and $H(y)$ for a logistic signal and Gaussian noise	123
4.20	$I(x, y)$, $H(y x)$ and $H(y)$ for a Cauchy signal and Gaussian noise	123
4.21	$I(x, y)$, $H(y x)$ and $H(y)$ for a Rayleigh signal and Gaussian noise	124
4.22	$I(x, y)$, $H(y x)$ and $H(y)$ for $N = 127$, various signals and Gaussian noise	125
4.23	Optimal σ for Gaussian, Laplacian and logistic signal and noise	128
4.24	Channel capacity for matched signal and noise	130
4.25	Coding efficiency for matched signal and noise	134
4.26	Random threshold distribution for Gaussian signal and noise	136
4.27	Average transfer function, and its variance for SSR	138
<hr/>		
5.1	Mutual information for $\sigma = 1$ and large N	153
5.2	Error in large N approximations to $I(x, y)$ for $\sigma = 1$	154
5.3	Error comparison for large N approximations to $I(x, y)$ for $\sigma = 1$	154
5.4	Mutual information for uniform signal and noise and large N	157
5.5	Error in mutual information for uniform signal and noise and large N	158
5.6	Error in the large N approximation to $\hat{H}(y x)$	163
5.7	Approximate $H(y x)$ for large N	164

5.8	Large N approximation to $P_y(n)$	169
5.9	Absolute error in output entropy using large N approximation to $P_y(n)$.	170
5.10	Large N approximation to output entropy	173
5.11	Large N approximation to mutual information	175
5.12	Approximation to $g(\sigma) = \int_{x=-\infty}^{x=\infty} P(x) \log_2(P_{1 x}) dx$	180
5.13	Finding channel capacity for large N	182
<hr/>		
6.1	Quantities for linear decoding	208
6.2	Linear decoding correlation coefficient	209
6.3	Linear decoding correlation coefficient for $N = 127$	210
6.4	Constant decoding MSE distortion for uniform signal and noise	212
6.5	MSE distortion for matched moments decoding	213
6.6	Reconstruction points for matched moments linear decoding, $N = 15$. .	215
6.7	MSE distortion for optimal linear (Wiener) decoding	216
6.8	Reconstruction points for optimal linear (Wiener) decoding, $N = 15$. . .	217
6.9	MSE distortion for maximum <i>a posteriori</i> decoding	220
6.10	Reconstruction points for maximum <i>a posteriori</i> decoding, $N = 15$	221
6.11	MSE distortion for MMSE decoding	224
6.12	Reconstruction points for MMSE decoding, $N = 15$	224
6.13	Correlation coefficient for MMSE decoding	225
6.14	MMSE decoding for uniform signal and noise	227
6.15	Comparison of MSE distortion for various decoding schemes, $N = 127$.	229
6.16	Comparison of correlation coefficients for $N = 127$	230
6.17	Decoding comparison for uniform signal and noise, $N = 127$	230
6.18	SQNR for each decoding for $N = 127$	231
6.19	Decoding transfer functions and variance	233
6.20	Average information bound	243
6.21	Percentage difference between MMSE distortion and average information bound	243
6.22	Mean square bias term	245

List of Figures

6.23	Average error variance term	245
<hr/>		
7.1	MSE distortion and correlation coefficient for Wiener decoding and large N	260
<hr/>		
8.1	Optimal thresholds for mutual information, $N = 2$	290
8.2	Optimal thresholds for mutual information, $N = 3$	291
8.3	Optimal thresholds for mutual information, $N = 4$	292
8.4	Optimal thresholds for mutual information, $N = 5$	293
8.5	Optimal reconstruction points for mutual information, $N = 5$	294
8.6	Optimal thresholds for mutual information, Gaussian signal and noise, $N = 15$	295
8.7	Optimal thresholds for linear decoding MSE distortion, $N = 5$	299
8.8	Optimal thresholds for MMSE distortion, $N = 2$	300
8.9	Optimal thresholds for MMSE distortion, $N = 5$	301
8.10	Optimal reconstruction points for minimised MMSE, $N = 5$	302
8.11	Optimal mutual information threshold clusters for various σ , $N = 5$	305
8.12	Locally optimal thresholds for maximum mutual information and $N = 3$	310
8.13	Locally optimal mutual information, $N = 3$	311
8.14	Difference in mutual information between local optima	312
8.15	Average Information Bound (AIB), Gaussian signal and noise, $N = 3$	316
8.16	Final bifurcation point, mutual information	318
8.17	Final bifurcation point, MMSE distortion	318
<hr/>		
9.1	Information efficiency as a function of σ , all thresholds equal	334
9.2	Information efficiency as a function of θ/σ_x , all thresholds equal	335
9.3	Information efficiency for fixed energy	336
9.4	Information efficiency for optimised thresholds	338
9.5	Optimal energy constrained thresholds	339

9.6	Minimising N subject to an information constraint	341
9.7	Operational rate-distortion tradeoff for SSR	349
9.8	Operational rate-distortion tradeoff for SSR, in dB	350
9.9	Comparison of $I(x, y)$ with $R_L(D)$ for SSR	351
9.10	Rate-distortion tradeoff for uniform signal and noise	351
9.11	Operational rate-distortion tradeoff for arbitrary thresholds	353

List of Tables

Table		Page
4.1	List of probability distributions and corresponding PDFs	101
4.2	List of ratios of signal PDF to noise PDF, $Q(\tau)$, for $\theta = 0$	103
5.1	Large N channel capacity and optimal σ for matched signal and noise .	181
8.1	Parameters for point density function, $N = 5$, Gaussian signal and noise, and maximised mutual information	308
9.1	Entropy and Shannon lower bound for various distributions	347

Chapter 1

Introduction and Motivation

WE begin by briefly outlining the background and motivation for the research described in this thesis, before giving an overview of each chapter, and pointing out the most significant contributions.

1.1 Background and Motivation

Although the methodology used to perform the research contained in this thesis is firmly within the fields of signal processing and applied mathematics, the motivation is interdisciplinary in nature.

The initial open questions that inspired this PhD research were (i) how might neurons make use of a phenomenon known as *stochastic resonance*, and (ii) how might a path towards engineering applications inspired by these studies be initiated?

Stochastic resonance (SR) is a counter-intuitive phenomenon where the presence of *noise* in a system is essential for optimal system performance. It is an effect that has been observed to occur in many systems, including both in neurons, and in electronic circuits. The term ‘resonance’ was originally used because the signature feature of SR is that a plot of output signal-to-noise ratio (SNR) has a single maximum, for some nonzero input noise intensity. Such a plot, as shown in Fig. 1.1, has a similar appearance to frequency dependent systems that have a maximum SNR, or output response, for some *resonant frequency*. However, in the case of SR, the resonance is ‘noise induced,’ rather than at a particular frequency.




Figure 1.1. Typical stochastic resonance plot. This plot shows the typical curve of output SNR vs input noise magnitude, for systems capable of stochastic resonance. For small and large noise values, the output SNR is very small, while some intermediate nonzero noise value provides the maximum output SNR.

The motivation for the research undertaken in this thesis is that, since we know the brain is far better at many tasks compared to electronic and computing devices, then maybe we can learn something from the brain. If we can ultimately better understand

the possible role of SR in the brain and nervous system, we may also be able to improve aspects of electronic systems.

1.1.1 Stochastic Resonance and Sensory Neural Coding

Although it is important to have an overall vision, in practical terms it is necessary to consider a concrete starting point. This research therefore began by investigating an exciting new development in the field of SR, known as *suprathreshold stochastic resonance* (SSR) (Stocks 2000a). Suprathreshold stochastic resonance occurs in an array of simple threshold devices. Each individual threshold device receives the same signal, but is subject to *independent* additive input noise. The output of each device is a binary signal, which is unity if the input is greater than the threshold value, and zero otherwise. The overall output of the *SSR model* is the *sum* of the individual binary signals. Originally, all threshold devices were considered to have the same threshold value.

Previous research into SSR considers the effect from the point of view of information transmission, and the model in which SSR occurs as a communication channel. Furthermore, the model in which SSR occurs was originally inspired by questions of sensory neural coding in the presence of noise. Unlike previously studied forms of SR, either in neurons or simple threshold based systems, SR can occur in such a model for signals that are not entirely or predominantly subthreshold. Although each threshold device in the SSR model is very simple, in comparison with more realistic neural models, such devices have actually been used to model neurons, under the name of the McCulloch-Pitts neural model (McCulloch and Pitts 1943).

The aim of the research undertaken for this thesis was to comprehensively investigate all known theoretical and numerical results on SSR, and extend this theory. It is anticipated that the results presented here will form a launching pad for future research into the specific role that SSR may play in real neural coding. Note that the goal is not to prove that living systems actively exploit SR or SSR—these are ongoing research areas, in the domain of neurophysiology and biophysics. Rather, our starting point is the theoretical mathematical underpinning of SR-type effects in the very simple McCulloch-Pitts model.

This is in-line with the time-honoured approach in physics and engineering: namely, to begin with the simplest possible model. As this thesis unfolds, it will be seen that the analysis of SR in arrays of such simplified neural models gives rise to rich complex

1.1 Background and Motivation

phenomena and also to a number of surprises. Explanation, or indeed, discovery, of these effects would not have been tractable if the starting point was an intricate neural model. Analysis of the simple model lays the foundation for adding further complexities in the future.

The intention is that the mathematical foundation provided by this thesis will assist future neurophysiologists into asking the right questions and in performing the right experiments when establishing if real neurons actively exploit SSR. In the meantime, this thesis also contributes to the application of SSR in artificial neural and electronic systems.

1.1.2 Low Signal-to-Noise Ratio Systems and Sensor Networks

Another motivation for this research is the important problem of overcoming the effects of noise in sensors and signal and data processing applications. For example, microelectronics technologies are shrinking, and are beginning to approach the nanoscale level (Martorell *et al.* 2005). At this scale, device behaviour can change and noise levels can approach signal levels. For such small SNRs, it may be impossible for traditional noise reduction methods to operate, and it may be that optimal circuit design needs to make use of the effects of SR.

A second area of much current research is that of distributed sensor networks (Akyildiz *et al.* 2002, Pradhan *et al.* 2002, Chong and Kumar 2003, Iyengar and Brooks 2004, Martinez *et al.* 2004, Xiong *et al.* 2004). Of particular interest to this thesis is the problem where it is not necessarily a network of complete sensors that is distributed, but it is actually the data acquisition or compression that is distributed (Berger *et al.* 1996, Draper and Wornell 2004, Pradhan *et al.* 2002, Xiong *et al.* 2004, Pradhan and Ramchandran 2005). A key aspect of such a scenario is that data is acquired from a number of independently noisy sources, that do not cooperate, and is then fused by some central processing unit. In the information theory and signal processing literature, this is referred to as *distributed source coding* or *distributed compression*.

Given that the SSR effect overcomes a serious limitation of all previously studied forms of SR, a complete theoretical investigation of its behaviour may lead to new design approaches to low SNR systems, or data acquisition and compression in distributed sensor networks.

1.2 Specific Research Questions

1.2.1 From Stochastic Resonance to Stochastic Quantisation

During the course of the research contained in this thesis, it became apparent that the SSR effect is equivalent to a noisy, or stochastic, quantisation of a signal (Widrow *et al.* 1996). Consequently, as well as describing its behaviour from the perspective of information transmission, it is equally valid to describe it from the perspective of *information compression*, or more specifically, *lossy compression*. Note that quantisation of a signal is a form of lossy compression (Berger and Gibson 1998, Gray and Neuhoff 1998).

The distinguishing feature of SSR that sets it apart from conventional forms of quantisation, is that conventionally, the rules that specify a quantiser's operation are considered to be fixed and deterministic. In contrast, when the SSR effect is viewed as quantisation, the governing rules are a set of independent random variables. Hence, we often refer to the SSR model's output as a *stochastic quantisation*.

Given this perspective, there are three immediate research questions that can be asked:

- Can we describe the SSR effect in terms of conventional quantisation and compression theory?
- Given that a question addressed in all SR research is that of finding the optimal noise conditions, what noise intensity optimises the performance of the SSR model, when it is described as a quantiser?
- How good is the SSR effect at quantisation when compared with conventional quantisers?

The underlying theme of this thesis is to address these three questions.

1.3 Thesis Overview and Original Contributions

This thesis consists of ten chapters, as follows:

- Chapter 1, the current chapter, provides the background and motivation for the work described in this thesis, and gives an overview of its original contributions to knowledge.

1.3 Thesis Overview and Original Contributions

- Chapter 2 contains a detailed overview of the historical landscape against which this thesis is set. It defines *stochastic resonance* and *quantisation* and gives a detailed literature review of SR, with particular emphasis on aspects relevant to quantisation. Chapter 2 is deliberately sparse in equations and devoid of quantitative results, but does provide qualitative illustrations of how SR works.
- Chapter 3 provides results that are somewhat peripheral to the main scope of this thesis, but that will prove useful for readers unfamiliar with, or confused about SR.
- Chapters 4 and 5 begin the main focus of the thesis, by defining the SSR model, giving a detailed literature review of all previous research on SSR, and replicating all the most significant theoretical results to date. These two Chapters consider only the original concept of the SSR model as a communications channel. In particular, we examine how the mutual information between the input and output of the SSR model varies with noise intensity. A subset of these results pertain specifically to a large number of individual threshold devices in the SSR model. Chapter 5 is devoted to replicating all such previous results, as well as developing new results in this area, while Chapter 4 focuses on more general behaviour.
- Chapters 6 and 7 contain original work on the description of the SSR model as a quantiser. There are two main aspects to such a description. Firstly, quantisers are specified by two operations: an encoding operation and a decoding operation. The encoding operation assigns ranges of values of the quantiser's input signal to one of a finite number of output states. The decoding operation approximately reconstructs the original signal, by assigning 'reproduction values' to each encoded state. In contrast to conventional quantisers, the SSR model's encoding is stochastic, as the output state for given input signal values is non-deterministic. However, it is possible to decode the SSR output in a similar manner to conventional quantisers, and we examine various ways to achieve this. The second aspect we consider is the performance of the decoded SSR model. Since the decoding is designed to approximate the original signal, performance is measured by the average properties of the error between this original signal, and the quantiser's output approximation. Conventionally, mean square error distortion is used to measure this average error, and we examine in detail how this measure varies with noise intensity, the decoding scheme used, and the number of threshold devices in the SSR model. As with Chapters 4 and 5, Chapter 6 focuses on

general behaviour, while Chapter 7 is devoted to discussion of the SSR model in the event of a large number of individual threshold devices.

- Chapter 8 contains original work that extends the SSR model beyond its original specification. We relax the constraint that all individual threshold devices must have identical threshold values, and allow each device to have an arbitrary threshold value. We then consider how to optimally choose the set of threshold values as the noise intensity changes. The most important result of this study is the numerical demonstration that the SSR model—where all threshold devices have the same threshold value—is optimal, for sufficiently large noise intensity.
- Motivated by recent neural-coding research, Chapter 9 further extends the SSR model, by including a constraint on the energy available to the system. The performance of a quantiser is characterised by two opposing factors: *rate*, and *distortion*. Chapter 9 also explores the SSR model, and its extension to arbitrary thresholds, from the point of view of rate-distortion theory.
- Finally, Chapter 10 concludes this thesis, by summarising its major results, and pointing out areas for possible further research.

This concludes Chapter 1. Without further ado, we now in Chapter 2 take a historical tour of prior research in the areas of SR and quantisation theory.

Chapter 2

Historical Landscape

BY definition, quantisation schemes are noisy in that some information about a measurement or variable is lost in the process of quantisation. Other systems are subject to stochastic forms of noise that interfere with the accurate recovery of a signal, or cause inaccuracies in measurements. However stochastic noise and quantisation can both be incredibly useful in natural processes or engineered systems. One way in which noisy behaviour can be useful is through a phenomenon known as stochastic resonance. The first main section of this chapter, Section 2.2, defines stochastic resonance and explores its history, with a particular emphasis on forms of stochastic resonance where the *thresholding* of random signals occurs. An important example where this occurs is in the generation of action potentials by spiking neurons. Thresholding a random signal can be described as a form of stochastic quantisation and hence Section 2.3 gives a brief history of standard quantisation theory. Such results and research have come mainly from the electronic engineering community, where quantisation needs to be understood for the very important process of analog-to-digital conversion—a fundamental requirement for the plethora of digital systems in the modern world. This chapter is entirely qualitative. Illustrative examples of stochastic resonance in threshold systems are given, but mathematical and numerical details are left for subsequent chapters.

2.1 Introduction

2.1.1 Chapter Structure

This Chapter is intended to set the context for the original work contained in this thesis. There are two sections; the first explores the history of stochastic resonance, and particularly the concept of stochastic resonance in threshold systems. The second section briefly outlines the concept of quantisation, and its role in digital systems and analog-to-digital conversion.

2.2 Stochastic Resonance

Stochastic Resonance (SR), although a term originally used in a very specific context, is now broadly applied to describe any phenomenon where the presence of internal noise or external input noise in a nonlinear system provides a better—according to some measure of performance—system response to a certain input signal than in the absence of noise. The key term here is *nonlinear*. SR cannot occur in a linear system—linear in this sense means that the output of the system is a linear transformation of the input of the system. A wide variety of performance measures have been used—we will discuss some of these later.

The term *stochastic resonance* was first used in this context in 1980 (Benzi 1980), and since then has been used—according to the ISI Web of Science database—in around 2000 publications, over a period of a quarter of a century. The frequency of publication, by year, of these papers is shown in Fig. 2.1. This figure illustrates how the use of the term expanded rapidly in the 1990s, and may possibly have plateaued in the 2000s.

SR has been the subject of many reviews, including full technical journal articles (Jung 1993, Moss *et al.* 1994, Dykman *et al.* 1995, Gammaitoni *et al.* 1998, Wiesenfeld and Jaramillo 1998, Luchinsky *et al.* 1999, Luchinsky *et al.* 1999, Anishchenko *et al.* 1999, Hänggi 2002, Harmer *et al.* 2002, Wellens *et al.* 2004, Shatokhin *et al.* 2004, Moss *et al.* 2004), editorial works (Bulsara *et al.* 1993, Astumian and Moss 1998, Petracchi *et al.* 2000, Abbott 2001), a book chapter (Wiesenfeld 1993) and magazine articles (Wiesenfeld 1993, Moss and Wiesenfeld 1995). There have been articles and letters on SR published in the prestigious journal, *Nature* (Douglass *et al.* 1993, Moss *et al.* 1993, Wiesenfeld and Moss 1995, Moss and Pei 1995, Bezrukov and Voydanoy 1995, Collins *et al.*

1995, Noest 1995, Collins *et al.* 1995b, Levin and Miller 1996, Bezrukov and Voydanoy 1997, Astumian *et al.* 1997, Dykman and McClintock 1998, Collins 1999, Bulsara 2005, Badzey and Mohanty 2005). A book has been written exclusively on SR (Ando and Graziani 2000), as well as a large section of another book (Anischenko *et al.* 2002), and sections on SR written in a popular science book (von Baeyer 2003).

SR has been widely observed throughout nature—it has been discovered in such diverse systems as climate models (Benzi *et al.* 1982), electronic circuits (Fauve and Heslot 1983), differential equations (Benzi *et al.* 1985), ring lasers (McNamara *et al.* 1988), semiconductor lasers (Iannelli *et al.* 1994), neural models (Bulsara *et al.* 1991, Longtin *et al.* 1991), physiological neural populations (Douglass *et al.* 1993), chemical reactions (Leonard and Reichl 1994), ion channels (Bezrukov and Voydanoy 1995), SQUIDs (Superconducting Quantum Interference Devices) (Hibbs *et al.* 1995), ecological models (Blarer and Doebeli 1999), cell biology (Paulsson and Ehrenberg 2000, Paulsson *et al.* 2000), financial models (Mao *et al.* 2002), and even social systems (Wallace *et al.* 1997). A slightly more exhaustive, albeit older, list is given in Luchinsky *et al.* (1998).

The first highly successful application inspired by SR involves the use of electrically generated subthreshold stimuli in biomedical prosthetics to improve human balance control and somatosensation (Priplata *et al.* 2004, Priplata *et al.* 2003, Collins *et al.* 2003, Harry *et al.* 2005). This work led to James J. Collins winning a prestigious MacArthur Fellowship (Harry *et al.* 2005).

SR is often described as a counter-intuitive phenomenon. This is largely due to its historical background, since in the first decade and a half since its 'discovery' in 1980, virtually all research into SR considered only systems driven by periodic—most often a sine-wave—input signals and broadband noise. In such systems, a natural measure of system performance is the output signal-to-noise ratio (SNR), or more precisely, often the ratio of the output Power Spectral Density (PSD) at the input frequency, to the output noise floor PSD. In linear systems driven by periodic input signals—in fact, any linear system—it is well known that the output SNR is maximised in the absence of noise. When systems are analysed in terms of SNR, it is the norm to implicitly assume that noise is a problem, usually with good reason. Hence, observations of the presence of noise in a system providing the maximum output SNR are often seen to be highly counter-intuitive.

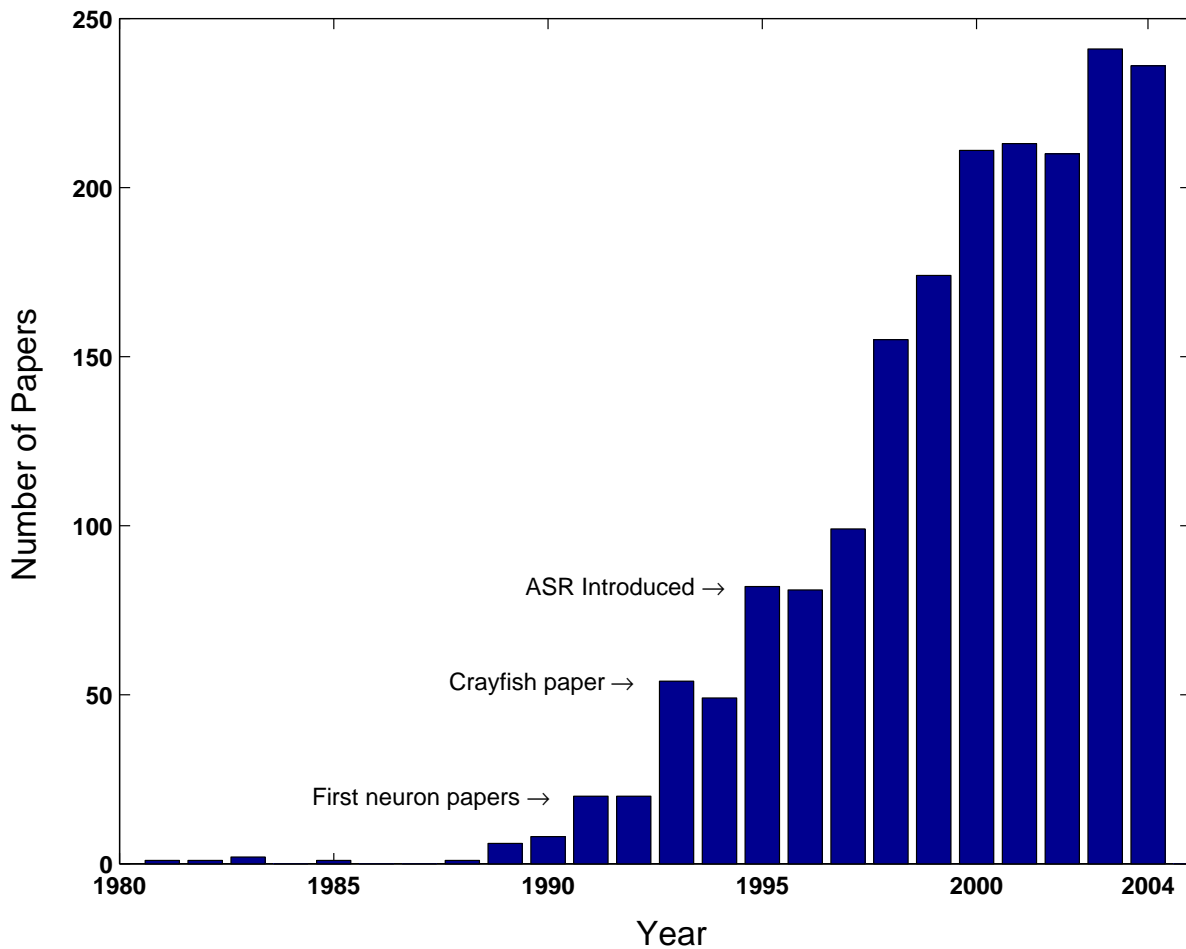


Figure 2.1. Frequency of stochastic resonance papers by year. This plot shows the frequency of stochastic resonance papers by year—between 1980 and 2004—according to the ISI web of knowledge database. There are several epochs in which large increases in the frequency of SR papers occurred. The first of these is between 1989 and 1992, when the most significant events were the first papers examining SR in neural models (Bulsara *et al.* 1991, Bulsara and Moss 1991, Longtin *et al.* 1991), and the description of SR by linear response theory—see Dykman *et al.* (1995) for a review. The second epoch is between about 1993 and 1996, when the most significant events were the observation of SR in physiological experiments on neurons (Douglass *et al.* 1993, Levin and Miller 1996), the popularisation of array enhanced SR (Linder *et al.* 1994), and of Aperiodic Stochastic Resonance (ASR) (Collins *et al.* 1995a). Around 1997, a steady increase in SR papers occurred, as investigations of SR in neurons and ASR became widespread. A rough plateau in growth appears to have formed between 2000 and 2004.

However, although not all noise can be described well as a random variable—it can, for example be constant or deterministic (even chaotic)—workers in the field have tended to focus on the stochastic case. The most common assumption is that the noise is Gaussian distributed, and white—that is, constant in power across all frequencies. When it is noted that there are many examples of systems or algorithms where randomness is of benefit, SR does not seem quite so counter-intuitive. Such examples include:

- Brownian ratchets (Doering 1995)—mechanical applications of this idea include self-winding (batteryless) wristwatches (Paradiso and Starner 2005);
- dithering in analog-to-digital conversion (Gray and Stockham 1993, Dunay *et al.* 1998, Wannamaker *et al.* 2000b);
- Parrondo's games—the random combination of losing games to produce a winning game (Harmer and Abbott 1999);
- noise induced linearisation (Yu and Lewis 1989, Dykman *et al.* 1994), noise induced stabilisation (Toral *et al.* 1999, Basak 2001), noise induced synchronisation (Neiman 1994), and noise induced order (Matsumoto and Tsuda 1985);
- the use of mixed (probabilistic) optimal strategies in game theory (von Neumann and Morgenstern 1944);
- random switching to control electromagnetic compatibility performance (Allison and Abbott 2000);
- random search optimisation techniques, including genetic algorithms (Gershenfeld 1999) and simulated annealing (Kirkpatrick *et al.* 1983);
- random noise radars—i.e. radars which transmit random noise waveforms in order to provide immunity from jamming, detection, and interference (Narayanan and Kumru 2005).
- techniques involving the use of Brownian motion for solving non-stochastic partial differential equations, such as the Dirichlet problem (Øksendal 1998);
- stochastic iterative decoding in error control coding applications (Gaudet and Rapley 2003)
- estimation of linear functionals (Plaskota 1996a, Plaskota 1996b).

2.2 Stochastic Resonance

Further discussion and other examples appear in Harmer and Abbott (2001) and Abbott (2001). Furthermore, the concept of making positive use of inherently undesirable phenomena is closely related to the idea of combining several faulty objects to obtain a working system (Challet and Johnson 2002).

The distinguishing feature of SR that sets it apart from most of the above list is that SR is usually¹ understood to occur in systems where there are well-defined input and output *signals* and the optimal output signal, according to some measure, occurs for some non-zero level and type of noise. As mentioned, SR was initially considered to be restricted to the case of periodic input signals. However, now it is used as an all encompassing term, whether or not the input signal is a periodic sine-wave, a periodic broadband signal, or aperiodic. An appropriate measure of output response depends on the task at hand, and the form of input signal. For example, for periodic signals and broadband noise, SNR is often used, but for random aperiodic signals, mutual information or correlation based measures are more appropriate.

The remainder of this section provides a more detailed history of the beginnings of SR research. It then proceeds to give the historical background to one of the simplest situations where SR can occur—the case where the nonlinearity involved is a simple threshold—and the extension of such simple situations to parallel arrays of thresholds. Finally, this section briefly explores some of the history of studies of SR in possibly the most important systems where SR is known to occur—that of neurons in the brain and nervous system.

2.2.1 A Brief History of Stochastic Resonance

The early years: 1980-1992

The term *Stochastic Resonance* was first used² in the open literature—at least, in the context of a noise-optimised system output SNR—by Benzi, Sutera and Vulpiani, in 1981, as a name for the mechanism they suggested is behind the periodic behaviour of the earth's major ice ages (Benzi *et al.* 1981, Benzi *et al.* 1982, Benzi *et al.* 1983). The term

¹An early paper on SR examined a system in which SR is said to occur that did not have any input signal (Gang *et al.* 1993)

²A search in the Inspec database for 'stochastic resonance' returns a number of published papers prior to 1981, commencing in 1973 (Frisch *et al.* 1973). However, these all use the term 'stochastic resonance' in the context of 'stochastic wave parametric resonance,' 'stochastic magnetic resonance,' or other stochastic systems, where the term 'resonance,' has nothing to do with the presence of noise.

was apparently also mentioned one year earlier by Benzi (Abbott 2001), in discussions at a workshop on climatic variations and variability in 1980 (Berger 1980). Climate records show that the period for the earth's climate switching between major ice ages and warmer periods is around 100,000 years. This also happens to be the period of the eccentricity of the earth's orbit. However current theories suggest that the eccentricity is not enough to cause such dramatic changes in climate. Benzi *et al.* (1981) —and, independently, Nicolis (1981), Nicolis (1982)—suggested that it is the combination of stochastic perturbations in the earth's climate, along with the changing eccentricity, which is behind the ice age cycle. Benzi *et al.* (1981) gave this mechanism the name *Stochastic Resonance*. In its early years, the term was defined only in the very specific context of a bistable system driven by a combination of a periodic force and random noise. Note that Benzi *et al.* (1981) considered the earth's orbital eccentricity to be a periodic driving signal, and the mentioned stochastic perturbations as the random noise. Interestingly, as pointed out in Hohn (2001), this theory for explaining the ice ages is still a subject of debate, even though SR is now well established as a *bona fide* phenomenon in a huge variety of other systems.

Further mathematical investigations of SR in the following few years included the demonstration of SR in a simple two-state model (Eckmann and Thomas 1982), the observation of SR in the Landau-Ginzberg equation (Benzi *et al.* 1985)—which is a partial differential equation—and the observation of SR in a system undergoing a Hopf bifurcation (Coullet *et al.* 1985).

Experimental observations of SR in physical systems also came quickly. In 1983 SR was reported in a Schmitt trigger electronic circuit (Fauve and Heslot 1983). This circuit is a bistable system; it was originally thought that bistability is a necessary condition for SR (McNamara and Wiesenfeld 1989). Three years later, SR was observed in a bidirectional ring laser, where the deliberate addition of noise was shown to lead to an improved output SNR (McNamara *et al.* 1988).

The ring laser paper brought about a large increase in interest in SR—approximately 50 published journal papers from 1989 to 1992—with a number of theoretical treatments being published in the next few years, for example, Gammaitoni *et al.* (1989a), Deb-nath *et al.* (1989), McNamara and Wiesenfeld (1989), and Gammaitoni *et al.* (1989b), including the important realisation that SR, in the limit of relatively small signal and relatively strong noise, can be described using linear response theory; see Dykman *et al.* (1995) for a review.

2.2 Stochastic Resonance

Incidentally, it has been pointed out that for bistable systems, the mechanism of SR has probably been known about for over 75 years. According to Dykman *et al.* (1995) and Luchinsky *et al.* (1998), the work of prolific Nobel prize winning chemist, Peter Debye, on the dielectric properties of polar molecules in a solid (Debye 1929), effectively shows SR behaviour. This fact is not entirely clear from a reading of Debye (1929), as it does not derive a formula for SNR, nor comment on the fact that any measure is optimised by a nonzero noise intensity; nevertheless, the relevant page is 105 (Dykman 2002). More recently, Kalmykov *et al.* (2004) points out that another work of Debye can be related to SR, stating that

“...it is possible to generalize the Debye-Fröhlich model of relaxation over a potential barrier ... and so to estimate the effect of anomalous relaxation on the stochastic resonance effect.” (Kalmykov et al. 2004)

However this does not imply that Debye knew of SR, only that his work has been generalised to show SR.

Expansion: 1993-1996

Stochastic resonance in excitable systems

The first important milestone in the period from 1993-1996 was the initial investigation into SR in neural and excitable systems. Previously, SR was only considered in bistable systems. An excitable system is one which has only one stable or rest state and a threshold, above which an excited state can occur, but which is not stable. The excited state eventually decays to the rest state (Gammaitoni *et al.* 1998). Neurons are a significant example of an excitable system. The first papers on the observation of SR in neural models were published by Bulsara *et al.* (1991), Bulsara and Moss (1991) and Longtin *et al.* (1991), although an earlier paper by Yu and Lewis (1989) effectively demonstrates how noise in a neuron model linearises the system response, a situation later described as SR. However research into SR in neurons and neural models only really took off in 1993, when a heavily cited *Nature* article reported the observation of SR in physiological experiments on crayfish mechanoreceptors (Douglass *et al.* 1993). In the same year a heavily cited paper by Longtin (1993) on SR in neuron models also became widely known and since then many published papers examine stochastic resonance in neurons—whether in mathematical models of neurons, or in biological experiments on the sensory neurons of animals—triggering a large expansion in research into SR. Stochastic resonance in neurons is discussed further in Section 2.2.3.

Aperiodic stochastic resonance

This important extension of SR was first addressed in Hu *et al.* (1992), which examines a bistable system driven by an aperiodic input signal consisting of a sequence of binary pulses, subject to noise. Extending SR to aperiodic signals is significant because while some important signals in nature and electronic systems are periodic, very many signals are not.

However further discussion of aperiodic SR was not undertaken until 1995, when Collins *et al.* (1995a) popularised the term Aperiodic Stochastic Resonance (ASR) in an investigation of an excitable system—a FitzHugh-Nagumo neuron model—subject to an aperiodic signal. Shortly afterwards, a letter to the journal *Nature* demonstrated the same behaviour in *arrays* of FitzHugh-Nagumo neuron models (Collins *et al.* 1995). These two papers are amongst the most highly cited of all papers on SR. The *Nature* paper is not without criticism however, and a letter to *Nature* by Noest (1995), followed by a reply from Collins *et al.* (1995b), only served to enhance the exposure of this work to the field.

As an aside, in the abstract of Collins *et al.* (1996a), the term *aperiodic* appears to be equated with the term *broadband*. However, of course, a signal can be periodic *and* broadband, for example a periodically repeated radar chirp pulse, or a simple square wave. Equating broadband with aperiodic appears to be due to the early work in SR considering only single frequency periodic signals (Gammaitoni *et al.* 1998), usually a sine wave of the form $x(t) = A \cos(\omega t)$, where $\omega = 2\pi f$ is the frequency of the sine wave. Thus, in general, an aperiodic signal can be considered to be broadband, but a broadband signal does not need to be aperiodic. The relevance to SR research is that the SNR measure used for single frequency signals is not applicable for either *broadband and periodic* signals, or aperiodic signals.

Shortly after these initial two papers by Collins *et al.*, the same authors also demonstrated ASR in three other theoretical models: a bistable-well system, an integrate-and-fire-neuronal model, and the Hodgkin-Huxley neuronal model (Collins *et al.* 1996a); as well as in physiological experiments on rat mechanoreceptors (Collins *et al.* 1996b). In these works, power-norm measures—both the correlation between the input and output signals, and the normalised power norm, or correlation coefficient—are used to characterise ASR, instead of the SNR, which is considered to be inappropriate for aperiodic signals.

2.2 Stochastic Resonance

Almost simultaneously with Collins *et al.*, an alternative approach to measuring the effects of SR on aperiodic signals was proposed by Kiss (1996). By contrast, this work employs an SNR-like measure based on cross-spectral densities to show the existence of SR in simple threshold based systems. The technique is demonstrated on a signal similar to that used by Hu *et al.* (1992), but can theoretically be applied to any broadband input. Although Collins *et al.* thought SNR measures to be inappropriate for aperiodic signals, as was soon demonstrated by Neiman *et al.* (1997), the correlation measures of Collins *et al.* can be derived from the cross-spectral density, which forms the basis of the SNR measure proposed in Kiss (1996).

Furthermore, the SNR measure used by Kiss can be rewritten in terms of a correlation-like measure. This issue is examined in detail in Section 3.4 of Chapter 3. Kiss considers that one of the advantages of his method is that it is robust to phase shifts between the input and output signal, whereas the correlation based measures of Collins *et al.* are not (Kish 2002)³. However this only applies to the first papers of Collins *et al.* since in Collins *et al.* (1996a)—unlike Collins *et al.* (1995a)—the cross-correlation between the input and output signals is calculated as a function of time delay, which is the more conventional way of calculating cross-correlation. Such a measure is indeed robust to phase shifts; the maximum cross-correlation simply occurs for a non-zero time lag. In fact, the cross-correlation is an ideal measurement of a time delay, and hence of phase shift.

Of these initial studies of ASR, it is the work by Collins *et al.* that gained the greatest exposure (Collins *et al.* 1995, Collins *et al.* 1995b), and thus many works on ASR have also used correlation based measures. However, also of great influence is a paper presented before the work of Collins *et al.* at a 1994 workshop on stochastic resonance, which shows that SR can occur for an aperiodic input signal, and the mutual information measure. This paper was subsequently published in a journal (DeWeese and Bialek 1995).

There are several reasons the work by DeWeese and Bialek (1995) has been significant for SR researchers:

- The point is made that one potentially universal characteristic of neural coding is that the

³Note that Kiss and Kish are the same person. All pre-1999 papers spell his name using ‘Kiss,’ whereas later papers use the spelling ‘Kish.’ The pronunciation is the same in both cases—‘Kish.’ This thesis will cite and refer to the spelling used in the corresponding publication.

“SNR (is) of order unity over a broad bandwidth.” (DeWeese and Bialek 1995)

Since this means that the environment in which sensory neural coding takes place appears to be very noisy, it is highly plausible that neural coding makes use of SR. This point is also made in Bialek *et al.* (1993) and DeWeese (1996).

- It is pointed out that measuring information transfer for single frequency sine waves by SNR is only really applicable in linear systems. Since SR cannot occur in linear systems, then the use of SNR only really applies to the case of small input signals so that the output exhibits a linear response.
- A proof is given of the fact that, for small signals in Gaussian noise, it is impossible for the output SNR to be greater than the input SNR. This fact has led some researchers to search for—and find (see Chapter 3)—circumstances in which the proof does not apply and that SNR gains can occur. These works did not pay attention to the previous point above, and therefore are possibly not really proving much.
- It is pointed out that sine waves do not carry information that increases with the time of observation:

“No information can be carried by the signal unless its entropy is an extensive quantity. In other words, if we choose to study a signal composed of a sine wave, the information carried by the signal will not grow linearly with the length of time we observe it, whether or not the noise is present. In addition to this, we would like to compare our results to the performance of real neurons in as natural conditions as possible, so we should use ensembles of broadband signals, not sine waves.” (DeWeese and Bialek 1995)

- It is demonstrated for the case of a subthreshold signal in a single threshold system that the information transferred through such a system can be optimised by modifying the threshold setting. With the optimal value for the threshold, the mutual information is strictly decreasing for increasing noise, and SR does not occur.

“So it seems that if you adapt your coding strategy, you discover that stochastic resonance effects disappear.. More generally, we can view the addition of noise to improve information transmission as a strategy for overcoming the incorrect setting of the threshold.” (DeWeese and Bialek 1995)

2.2 Stochastic Resonance

Thus, the conclusions drawn are that single frequency periodic signals are not relevant for information transfer, and particularly not relevant for neurons. At the time, with a growing interest into SR in neurons, such a realisation led researchers away from studying conventional single frequency SR in neurons, to more realistic broadband and aperiodic signals. For most authors, this naturally led to using measures other than SNR. Note however, that one important exception to this idea is in the encoding of sound by the cochlear nerve. The mechanism by which audio signals are encoding is essentially a biological Fourier transform; different spatial regions in the *organ of Corti*—that is, the part of the inner ear that contains sensory neurons—are sensitive to different frequencies of sound waves (Kandel *et al.* 1991).

Furthermore, the conclusion that SR in threshold systems is simply a way to overcome the incorrect threshold setting seems to have led many to think that making use of noise is a sub-optimal means of designing a system. The contrasting viewpoint is that noise is ubiquitous; since it is virtually impossible to remove all noise completely from systems, design methods should consider the effects of SR, and that various design parameters, such as a threshold value, may in some circumstances need to be set in ways that make use of the inherent noise to obtain an optimal response. We discuss exactly this situation in Chapter 8.

Subsequent to DeWeese and Bialek (1995), the next paper to use mutual information as a measure of ASR is a highly cited paper published in the journal *Nature*, which uses mutual information to experimentally show SR occurring in the cercal sensory system of a cricket (Levin and Miller 1996). As discussed in DeWeese and Bialek (1995), Levin and Miller (1996) point out that rather than using SNR,

“It is the total information encoded about a signal that is the biologically relevant quantity to consider.” (Levin and Miller 1996)

In the same year, two articles were published in the same issue of *Physical Review E*—those of Bulsara and Zador (1996) and Heneghan *et al.* (1996)—which theoretically examine the use of mutual information to measure SR for aperiodic signals. These papers paved the way for the use of information theory in SR research. The same year also saw a paper that applies other information theoretic measures—dynamical entropies and Kullback entropy—to measure SR, however unlike ASR, the system considered is driven by a periodic signal (Neiman *et al.* 1996).

Array enhanced stochastic resonance

The bulk of this thesis concerns a form of SR that occurs in arrays of parallel threshold devices. We briefly mention some of the history of what is sometimes known as Array Enhanced Stochastic Resonance (AESR). The first paper to use this term is Lindner *et al.* (1995), which shows that a chain of coupled nonlinear oscillators can provide an enhanced SR effect, when compared with a single oscillator. A very similar effect is discussed in Collins *et al.* (1995) where ASR is studied in an array of FitzHugh-Nagumo neuron models. Each neuron is considered to receive the same subthreshold aperiodic input signal, but independent noise, and the overall output is the summed response from all neurons. This result is also discussed in Moss and Pei (1995).

Other early papers showing the effect of AESR include Bezrukov and Voydanoy (1995), Pei *et al.* (1996), Gailey *et al.* (1997), Neiman *et al.* (1997), and Chialvo *et al.* (1997). An unpublished preprint gives a more detailed history (Góra 2003). The main point in these works is that the magnitude of SR effects can be enhanced by combining the outputs of more than one single SR-capable component.

Consolidation 1997-2005

Between 1997 and 2000 there was a large increase in the number of published papers either directly about SR, or listing SR as a keyword. The main development is that starting in 1997, a large number of papers examine systems showing ASR for aperiodic input signals. For example some of the papers in 1997-98 are Chialvo *et al.* (1997), Gailey *et al.* (1997), Eichwald and Walleczek (1997), Neiman *et al.* (1997), Vaudelle *et al.* (1998), Fakir (1998a), Fakir (1998b), and also Godivier and Chapeau-Blondeau (1998).

The popularisation of SR as a phenomenon that is not restricted to periodic signals has seen its original definition expanded to encompass almost any system in which input and output signals can be defined, and in which noise can have some sort of beneficial role. This period of growth however appears to have plateaued a little between 2001 and 2004. The general consensus appears to be that the most recent highly significant result in SR research is its expansion to aperiodic input signals. It could be said however, due to the number of papers investigating SR in neurons, that the most significant result of SR may be yet to come. If it can be established that SR plays an important role in the encoding and processing of information in the brain, and that it somehow provides part of the brain's superior performance to computers and artificial intelligence

2.2 Stochastic Resonance

in some areas, then using this knowledge in engineering systems may revolutionise the way we design computers, sensors and communications systems.

2.2.2 Stochastic Resonance in Threshold Systems

The prime objective of this thesis is to examine *stochastic quantisation*⁴. We define stochastic quantisation to mean:

The quantisation of a signal by thresholds that are independent and identically distributed random variables.

The starting point of this work is a particular form of SR known as Suprathreshold Stochastic Resonance (SSR). To set the context for SSR, it is necessary to briefly review SR in threshold based systems, a story that began in 1994. In a decade, the phenomenon is now known to be so widespread that a recent paper has made the point that almost all threshold systems exhibit stochastic resonance (Kosko and Mitaim 2004).

A single threshold

The first paper to consider SR in a system consisting of a simple threshold, which when crossed by an input signal gives an output pulse, was published in April 1994 under the curious title of *Stochastic resonance on a circle* (Wiesenfeld *et al.* 1994). In this paper, the authors note that all previous treatments of SR are based on the classical motion of a particle confined in a monostable or multistable potential. This appears to not be strictly true, as in 1991 and 1993 Bulsara *et al.* (1991), Bulsara and Moss (1991), Longtin *et al.* (1991), Longtin (1993), Douglass *et al.* (1993) and Moss *et al.* (1993) all analysed SR in neurons, which as excitable systems are not bistable systems in the classical sense. Additionally, although the work on SR in Schmitt trigger circuits (Fauve and Heslot 1983, Melnikov 1993) has been put into the bistable system category, a Schmitt trigger is closely related to the simple threshold crossing system—see the discussion on page 1219 of Luchinsky *et al.* (1999).

Nevertheless, the paper by Wiesenfeld *et al.* (1994) does appear to be the first paper to show the presence of SR in a system consisting of a simple threshold and the sum of

⁴Note that this is a different usage of the term ‘stochastic quantisation’ than that prevalent in areas of quantum physics, stochastic differential equations, and path-integrals such as the Parisi-Wu stochastic quantisation method (Damgaard and Huffel 1988).

a subthreshold signal and noise. Unlike previous work on SR, this paper explores SR with

“...a different class of dynamical systems based not on bistability but rather an excitable dynamics.” (Wiesenfeld et al. 1994)

The system considered consists of a weak subthreshold periodic signal, a potential barrier, and zero-mean Gaussian white noise. The output is a “spike”, which is a short duration pulse, with a large amplitude and deterministic refractory time. Note that a short refractory time for the pulse is important; the refractory time must be much less than the correlation time of the noise, or threshold crossings will occur that cannot give an output pulse. This system is described as

“...a simple dynamical process based on a single potential well, for which SR can be observed.” (Wiesenfeld et al. 1994)

It is also observed to represent the process of action potential events in sensory neurons.

Fig. 2.2 illustrates qualitatively why SR occurs in such a system. In each subfigure, the lower plot indicates the input signal’s amplitude against increasing time, with the straight line being the threshold value. The upper trace is the output signal plotted against the same time scale as the input signal. This output signal is simply a narrow pulse, or ‘spike,’ every time the input signal crosses the threshold with positive slope.

Fig. 2.2(a) shows a noiseless subthreshold periodic input signal. Such a signal never induces a threshold crossing, and the output remains constant and spike-less. Clearly nothing at all can be said about the input signal by examining the output signal, except that the input is entirely subthreshold. Fig. 2.2(b) indicates how in the absence of a signal, threshold crossings will occur randomly. The input signal in this plot is bandlimited—i.e. coloured—Gaussian noise. The statistics of this case are required for obtaining output SNRs when a signal is present, although such formulas will only hold when the input signal is small compared to the noise. In Fig. 2.2(c), the lower trace shows the sum of a periodic input signal, and bandlimited Gaussian noise. Unlike Fig. 2.2(a), for this signal some threshold crossings do occur. The probability of an output pulse is dependent on the amplitude of the noiseless input signal. This is clear visually and intuitively; a threshold crossing is more likely to occur when the input

2.2 Stochastic Resonance

signal is close to the threshold, than when it is not. Thus, compared with the absence of noise, some information about the original noiseless input signal is now available at the output, and hence removing the noise is counterproductive in this system. Note that the output pulse ‘train’ does not ‘look’ anything like the input signal, but in fact encodes the input by a form of stochastic frequency modulation. If the input signal is to be recovered, one must obtain an ensemble average to make sense of this, as is demonstrated shortly, and described in Sulcs *et al.* (2000).

Another way of understanding this effect is to realise that adding noise to a subthreshold periodic signal, and then thresholding the result, is equivalent to thresholding the noiseless signal with a time varying stochastic threshold value. This is equivalent to having a signal that is thresholded at a random amplitude, where sometimes the random amplitude is greater than the signal’s maximum amplitude. Hence, the periodic signal is thresholded at a random phase in its period, and sometimes not at all.

Returning to Wiesenfeld’s work, his paper notes that such a threshold-based system can be realised by a Josephson junction biased in its zero voltage state. It then provides a general theoretical approach for what it calls a “generic threshold-plus-reinjection dynamics” and shows how it can lead to SR. The main theoretical result given is a derivation of the SNR at the output of the generic system, using the autocorrelation of the output. This formula for SNR is shown to have a maximum for a non-zero value of noise strength, and hence shows that SR exists in such a system.

Shortly after this initial paper, several other authors also analysed SR in threshold systems subject to periodic input signals and additive noise. The first of these is an excellent paper by Jung (1994), which extends the theory of Wiesenfeld *et al.* (1994) by removing an assumption of Poisson statistics, and using the decades old work of S.O. Rice, who analysed the expected number of threshold crossings by either broadband noise alone, or the sum of a sinusoid and broadband noise (Rice 1944, Rice 1945, Rice 1948). Jung (1994) shows mathematically that the correlation coefficient between a subthreshold periodic input signal and its corresponding random output pulse train possesses a noise induced maximum.

Jung (1994) also gives a convincing argument, which has been much used in SR research, for why such simple threshold crossing systems can be used to provide simple models of neurons, particular in studies of neural networks, with large numbers of neurons. As noted by Jung (1994), such a simple two-state model of a neuron was first proposed over fifty years ago by McCulloch and Pitts (1943).

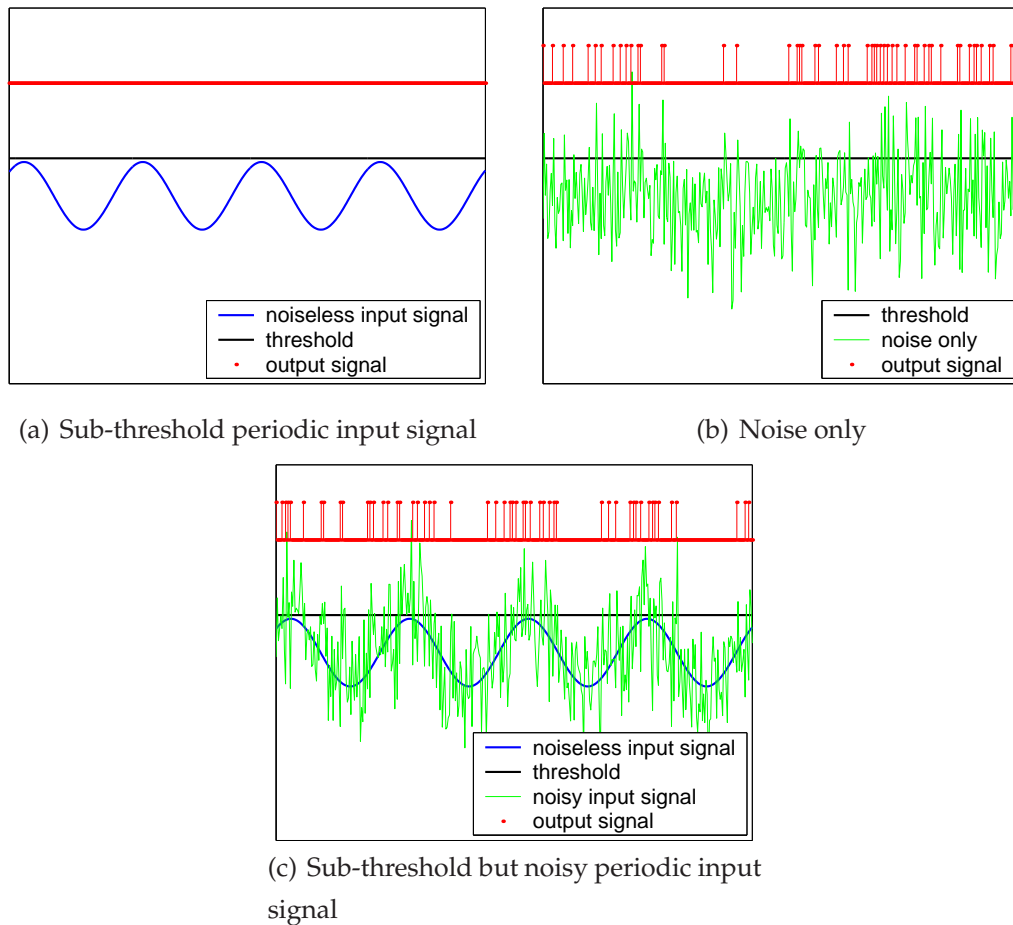


Figure 2.2. Qualitative illustration of SR occurring in a simple threshold based system. In each subfigure, the lower trace indicates the input signal's amplitude against increasing time, with the black line being the threshold value. The upper plot is the output signal plotted against the same time scale as the input signal. The output signal is a 'spike'—that is, a short duration pulse—every time the input crosses the threshold with positive slope, i.e. from subthreshold to suprathreshold but not *vice versa*. Fig. 2.2(a) shows an entirely sub-threshold periodic input signal. Such a signal will never cause an output spike. Fig. 2.2(b) shows the case of threshold crossings occurring randomly due to coloured Gaussian noise. Fig. 2.2(c) illustrates how the presence of noise enables some information about a subthreshold input signal to be present at the output. In the absence of noise, the output will always remain constant. When noise is added to the input, threshold crossings can occur, with a probability related to the amplitude of the input signal.

2.2 Stochastic Resonance

Another observation made is that SR occurs for subthreshold signals since, in the absence of noise, no threshold crossings can occur, whereas

“...in the presence of noise, there will be noise-induced threshold crossing(s), but at preferred instants of time, i.e. when the signal is larger.” (Jung 1994)

In addition, the conclusion is made that suprathreshold signals will never show stochastic resonant behaviour. We will see later in this Chapter that this is not true, once some of the assumptions used by Jung are discarded. Jung (1994) was followed up shortly afterwards in two more papers on the same topic (Jung and Mayer-Kress 1995, Jung 1995).

Simultaneously with Jung, an alternative approach was taken by Gingl *et al.* (1995b)—aspects of which also appeared in Kiss (1996) and Gingl *et al.* (1995a)—who consider a threshold based system to be a “Level Crossing Detector” (LCD), that is, a system that detects whether or not an input signal has crossed a certain voltage level. Such a system is described as “non-dynamical,” to differentiate it from the bistable systems used in ‘classical’—i.e. pre 1994—SR studies. Like Jung (1994), this work also uses the established work of Rice. Formulas first given in Rice (1944) are applied to derive a formula for the SNR in the linear response limit at the output of a LCD, a formula that is quite similar to that of Wiesenfeld *et al.* (1994). The equation obtained is verified by simulation in Gingl *et al.* (1995b).

These works on SR in threshold based systems led Gammaitoni to publish two separate papers illustrating his view that SR in threshold systems is equivalent to the effect of *dithering*—a technique used in the process of analog-to-digital conversion and image processing (Gray and Stockham 1993). This point is concisely made in the abstract of Gammaitoni (1995a), where it is stated that

“...the use of the term resonance is questionable and the notion of noise induced threshold crossings is more appropriate.” (Gammaitoni 1995a)

As commented on by Gammaitoni (1995b), such a question of nomenclature had arisen previously in the history of SR research, namely that the term ‘resonance’ had attracted some early criticism, since resonance is usually thought of in the sense of a resonant frequency. We take the point of view that such questions of nomenclature are no longer relevant, as the widely accepted definition of SR is now broad enough to cover dithering. A discussion of this topic is given in Ando (2002). The important lesson that

must be taken from Gammaitoni though is that analysis of systems that show SR can be aided by a thorough knowledge of pre-existing theories of dithering and associated research—namely the fields of lossy source coding, and quantisation. These fields are both briefly discussed in Section 2.3, and specific theoretical results stated in Chapters 4 to 9.

As with the previously mentioned papers in this Section, both Gammaitoni (1995a) and Gammaitoni (1995b) consider a periodic subthreshold input signal subject to additive white noise. In the initial paper, only a single threshold is considered at first. In subsequent sections and in the follow-up paper, a system in which more than one threshold is used to quantise a subthreshold signal is examined. However in both cases, the main point made is that SR in such systems can be considered as a special case of dithering.

Also of interest to this thesis is a formula for measuring the dithering effect proposed in Gammaitoni (1995a). This can be written as

$$D = \sqrt{\int_x (\mathbb{E}[y|x] - x)^2 dx}, \quad (2.1)$$

where x is the input signal to a threshold system, y is the output signal, and $\mathbb{E}[\cdot]$ indicates the expected value. Thus, D is the root mean square (rms) error between the input signal and the average output signal. An undiscussed assumption built into this equation is that a uniform weighting is given to each possible input signal value. In general, if the input signal is taken from a random distribution with Probability Density Function (PDF) $P(x)$, Eqn. (2.1) can be rewritten to take into account the varying probabilities of each value of x occurring as

$$D = \sqrt{\int_x (\mathbb{E}[y|x] - x)^2 P(x) dx}, \quad (2.2)$$

where the integration is over the support of $P(x)$. This formula now gives the *root mean square bias* of the system, a term that will be discussed in Chapter 6, where we will see that a performance measure of a quantiser should, in contrast to Gammaitoni (1995b) and Gammaitoni (1995a), also take into account the average conditional error variance as well as the bias.

A more recent paper by an experienced worker in the area of dithering in analog-to-digital conversion has also compared SR in threshold systems with non-subtractive dithering (Wannamaker *et al.* 2000a). This paper demonstrates that

2.2 Stochastic Resonance

“...the existence or absence of stochastic resonance in [static nonlinear systems] can be predicted from the effects of “dither averaging” upon their transfer characteristics.” (Wannamaker et al. 2000a)

However, Wannamaker *et al.* (2000a) considers only small periodic input signals, and *not* suprathreshold, or random input signals, which are the situations that will be dealt with in this thesis.

Other early papers on SR in threshold-based systems appeared in 1996 from Bulsara and Zador (1996) and Bulsara and Gammaitoni (1996), who use mutual information to measure SR for a sub-threshold aperiodic input signal. Shortly after this, Chapeau-Blondeau and co-authors (Chapeau-Blondeau 1996, Chapeau-Blondeau and Godivier 1996, Chapeau-Blondeau and Godivier 1997, Chapeau-Blondeau 1997b, Godivier and Chapeau-Blondeau 1997, Chapeau-Blondeau 1997a, Godivier *et al.* 1997), published several papers examining threshold-based SR from a wide variety of new angles.

Multiple thresholds and soft thresholds

The first paper to consider systems consisting of more than one threshold is Gammaitoni (1995b), which, as mentioned, considers SR in threshold based systems to be equivalent to dithering. The second is Gailey *et al.* (1997), which analyses an ensemble of N threshold elements, using the classical theory of nonlinear transformations of a Gaussian process to obtain cross correlations. Bezrukov and Voydanoy (1997) were the first to address a

“...class of non-dynamical and threshold-free systems that also exhibit stochastic resonance” (Bezrukov and Voydanoy 1997),

thus showing that a ‘hard’ threshold—that is, a threshold that divides its inputs into exactly two states, rather than a continuum of states—is not a necessary condition for SR to occur in non-dynamical systems.

Stochastic resonance for suprathreshold signals

The common aspect of all the above cited works is that for a single threshold, SR is only shown to occur for subthreshold input signals. It has been discussed several times that SR cannot occur in a threshold based system for an optimally placed threshold. However this is only true for certain situations. One such situation is where only the

output SNR at the frequency of a periodic signal is measured. For example, consider the case of a threshold system where the output is a pulse whenever the input signal crosses the threshold with positive slope. In the absence of noise, placing the threshold at the mean of the input signal will cause output pulses to occur once per input period. In the presence of noise, the output signal will be noisy, since spurious pulses, or jitter in the timing of the desired pulse, will occur. Thus the absence of noise is desirable in this case. This situation is illustrated in Fig. 2.3.

Another such situation, as already discussed, is given in DeWeese and Bialek (1995). When the input signal is aperiodic, and the measure used is mutual information, an optimally placed threshold also precludes SR from occurring.

However, one example that shows that SR can indeed occur for a signal that is not entirely subthreshold⁵, is when averaging is allowed for a periodic—but not single frequency—signal. Figs. 2.4 and 2.5 illustrate that the ensemble average of a noisy thresholded periodic signal—with an optimal threshold—can be better in a certain sense, in the presence of noise than without it.

Figs. 2.4(b) and 2.4(c) show a periodic but *not single frequency* signal being thresholded at its mean. Fig. 2.4(a), which shows the same signal completely subthreshold, is provided for comparison. In Fig. 2.4(b) where noise is absent, the output is identical to that of the single frequency sine wave of Fig. 2.3(a). Thus, no amount of averaging can differentiate between the two output signals. However, in Fig. 2.4(c), where noise is present, the shape of the periodic signal can be recovered upon averaging. This is illustrated in Figs. 2.5(a) and 2.5(b), which show that the shape of both subthreshold and suprathreshold input signals can be recovered in the presence of noise by ensemble averaging the output, albeit with a certain amount of distortion.

The use of ensemble averaging to increase the SNR of a noisy signal is a well known technique in areas like sonar signal processing—see also Section 6.2 in Chapter 6. When N independently noisy realisations of the same signal are averaged, the resulting SNR is known to be increased by a factor of N . This is illustrated in Figs. 2.5(c) and 2.5(d), where the input signals shown in Figs. 2.2 and 2.4 respectively have been ensemble averaged for 1000 different noise realisations. The difference between this technique, and the situation described in Figs. 2.5(a) and 2.5(b), is that in one case the

⁵In this thesis, the term ‘suprathreshold’ is not intended to mean that a signal is always entirely greater than a certain threshold. Instead, it refers to a signal that is allowed to have values that are both above and below a threshold. This includes the special case of entirely suprathreshold signals.

2.2 Stochastic Resonance

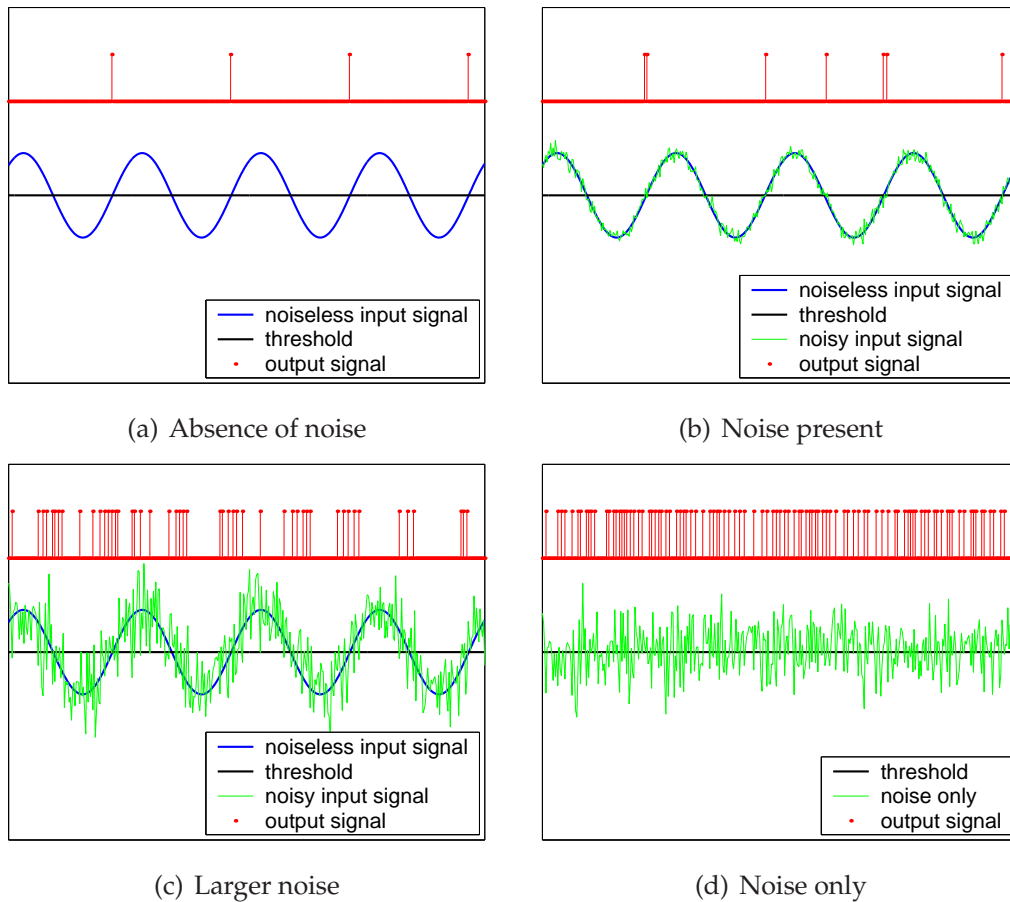


Figure 2.3. Thresholding a periodic signal at its mean. This plot illustrates that for an optimally placed threshold, at least as far as determining the frequency of a sine wave, the optimal output occurs in the absence of noise. In each subfigure, the lower trace indicates the input signal's amplitude against increasing time, with the straight line being the threshold value. The upper plot is the output signal plotted against the same time scale as the input signal. The output signal is a 'spike'—that is, a short duration pulse—every time the input crosses the threshold with positive slope—that is, from subthreshold to suprathreshold but not *vice versa*. Fig. 2.3(a) shows that in the absence of noise, there is exactly one threshold crossing per period. Fig. 2.3(b) shows that as soon as some small amount of additive noise is present, extra threshold crossings occur. This results in some noise being present in the output, as there is no longer exactly one output pulse per period. Fig. 2.3(c) shows that as the noise becomes larger, many more output pulses occur, although with a higher frequency when the input signal is close to the threshold, than when it is near its maximum or minimum. This indicates that averaging should recover the signal period. Fig. 2.3(d) shows that in the absence of a signal, the occurrence of output pulses are completely random.

signal being averaged is continuously valued, whereas in the thresholded signal case, the signal being averaged is binary. This distinction is crucial. Different techniques must be used to properly analyse each case. Furthermore, in practice, ensemble averaging a continuously valued signal to increase its SNR is simply not practical, due to the difficulty of precisely storing an analog quantity. Instead, such a signal is converted to a digital signal by quantisation, prior to averaging, giving rise to a highly analogous situation to the scenario in Figs. 2.5(a) and 2.5(b). The difference is that the quantisation is performed by a quantiser with a number of different threshold values, on each signal realisation. In the case of Figs. 2.5(a) and 2.5(b), a one bit quantisation is performed. The presence of noise then allows ensemble averaging to improve the resulting output SNR.


Thus, it seems that, at least in terms of efficiency, if the input SNR is very large, it might be worthwhile to perform the one bit quantisation N times—provided the noise is independent in each realisation—rather than try to obtain a high precision quantisation of a very noisy signal—which is stored in a multi-bit binary number—and then average the result N times. Such a technique has indeed been performed in sonar signal processing, in a method known as DIgital MUltibeam Steering (DIMUS), employed in submarine sonar arrays (Rudnick 1960)—see Section 4.2.5 of Chapter 4

Hence, SR can occur for non-subthreshold signals; one simply needs to clarify what it is that is being measured! This fact was perhaps overlooked due to an ingrained emphasis on measuring SR by the output SNR at the fundamental frequency of the input periodic signal. At the point in time where this was starting to be questioned, the emphasis switched to aperiodic input signals. As discussed below, the ensemble averaging performed here cannot be carried out for aperiodic signals in the same way.

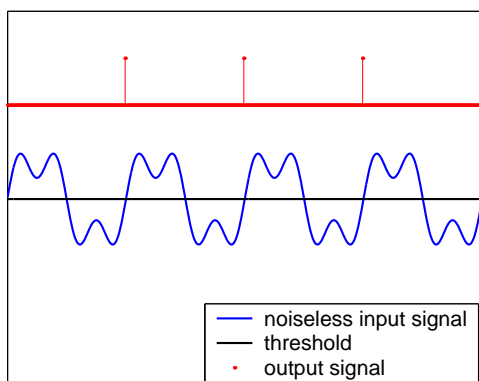
Suprathreshold stochastic resonance

The fact that ensemble averaging a thresholded periodic signal can provide a better response in the presence of noise rather than its absence leads, almost, but not quite, to the concept of Suprathreshold Stochastic Resonance (SSR) (Stocks 2000a).

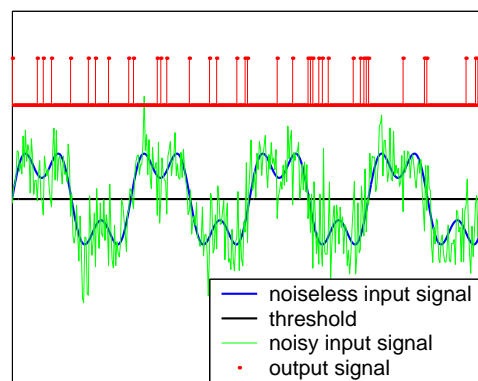
In the situations illustrated above, it is the periodicity of the input—regardless of the shape—that allows ensemble averages, taken over a period of time, to increase the output SNR. Since one knows that the signal is periodic, provided that separate ensembles of the input are all mutually in phase, each signal segment that is averaged can be collected at any point in time. This allows independent amplitudes of noise to be added



(a) Subthreshold signal



(b) Suprathreshold, noise absent



(c) Suprathreshold, noise present

Figure 2.4. Threshold SR for a periodic, but not sinusoidal signal. This figure shows an input signal that while periodic, is not a single frequency sine wave. Instead, the input is the sum of two sine waves of different frequencies. In each subfigure, the lower trace indicates the input signal's amplitude against increasing time, with the straight line being the threshold value. The upper plot is the output signal plotted against the same time scale as the input signal. The output signal is a 'spike'—that is, a short duration pulse—every time the input crosses the threshold with positive slope—that is, from subthreshold to suprathreshold but not *vice versa*. Fig. 2.4(a) shows the output signal for the case where the input signal is subthreshold, but additive noise causes threshold crossings. As with the single frequency input signal of Fig. 2.2, the noise causes output pulses to occur. The probability of a pulse is higher when the signal is closer to the threshold. Fig. 2.4(b) illustrates how, in the absence of noise, thresholding this signal at its mean will provide exactly the same output as a single frequency sine wave. Fig. 2.4(c) shows the output signal when the input is thresholded at its mean and additive noise is present. Unlike in the absence of noise, more than one pulse per period can occur. Thus, the presence of input noise creates output noise if it is only the input signal's period that is to be recovered. On the other hand, the presence of noise allows the shape of the input signal to be recovered by ensemble averaging—see Fig. 2.5.

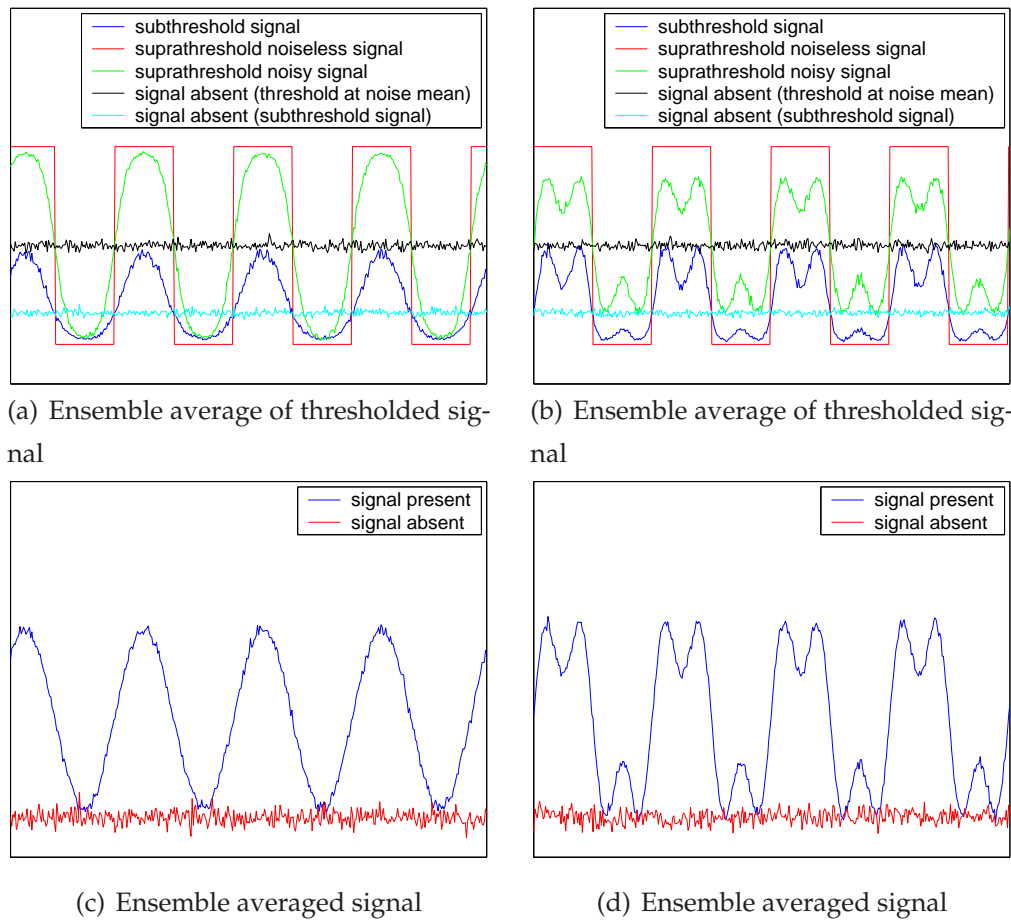


Figure 2.5. Ensemble averages of thresholded and unthresholded signals. This figure shows the result of ensemble averaging 1000 realisations of the outputs shown in Figs. 2.2, 2.3 and 2.4, as well as ensemble averages of the input signals. Fig. 2.5(a) shows ensemble averages when the input is the single frequency sine wave of Figs. 2.2 and 2.3. In the absence of noise, the ensemble average is a square wave. However, in the presence of noise, the ensemble average is clearly closer in shape to the original input signal, although also somewhat distorted. This is true for both subthreshold and suprathreshold input signals. The average of the noise is also shown for comparison. Fig. 2.5(b) shows the ensemble averages for the periodic but not single frequency input signal of Fig. 2.4. The absence of noise gives the same signal as for the single frequency input, whereas again the presence of noise provides a signal that has a shape close to the input signal's. It is well known that averaging a noisy signal N times reduces the SNR by a factor of N . This is illustrated for the unthresholded input signals in Figs. 2.5(c) and 2.5(d). Clearly the ensemble averaging of the thresholded signals in the presence of noise provides a similar effect to this, although since distortion is introduced due to the discrete nature of the output, the output SNR is greater.

2.2 Stochastic Resonance

to each amplitude of the signal, since each amplitude of signal is periodically repeated, whereas the noise is not. This situation is described in Gammaitoni *et al.* (1998). In such a situation, if the input signal is aperiodic, ensemble averaging in this fashion would not work.

By contrast, we will see in Chapter 4 that SSR refers to the *instantaneous* averaging of the outputs from an *array* of independently noisy threshold devices that all receive the same signal.

Further discussion of SSR is deferred until Chapters 4–9. However we demonstrate in these Chapters that SSR can be described as a form of non-deterministic quantisation. Hence, the next section of this chapter discusses quantisation theory. Before then however, we give a brief pointer to some of the most important results on SR in neurons and neural systems.

2.2.3 Stochastic Resonance in Neural Systems

According to the ISI Web of Science database, about 20% of SR papers also contain a reference in the title, abstract, or keywords to the words *neuron* or *neural*.

As mentioned previously, the first papers investigating SR in neuron models appeared in 1991 (Bulsara *et al.* 1991, Bulsara and Moss 1991, Longtin *et al.* 1991) with such research accelerating—for example Longtin (1993)—after the observation of SR in physiological experiments on crayfish mechanoreceptors in 1993 (Douglass *et al.* 1993). A good history of this early work on SR in neurons is given in Hohn (2001), and a recent summary of progress in the field was published in Moss *et al.* (2004).

However, there are some other published works that indicate that the positive role of noise in neurons was noticed prior to 1991. For example, Horsthemke and Lefever (1980) discuss noise-induced transitions in neural models, and in particular Yu and Lewis (1989) advocate noise as being an important element in signal modulation by neurons.

One of the most intriguing proposed applications inspired by SR, is that of enhanced cochlear implant signal encoding; various authors have advocated the exploitation of SR in this area (Morse and Roper 2000, Morse and Meyer 2000, Hohn and Burkitt 2001, Stocks *et al.* 2002, Chatterjee and Robert 2001, Rubinstein and Hong 2003, Behnam and Zeng 2003, Chatterjee and Oba 2005). A more exhaustive list is given in Hohn (2001).

2.3 Information and Quantisation Theory

Analog-to-digital Conversion (ADC) is a fundamental stage in the electronic storage and transmission of information. This process involves obtaining a sample of a signal, and its quantisation⁶ to one of a finite number of levels.

According to the Australian Macquarie Dictionary (Delbridge *et al.* 1997), the definition of the word ‘*quantise*’ is

“1. *Physics*: a. to restrict (a variable) to a discrete value rather than a set of continuous values. b. to assign (a discrete value), as a quantum, to the energy content or level of a system. 2. *Electronics*: to convert a continuous signal waveform into a waveform which can have only a finite number (usually two) of values.”
(Delbridge *et al.* 1997)

One of the aims of this thesis is to consider theoretical measures of the performance of a *stochastic* quantisation method—that is, a method that assigns discrete values in a non-deterministic fashion—and compare its performance with some of the conventional quantisation schemes that are often used in ADCs.

This section describes the basic ideas of quantisation and then briefly lists some important results in quantisation theory. However, firstly we touch briefly on information theory, and the concepts of entropy and mutual information. These ideas are required for most other chapters in this thesis.

2.3.1 Definition of Mutual Information

Consider two correlated random variables, X and Y . If X has a PDF $P(X)$ and Y has a PDF $P(Y)$, and the joint PDF is $P(X, Y)$, then Shannon’s mutual information between X and Y is defined as the *relative entropy* between the joint PDF, $P(X, Y)$, and the product of the PDFs $P(X)$ and $P(Y)$, where the relative entropy (in bits) (Cover and Thomas 1991) is

$$D(P(X, Y) || P(X)P(Y)) = \int P(\eta) \log_2 \left(\frac{P(\eta)}{P(X)P(Y)} \right) d\eta. \quad (2.3)$$

⁶The Australian Macquarie Dictionary (Delbridge *et al.* 1997) lists the spellings of ‘*quantise*’ and ‘*quantisation*’, as being equally frequent in usage in Australia, with the U.S. English spelling of ‘*quantize*’ and ‘*quantization*’. Although the ‘*iz*’ spelling is the most commonly used in the literature, the convention in this thesis is to take the first listed spelling from the Australian Macquarie Dictionary (Delbridge *et al.* 1997), that is, the ‘*is*’ spelling.

2.3 Information and Quantisation Theory

Thus, the mutual information is

$$I(X, Y) = \int_X \int_Y P(X, Y) \log_2 \left(\frac{P(X, Y)}{P(X)P(Y)} \right) dXdY. \quad (2.4)$$

It can be shown (Cover and Thomas 1991) that mutual information can be expressed as the difference between the entropy of X , $H(X)$, and the average conditional entropy, or equivocation, $H(Y|X)$, as

$$I(X, Y) = H(Y) - H(Y|X). \quad (2.5)$$

For a discrete random variable, X_d , with n possible states that occur with probabilities $p_i, i = 1, \dots, n$ where $0 \leq p_i \leq 1$ and $\sum_{i=1}^n p_i = 1$, the entropy of X_d is defined as

$$H(X_d) = - \sum_{i=1}^n p_i \log_2 p_i. \quad (2.6)$$

For a continuous random variable, X_c , with probability density function $p_{X_c}(x)$ where, $q \leq x \leq r$, the entropy of X_c is defined as

$$H(X_c) = - \int_q^r p_{X_c}(x) \log_2 p_{X_c}(x) dx. \quad (2.7)$$

Note that in this case the entropy is known as differential entropy, and is subtly different from the discrete entropy (Cover and Thomas 1991). For example, discrete entropy is always non-negative where as differential entropy can be positive, zero or negative. We also have the conditional entropy of two discrete random variables, X (n states) and Y (m states) given by

$$H(Y|X) = - \sum_{i=1}^n \sum_{j=1}^m P(X_i, Y_j) \log_2 P(Y_j|X_i), \quad (2.8)$$

where $P(X_i, Y_j)$ is the joint probability mass function of X and Y and $P(Y_j|X_i)$ is the conditional probability that Y is Y_j given X is X_i . If Y is a continuous random variable with probability density $p_Y(y)$ and X (n states) is a discrete random variable, the conditional entropy of Y given X is

$$H(Y|X) = \sum_{i=1}^n \int_y P(X_i, Y) \log_2 P(Y|X_i) dy, \quad (2.9)$$

where $p(X_i, Y)$ is the joint probability density of X_i and Y and $P(Y|X_i)$ is the conditional probability density of Y given X_i .

2.3.2 The Basics of Quantisation Theory

Quantisation of a signal or source consists of the partitioning of the signal into a discrete number of intervals, or cells. Certain rules specify which range of values of the signal get assigned to each cell. This process is known as *encoding*. If an estimate of the original signal is required to be made from this encoding, then each cell must also be assigned a reproduction value. This process is known as *decoding*.

Encoding

Although quantisation is an integral part of an ADC, it is of course not restricted to such a narrow scope; the input does not need to be an electronic signal, or a continuously valued variable. A basic example of quantisation is in the formation of a histogram for some real valued data set. For example, consider a study that measures the heights of 1000 people. The researcher may decide to divide her measurements up into ten bins, of which eight are equally spaced with length 5 cm starting from 160 cm. The other two bins—overflow bins—are for measured heights of less than 160 cm, and more than 200 cm. To obtain a histogram, the number of measurements that fall into each bin are then counted and plotted against the index of the bin.

What information can then be gleaned from the histogram about the heights of the people in the study? The researcher will look at statistical measures such as percentage frequency of each bin, mean, mode, median and variance. If all the measurements fall in just one or two bins, the researcher will realise that the bin spacing, and the difference between the maximum and minimum bins is too wide to obtain any detail about the distribution of heights, or that more bins are required. Another similar problem will occur if most of the measured heights are in one or both of the overflow bins.

Decoding

Such questions of bin number, size and placement are precisely those faced by the designer of a quantiser. The difficulty of finding a good design may also be compounded by the fact that the signal being quantised may not be as stationary as the heights of people. Furthermore, not only does the quantisation of a signal usually require the *encoding* into bins, it also requires the *decoding* operation to be specified. In the histogram binning analogy, the need for decoding is probably not all that interesting to

2.3 Information and Quantisation Theory

the anatomy researcher measuring people's heights, but could be understood as follows. Select randomly one of the 1000 people whose heights were measured and ask that person to specify which bin his height falls in, without specifying his exact height. What can the researcher then say about the height? If the bin is the one from 180–185 cm, only that the person is no shorter than 180 cm and no taller than 185 cm. If asked to guess the height, the researcher would probably guess 182.5 cm, knowing that the maximum error in the guess would be 2.5 cm.

It is this question of assigning an estimated height to a bin that is exactly the problem of decoding in quantisation. In the case of a signal being quantised, the value assigned as the decoding for each quantisation bin is sometimes known as the *reproduction point*.

Another example of quantisation is the representation of real numbers in a computer's architecture. Examples of such quantisation schemes include the IEEE floating-point and fixed point standards (Widrow *et al.* 1996).

Measures of a quantiser's performance

One important measure of a quantiser is its *rate*. In this context, 'rate' does not necessarily refer to a quantity that is defined in terms of 'per unit time.' For example mutual information, entropy, or the number of output bits have all been used as the definition of 'rate.' The idea is that rate provides a measure of how many bits—that is, how many binary symbols—are required to represent information. Thus, the rate of a quantiser is usually meant as the (average) number of bits per sample that the quantiser output consists of, or contains about the input. If mutual information is used as the definition of rate, then in general, the rate will depend on the statistics of the input signal as well as the encoding process. However, for a deterministic encoding, the rate is simply the average entropy of the output encoding, which will be the same as the mutual information, and is given by

$$I(x, y) = H(y) = - \sum_{n=0}^N P_y(n) \log_2 P_y(n), \quad (2.10)$$

where $P_y(n)$ is the probability of output state n occurring. The maximum rate occurs when all output states are equally likely and is given by $I(x, y)_{\max} = \log_2 (N + 1)$, that is, the number of bits at the output of the quantiser.

Information theory also tells us that the quantisation of a signal will always cause some error in a reproduction of the original signal. This error is known as the distortion, and

is most commonly measured by the mean square error between the original signal, and the reproduced signal (Gray and Neuhoff 1998). Thus, if the encoding is decoded to a signal, z , the error is given by

$$\epsilon = x - z, \quad (2.11)$$

and the mean square distortion is

$$D_{\text{ms}} = \text{E}[(x - z)^2]. \quad (2.12)$$

A commonly used measure of a quantiser's performance is its Signal-to-Quantisation-Noise Ratio (SQNR), which is the ratio of the input signal's power to the mean square distortion power. If the input signal has power σ_x^2 , then this can be expressed as

$$\text{SQNR} = 10 \log_{10} \left(\frac{\sigma_x^2}{D_{\text{ms}}} \right). \quad (2.13)$$

2.3.3 Quantisation Literature Review

Comprehensive reviews of the history of quantisation theory, and the closely related topic of lossy source coding can be found in Gray and Neuhoff (1998) and Berger and Gibson (1998). These papers provide a detailed history of early practically motivated quantisation work, such as pulse code modulation in the 1950s, as well as early theoretical work on lossy source coding, including Shannon's initial formulation of the problem of minimising the rate required to achieve a given distortion. As well as setting the historical context for these fields, Gray and Neuhoff (1998) and Berger and Gibson (1998) also work through the state of the art, and future directions for research. Another excellent reference is the textbook Gersho and Gray (1992).

We will now briefly discuss some of the aspects of quantisation theory most pertinent to the work in this thesis.

Optimal quantisation

We have seen that the design of a quantiser reduces to choosing how to partition an input into bins, and the selection of reproduction points for each bin. Selecting the bins requires choosing the number of bins and the values of the thresholds that define which input values go in to each bin. Reproduction points should be selected that provide a good estimate of all the values represented by the corresponding bin.

So how should a quantiser be optimally designed? Suppose that only N -bit quantisations can be performed. Then $N - 1$ threshold values and N reproduction points are required.

Given a measure of distortion, the reproduction points can be chosen as the values that, on average, minimise that distortion measure, for a given input signal. For continuously valued random source distributions—other than the uniform distribution—analytic expressions for the optimal thresholds are rare, at least for the mean square error distortion. However, standard numerical algorithms such as the Lloyd-Max method (Max 1960, Lloyd 1982, Gersho and Gray 1992) can be applied to find the optimal partition and reproduction points for a given specified source distribution, or a set of training data.

However, these algorithms appear to never have been extended to consider the situation where the threshold values that form the partitions are independently noisy, or random variables. Such a situation is discussed in this thesis.

Dithering and stochastic quantisation

Although we have defined ‘stochastic quantisation’ to mean quantisation by thresholds that are independent random variables, the concept of ‘dithering’ could also be understood as stochastic quantisation, in that it also results in an output signal which is not deterministic.

For technical references on dithering, see for example Chou and Gray (1991), Gray and Stockham (1993), Kikkert (1995), Zamir and Feder (1995) Wannamaker (1997), Carbone and Petri (1998), Dunay *et al.* (1998), Carbone and Petri (2000), and Wannamaker *et al.* (2000b).

Work commenting on the relationship between SR and dithering includes Gammaitoni (1995b) Gammaitoni (1995b), Wannamaker *et al.* (2000a), Ando and Graziani (2000) and Ando (2002). Note that Vaudelle *et al.* (1998) discusses SR in images—in a way that appears to be identical to dithering in images—without citing dithering.

Recall from earlier discussion in this Chapter how a subthreshold signal subject to a threshold is non-detectable at the output. Adding noise to the input signal to allow threshold crossings is effectively the same as dithering in a single bit quantiser, when the signal amplitude is smaller than the quantiser’s bin size. This illustrates how SR can be related to the use of dithering to increase the dynamic range of an ADC. See,

for example, Wannamaker *et al.* (2000a) and Lim and Saloma (2001) for discussions on this.

However, there are several crucial differences between dithering as it is usually understood, and the concept we have called ‘stochastic quantisation.’ These are that dither signals:

- are usually considered to have a dynamic range smaller than the width of a quantiser’s bin size, and therefore are small compared to the signal’s dynamic range;
- are usually applied to quantisers with widely spaced thresholds;
- and usually have PDFs with finite support, like the uniform distribution, rather than infinite support, like the Gaussian distribution.

By contrast, we will consider stochastic quantisation in the following scenarios:

- large dither—or noise—amplitudes, compared to the signal’s amplitude;
- the case where all thresholds in a quantiser have identical values, but becomes independent random variables due to the addition of noise;
- and noise signals with PDFs with infinite support.

Estimation theory

We will also in this thesis touch on areas of *point estimation theory*. An excellent technical reference encompassing this field is Lehmann and Casella (1998). The main estimation topic we will look at is that of minimising mean square error distortion between the input and output signals of a nonlinear system. Such a goal also appears in quantisation theory, however we will find it useful to use ideas from estimation theory, such as Fisher information, and the Cramer-Rao bound, which are not generally used in conventional quantisation theory.

There are also a number of papers in the SR literature which tackle SR from this point of view (Greenwood *et al.* 1999, Greenwood *et al.* 2000, Chapeau-Blondeau and Rojas-Varela 2001, Greenwood *et al.* 2003, Chapeau-Blondeau and Rousseau 2004).

2.4 Chapter Summary

This chapter reviews the two main areas in the scientific literature that are relevant to the new work presented in this thesis. Section 2.2 uses a historical perspective to review stochastic resonance, and the main sub-areas of stochastic resonance important to this thesis. In particular we define stochastic resonance, and discuss its occurrence in simple threshold based systems, and neurons.

Section 2.3 then introduces the concept of quantisation, briefly indicates some of the most commonly used measures of a quantiser's performance, and points out references in the literature that discuss quantisation in full technical detail.

2.4.1 Original Contributions for Chapter 2

This chapter included the following original contributions:

- A discussion of the evolution of the term 'stochastic resonance,' and—in the first sentence of Section 2.2—a statement defining stochastic resonance as it is currently widely understood.
- An argument explaining why SR need not necessarily be thought of as 'counter-intuitive,' based on the fact that there are many systems known where randomness can be useful.
- A historical review and elucidation of the major epochs in the history of stochastic resonance research.
- A qualitative demonstration that SR can occur in a single threshold device, where the threshold is set to the signal mean. Stochastic resonance will not occur in the conventional SNR measure in this situation, but only in a measure of distortion, after ensemble averaging.

2.4.2 Further Work

Possible future work and open questions arising from this chapter might include:

- Further investigation of the simple example illustrating the way that stochastic resonance can be understood to occur in a threshold system for a signal with a mean value equal to the threshold.

This concludes Chapter 2, which sets the historical context for the work in this thesis. The following chapter looks in more detail at some of aspects of SR that seem confusing to the uninitiated. It also provides the first analytical and numerical results of this thesis.

Chapter 3

How Should Stochastic Resonance be Measured?

INITIAL studies of stochastic resonance focused on systems driven by a periodic signal, and hence used a signal-to-noise ratio based measure for comparison between the input and output of the system. It has been pointed out that for the more general case of aperiodic signals other performance measures are necessary, such as cross-correlation or information theoretical tools. This Chapter analyses the application of such measures to stochastic resonance, and discusses situations where signal-to-noise ratio based measures have been used. The relevance (or lack thereof) of these measures to engineering problems is discussed.

3.1 Introduction

There can sometimes be a certain amount of skepticism about claims made in the Stochastic Resonance (SR) literature from researchers who have never seriously performed research in this field. Partly this is due to the emphasis often put on the counter-intuitive notion of noise being “useful,” rather than “a nuisance.”⁷ As discussed in Chapter 2, the skeptics may find it easier to come to grips with SR when it is described as a way to make use of randomness to overcome certain limitations of a system. The second source of skepticism stems from the extensive use of techniques borrowed from other fields, which are then used slightly out of their original context. A third source is due to confusion about the definition of SR and how it has changed. Originally it was said to only occur for periodic input signals in bistable systems subject to broadband noise. At present, some authors still use only this definition, whereas others are familiar with SR in a myriad of other forms.

The aim of this Chapter is to investigate several sources of such confusion, in particular, the use of signal-to-noise ratio (SNR) measures to quantify SR, the debate about SNR gains due to SR, and the relationship between SNRs and information theory.

Section 3.3 contains original work that uses the *data processing inequality* to investigate certain results on channel capacity in a system capable of SR (McDonnell *et al.* 2003b, McDonnell *et al.* 2003c).

Section 3.4 contains original work exploring the theory behind a frequency dependent SNR measure, and its application to the measurement of SR (McDonnell *et al.* 2004a).

3.1.1 How Do I Measure Thee? Let Me Count The Ways

The reason for the name of this Section⁸ is to draw attention to the fact that SR has been measured in many different ways. Examples include SNR (Benzi *et al.* 1981), spectral power amplification (Rozenfeld and Schimansky-Geier 2000, Imkeller and Pavlyukevich 2001, Drozhdin 2001), correlation coefficient (Collins *et al.* 1995a), mutual information (Levin and Miller 1996), Kullback entropy (Neiman *et al.* 1996), channel capacity (Chapeau-Blondeau 1997b), Fisher information (Greenwood *et al.* 1999), ϕ -divergences (Inchiosa *et al.* 2000, Robinson *et al.* 2001), and mean square distortion (McDonnell *et al.* 2002a). It has also been analysed in terms of residence time

⁷The phrase “noise is often thought of as a nuisance” can be found in numerous papers on SR.

⁸Apologies to Elizabeth Barrett Browning (1998)

distributions—see Gammaitoni *et al.* (1998) for a review—and Receiver Operating Characteristic (ROC) curves (Robinson *et al.* 1998, Galdi *et al.* 1998, Zozor and Amblard 2002), which are based on probabilities of detecting a signal to be present, or falsely detecting a non-existing signal (Urick 1967).

The key point is that the measure appropriate to a given task should be used. Unfortunately, due to historical reasons, many authors tend to use the original measure used, SNR, in contexts where it is effectively meaningless. This Section gives a brief history of the use of SNRs in SR research, and a discussion of some of the criticisms of its use.

It was first thought that SR occurs only in bistable dynamical systems, generally driven by a periodic input signal, $A \sin(\omega_0 t + \phi)$, and broadband noise. Since the input to such systems is a simple sinusoid, the SNR at the input is a natural measure to use, with the following definition most common,

$$\text{SNR} = \frac{P(\omega_0)}{S_N(\omega_0)}. \quad (3.1)$$

In Eqn. (3.1), $P(\omega_0)$ is the input signal power and $S_N(\omega_0)$ is the Power Spectral Density (PSD) of the noise at frequency ω_0 . Stochastic resonance occurs when the ratio of the output power at frequency ω_0 to the background noise PSD at ω_0 , is maximised by a nonzero value of noise intensity.

It is well known in electronic engineering that nonlinear devices cause output frequency distortion—that is, for a single frequency input, the output will consist of various harmonics of the input (Cogdell 1996). Hence, basic circuit design requires the use of filters that remove unwanted output frequencies. For example, this *harmonic distortion* in audio amplifiers is very undesirable. On the other hand, high frequency oscillators make use of this effect, by starting with a very stable low frequency oscillator, and sending the generated signal through a chain of frequency multipliers. The final frequency is harmonically related to the low frequency source.

For more than one input frequency, the output of the nonlinear device will contain the input frequencies, as well as integer multiples of the sum and difference between all frequencies (Cogdell 1996). This effect of creating new frequencies is known as *intermodulation distortion*. In the field of optics this phenomenon can be used to generate lower frequency signals—for example, T-rays (i.e. terahertz radiation)—from different optical frequencies by a method known as *optical rectification* (Mickan *et al.* 2000, Mickan and Zhang 2003).

3.1 Introduction

A study of such higher harmonics generated by a nonlinear system exhibiting SR has been published (Bartussek *et al.* 1994), and the phenomenon is also discussed in subsequent works (Bulsara and Inghiosa 1996, Inghiosa and Bulsara 1998). However, much research into SR has only been interested in the output frequency component that corresponds to the fundamental frequency of the periodic input signal, in which case the output SNR is given by Eqn. (3.1) and ignores all other output harmonics.

More recently, attempts have been made to overcome this, by defining the output SNR as a function of all input frequencies (Kiss 1996, Gingl *et al.* 2001, Mingesz *et al.* 2005). However, while such formulations may have some uses—see Section 3.4—there has been much discussion regarding the inadequacies of SNR as an appropriate measure for many signal processing tasks (DeWeese and Bialek 1995, Galdi *et al.* 1998).

One of the most important objections can be illustrated as follows. Consider a periodic, but broadband input signal, such as a regularly repeated radar chirp signal. The use of SNR as the ratio of the output power of the fundamental frequency to the background noise PSD is meaningless for signal recovery here, unless only the fundamental period is of interest. This output SNR measure only provides information about the period of the signal—the output SNR at that frequency—and nothing about the shape of the chirp in the time domain.

This inadequacy was recognised when researchers first turned their attention to ASR, who ushered in the widespread use of cross-correlation and information theoretic measures, which can, in some sense, describe how well the shape of the output signal is related to the input signal. Inghiosa and Bulsara (1995) give an excellent description of the issue,

“...a nonlinear signal processor may output a signal that has infinite SNR but is useless because it has no correlation with the input signal. Such a system would be one which simply generates a sine wave at the signal frequency, totally ignoring its input.” (Inghiosa and Bulsara 1995)

While many SR researchers realised that studying aperiodic input signals substantially increases the relevance of SR to applications such as studies of neural coding, some did not get past the need to move on to measures other than SNR in such circumstances. This has led to a somewhat strange debate about whether or not ‘SNR gains’ can be made to happen in a nonlinear system by the addition of noise.

3.1.2 Chapter Structure

This Chapter is divided into three—mostly self-contained—main sections. Firstly, Section 3.2 discusses the SNR gain debate, and the main questions on this issue in the literature. Next, Section 3.3 considers an apparent contradiction in the SR literature, regarding whether the addition of noise to a system can increase the information available at the output. The conclusion drawn is that, as we might expect, information cannot be increased by any form of signal processing, although the right amount of noise might minimise the amount of information lost. Thirdly, Section 3.4 examines the theory behind a frequency dependent SNR formula introduced to the SR literature in Kiss (1996). Some results relevant to this chapter are given in more detail in Appendix A.

3.2 The SNR Gain Debate

In the last decade, a number of researchers have reported results claiming that it is possible to obtain an SNR gain in some nonlinear systems by the addition of noise (Kiss 1996, Loerincz *et al.* 1996, Vilar and Rubí 1996, Chapeau-Blondeau 1997a, Chapeau-Blondeau and Godivier 1997, Chapeau-Blondeau 1999, Gingl *et al.* 2000, Liu *et al.* 2001, Gingl *et al.* 2001, Makra *et al.* 2002, Casado-Pascual *et al.* 2003). There has been some criticism of these works, for example the comment on Liu *et al.* (2001) given in Khorvanov and McClintock (2003).

When only the frequency of an input periodic signal is of interest, the SNR gain in a nonlinear system is defined as

$$G = \text{SNR}_o / \text{SNR}_i = \frac{P_o(\omega_0)}{P_i(\omega_0)} \frac{S_{N,i}(\omega_0)}{S_{N,o}(\omega_0)}. \quad (3.2)$$

In the SR community, such results have been seen as fairly controversial, for two reasons, as discussed in the remainder of this Section.

3.2.1 Can SNR Gains Occur At All?

Firstly, initially it seemed that SNR gains contradicted proofs that SNR gains cannot occur. For example DeWeese and Bialek (1995)—see also Dykman *et al.* (1995)—show for stationary Gaussian noise and a signal that is small compared to the noise, that

3.2 The SNR Gain Debate

for nonlinear systems the gain, G , must be less than or equal to unity, and that hence no SNR gain can be induced by utilising SR. This proof is based on the use of linear response theory, where, since the signal is small compared to the noise, both the signal and noise are transferred linearly to the output, and as in a linear system, no SNR gain is possible. Much attention has been given to this fact, since most of the earlier studies on SR were kept to cases where the linear response limit applies, to ensure that the output is not subject to the above mentioned harmonic distortion (Gingl *et al.* 2000).

Once this fact was established, researchers still hoping to be able to find systems in which SNR gains due to noise could occur turned their attention to situations not covered by the proof—that is, the case of a signal that is not small compared to the noise, or broadband signals or non-Gaussian noise.

For example, Kiss (1996) considers a broadband input signal, and being broadband, the conventional SNR definition cannot be used. Instead, a new frequency dependent SNR measure is derived, a measure with which an SNR gain is shown to occur. Further examples are Chapeau-Blondeau and Godivier (1997) and Chapeau-Blondeau (1997a), which use the conventional SNR definition, but the large signal regime to show the existence of SNR gains. Furthermore, Chapeau-Blondeau (1999) also considers the case of non-Gaussian noise.

However, the interpretation of some of this work can be a little fuzzy. For example, it is stated in Makra *et al.* (2002) that

“...it is a legitimate question whether systems showing SR can possibly function as filters, making the signal passing through them less noisy. A positive answer to this question might have great significance: it might establish the theoretical basis for technical applications and might lead us to a better understanding of several (especially biological) systems.” (Makra et al. 2002)

The title of Makra *et al.* (2002) is “Signal-to-noise ratio gain by stochastic resonance in non-dynamical and dynamical bistable systems” and implies that the SNR gain is *due to* SR, and hence implies that the gain is due to the addition of noise to an already noisy signal, whilst keeping the input SNR constant.

However, what is actually performed, is that the SNR gain is plotted against input noise, where modifying the input noise simultaneously changes the input SNR—that is, an increase in input noise causes a decrease in input SNR. Since it is known that SR occurs in the systems analysed, the output SNR initially increases with increasing

noise, reaches a maximum, and then decreases again. Hence with increasing input noise, the output SNR has a single-peaked curve, and the input SNR has a linearly increasing curve. Therefore a general statement is that the SNR gain also has a single-peaked curve, when plotted against input noise, and shows a maximum with nonzero noise. What is described in Makra *et al.* (2002) is effectively exactly the same thing as SR—that is, a plot of the output SNR against input noise.

If an SNR gain was to be “caused by SR”, then it would need to be *caused by* the addition of noise to an *already noisy* signal, where the input SNR gain is unaffected by the addition of more noise. Hence, the statements implying that the gain is due to SR are a little misleading.

Instead, the SNR gain does not occur due to the addition of noise to an already noisy signal. An increase in input noise—which causes a decrease in input SNR—can increase the SNR gain, due to the same mechanism that causes the SR itself. This means that the gain is due to the characteristics of the system itself rather than the addition of noise.

However, the main point emphasised is that the SNR gain can be greater than one, which does not occur for the linear response regime, and is a valid point. Thus, the answer to the title of this subsection is “yes, SNR gains can occur.” The more important question is whether such gains are meaningful.

3.2.2 Are SNR Gains Meaningful?

By looking outside the conditions of the proof that SNR gains cannot occur in the linear response limit, SNR gains can be found. However, the second reason that an emphasis on SNR gains due to SR are seen to be controversial is that the definitions of SNR used in cases where SNR gains occur are not always particularly meaningful. Taking the approach of looking outside the parameters of the proof assumes that SNR is still a useful measure outside these parameters. A strong argument against this assumption, and for the use of information theory, rather than SNRs, is given in DeWeese and Bialek (1995), as discussed in Section 2.2.1 of Chapter 2.

Intelligent discussions of this point are given in Inchiosa and Bulsara (1995), Galdi *et al.* (1998), Robinson *et al.* (1998), Petracchi (2000), Hänggi *et al.* (2000), and Robinson *et al.* (2001). For example, Hänggi *et al.* (2000) give a general investigation of SNR gains due to noise, and are highly critical of the use of SNR in such systems, and indicate more

3.3 The Data Processing Inequality and SR

appropriate measures to use, at least for signal detection or estimation problems. Thus, the answer to the title of this subsection is “probably not, for most tasks.”

Recall the quote given in Section 3.1.1; Inchiosa and Bulsara (1995) recognise what is well known to electronic engineers—that an SNR gain is not in itself a remarkable thing, and that SNR gains are routinely obtained by filtering—for example, the band-pass filter. The reason that more is made of such phenomena in the SR literature, is that the reported SNR gains are said to be due to the *addition of noise* to an already noisy signal, rather than a deliberately designed filter. Another paper by the same authors also discusses this topic (Inchiosa and Bulsara 1996). As discussed in Subsection 3.2.1, the view that SNR gains occur *due to* SR, can be misleading.

Despite the objections raised, work reporting SNR gains continue to be churned out without regard to these criticisms. An associated problem is that of relating SNRs to information theory. For example, it has sometimes been stated that an SNR gain in a periodic system is analogous to an increase in information. The following Section investigates such a claim.

3.3 The Data Processing Inequality and SR

The search for SNR gains due to SR in the case of periodic input signals naturally led some authors to look for an analogy to compare input performance to output performance for Aperiodic Stochastic Resonance (ASR). As mentioned, Kiss (1996) defined a frequency dependent SNR measure based on cross-spectral densities, which he considered to be valid for such aperiodic input signals. This method is discussed further in Section 3.4.

An alternative approach for measuring ASR is mutual information. A special case of mutual information is known as *channel capacity* (Cover and Thomas 1991). Channel capacity is simply defined as the maximum possible mutual information through a ‘channel’ or system. It is usually defined in terms of the input probability distribution that provides the maximum mutual information, subject to certain constraints. For example, the input signal may be restricted to two states—that is, a binary signal—or to be a continuously valued random variable, but with a specified power.

The most widely known formula describing channel capacity is the Shannon-Hartley formula, which gives the channel capacity for the transmission of a power-limited and

band-limited signal through an additive, signal-independent, Gaussian white noise channel. As mentioned in Berger and Gibson (1998), this formula is often misused in situations where it does not apply, including one could argue, in the SR literature, for example, Kish *et al.* (2001).

Channel capacity as a measure of SR is discussed in Chapeau-Blondeau (1997b), Godivier and Chapeau-Blondeau (1998), Goychuk and Hänggi (1999), Kish *et al.* (2001), Goychuk (2001) and Bowen and Mancini (2004), all of which show that the right level of noise can provide the maximum channel capacity. However, in general this only means that the right level of input noise optimises the channel, that is, SR occurs.

Of more interest to us here is a way of comparing the input signal to the output signal in a way analogous to SNR gains for periodic signals. For example, in Chapeau-Blondeau (1999) it is considered that comparing the channel capacity at the input and the output of a system for an aperiodic input signal is analogous to a comparison of the input and output SNRs for periodic input signals. This Section investigates the use of channel capacity in simple threshold based systems where SR can occur, and shows by use of a well known theorem of information theory, that such an analogy is a false one.

3.3.1 The Data Processing Inequality

The Data Processing Inequality (DPI) of information theory proves that no more information can be obtained out of a set of data than is there to begin with. It states that given random variables X, Y , and Z that form a Markov chain in the order $X \rightarrow Y \rightarrow Z$, then the mutual information between X and Y is greater than or equal to the mutual information between X and Z (Cover and Thomas 1991). That is

$$I(X, Y) \geq I(X, Z). \quad (3.3)$$

In practice, this means that no signal processing on Y can increase the information that Y contains about X .

It should be noted that the terminology *Markov chain* used in connection with the DPI is somewhat more inclusive than that prevalent in applied probability. DPI usage requires only the basic Markov property that Z and X are conditionally independent given Y . By contrast the usage in applied probability requires also that X, Y , and Z range over the same set of values and that the distribution of Z given Y be the same as that of Y given X .

3.3 The Data Processing Inequality and SR

Generic nonlinear noisy system

To illustrate the arguments in this Section, consider a generic system where a signal, $s(t)$, is subject to independent additive random noise, $n(t)$, to form another random signal, $x(t) = s(t) + n(t)$. The signal $x(t)$ is then subjected to a nonlinear transformation, $T[\cdot]$ to give a final random signal, $y(t) = T[x(t)]$. A block diagram of such a system is shown in Fig. 3.1.

As noted in the previous Section, many papers have demonstrated that SNR gains can occur due to SR. Such a system as that of Fig. 3.1 describes many of those reported to show SNR gains for both periodic or aperiodic input signals.




Figure 3.1. Generic nonlinear system. A schematic diagram of a generic noisy nonlinear system. The input signal, $s(t)$, is subject to additive noise, $n(t)$. The sum of the signal and noise, $x(t)$ is subjected to the nonlinear transfer function, $T[\cdot]$, to give $y(t) = T[x(t)]$.

Channel capacity, information, and SNR gains

Ignoring for now the question over whether the SNR measure has much relevance in such cases, the observation of SNR gains can appear on the surface to contradict the DPI. The reason for this is that one could be led to believe that information can always be related to SNR by the Shannon-Hartley channel capacity formula⁹,

$$C = 0.5 \log_2 (1 + \text{SNR}) \quad \text{bits per sample.} \quad (3.4)$$

Clearly, when this formula applies, an increasing SNR leads to an increase in the maximum possible mutual information through a channel. Suppose that Eqn. (3.4) does apply in Fig. 3.1 and that the SNR of $s(t)$ in $x(t)$ is SNR_1 . Then the maximum mutual information between $s(t)$ and $x(t)$ is $I(s, x) = 0.5 \log_2 (1 + \text{SNR}_1)$ bits per sample.

Suppose also that the operation $T[\cdot]$ filters $x(t)$ to obtain $y(t)$, such that the filtering provides an output SNR for $s(t)$ in $y(t)$ of SNR_2 . If the filtering provides an SNR gain

⁹Some references instead refer to this formula as the “Hartley-Shannon formula”, the “Shannon-Hartley-Tuller law”, or simply as “Shannon’s channel capacity formula.” It is also variously known as a ‘theory,’ ‘law,’ ‘equation,’ ‘limit,’ or ‘formula.’

then $\text{SNR}_2 > \text{SNR}_1$. Now, consider the overall system that has input, $s(t)$, and output, $y(t)$. If Eqn. (3.4) applies for this whole system, then the mutual information between $s(t)$ and $y(t)$ is $I(s, y) = 0.5 \log_2(1 + \text{SNR}_2) > I(s, x)$. This is clearly a violation of the DPI. Thus, an SNR gain either cannot occur in a system in which Eqn. (3.4) applies, or Eqn. (3.4) does not apply. If one believes that Eqn. (3.4) always applies, then skepticism about the occurrence of SNR gains can be forgiven.

However, it is instead the validity of Eqn. (3.4) that needs consideration. As mentioned, this formula is often widely misused (Berger and Gibson 1998), as it only applies for additive Gaussian white noise channels, where the signal is independent of the noise. Of particular relevance here is the *additive noise* part. No SNR gain such as that from SNR_1 to SNR_2 can be achieved in such an additive noise channel. Hence, Eqn. (3.4) can never apply to the situation mentioned above between signals $s(t)$ and $y(t)$, since even if it applies between $s(t)$ and $x(t)$, the SNR gain required in the filtering operation rules it invalid. Hence, there is no reason to be skeptical about SNR gains, except for cases where the Shannon-Hartley channel capacity formula is actually valid.

However, such a discussion does indicate that any analogy between SNR gains and mutual information is fraught with danger. In the remainder of this Section we will examine this in further detail. Using simple examples, we show that such an analogy is generally false. In particular, in Chapeau-Blondeau (1999), a calculation of the maximum mutual information between the output and input—labelled as C_{out} —is made, as well as what appears to be a calculation of the maximum mutual information between the signal by itself and the signal plus noise—labelled as C_{in} . It is shown that it is possible for $C_{\text{out}} > C_{\text{in}}$. This appears on the surface to be a clear contradiction of the DPI. This contradiction is investigated in this Section.

Binary signals

We consider in this investigation only random input signals consisting of a random pulse train with two discrete levels, such as 0 and 1 or -1 and $+1$. With this signal, the model shown in Fig. 3.1 is equivalent to the binary memoryless channel often considered in information theory textbooks. The noise and the threshold are characteristics of the channel, and the addition of noise to the signal causes the input signal to be transmitted through the channel with errors. For the particular case of a binary *symmetric* memoryless channel, the typical analysis is to calculate the probability p that the input is inverted—that is, the probability of error—so that $p(y = 1|s = 0) = p(y = 0|s =$

3.3 The Data Processing Inequality and SR

1) = p . If we had $p(y = 1|s = 0) \neq p(y = 0|s = 1)$ then the channel would not be symmetric.

The channel capacity of a discrete memoryless channel is defined as

$$C = \max_{p(x)} I(X, Y) \quad (3.5)$$

where the maximum is taken over all possible input Probability Density Functions (PDFs), $p(x)$. Therefore, for a binary channel, the channel capacity occurs for the value of $p(y = 0)$ —since $p(y = 1) = 1 - p(y = 0)$ —that maximises $I(X, Y)$.

Does this model form a Markov chain?

The DPI states that $I(s, x) \geq I(s, y)$ provided s, x and y form a Markov chain $s \rightarrow x \rightarrow y$. This means that the conditional distribution of y depends only on x and is conditionally independent of s (Cover and Thomas 1991). For the model illustrated in Fig. 3.1, since $y = T[x]$, y is a function of x and thus y is conditionally independent of s . Hence the model forms a Markov chain $s \rightarrow x \rightarrow y$.

We are now ready to begin our investigation of the implications of the DPI in systems like that shown in Fig. 3.1, subject to binary signals. We provide two examples, firstly the asymmetric binary channel, and secondly, a form of binary erasure channel.

3.3.2 Example 1: Asymmetric Binary Channel

Let $T[\cdot]$ be defined such that

$$y(t) = T[x(t)] = \begin{cases} 1 & \text{if } x(t) = s(t) + n(t) \geq \theta, \\ 0 & \text{otherwise.} \end{cases} \quad (3.6)$$

Let the input signal, $s(t)$, be a discrete time random binary signal such that both possible values of s are equally likely. Hence, s is a discrete random variable with probability mass function $P(s) = 0.5, s \in \{0, 1\}$. The noise is assumed to be continuously valued.

Noiseless case

In the noiseless case, $P(n = 0) = 1$, and if $0 < \theta \leq 1$, then the output, y , will be identical to the input. Also, $P(y = s) = 1$, and the system is a symmetric binary channel with $p = 0$. But if $\theta > 1$ then the output will be always zero—that is, $P(y =$

0) = 1. Hence, for a threshold greater than the maximum input signal—or less than the minimum input signal—then all information is lost, since it is impossible to infer anything at all about the input signal, other than that it is always subthreshold.

Finite nonzero noise and subthreshold signal

Let the threshold, $\theta = t > 1$, so that the maximum value of the signal is below the threshold by $t - 1$. Let the noise be uniformly distributed between 0 and b . Hence, the PDF of the noise is

$$p_n(n) = \begin{cases} \frac{1}{b} & 0 \leq n \leq b, \\ 0 & \text{otherwise.} \end{cases} \quad (3.7)$$

Consider only the case of $t - 1 \leq b < t$, so that the output, y , will sometimes be equal to 1 when $s = 1$, but never equal to 1 when $s = 0$. If $b < t - 1$, then y will always be zero, since the threshold is never crossed. If $b \geq t$ then sometimes the threshold will be crossed when $s = 0$. Thus the system is an asymmetric binary channel. Hence, by the addition of noise to the input signal, the output becomes correlated with the input signal—some information is conveyed; the output is only ever 1 when the input is 1—an improvement over the case of zero noise, where the output conveys no information about the input.

Such a channel is a specific case of those already analysed in Chapeau-Blondeau (1997b) in the SR context, where a calculation of the mutual information for the general case of an arbitrary noise distribution and arbitrary $P(s)$ is made. Chapeau-Blondeau (1997b) also calculates the channel capacity and shows the existence of a noise induced maximum. Here we consider a specific case simply as a means of illustrating the validity of the DPI.

Comparing $I(s, x)$ and $I(s, y)$

It is shown in Section A.1 of Appendix A that the mutual information between s and x is always one bit per sample, that is, $I(s, x) = 1$. This is consistent with the fact that the noise is small enough to allow $b < 1$ always, and therefore the input signal can be determined from the noisy signal without error, if a threshold is placed between b and 1.

It is also shown in Section A.1 of Appendix A that

$$I(s, y) = 0.5 + P_e \log_2 P_e - (0.5 + P_e) \log_2 (0.5 + P_e) \quad \text{bits per sample,} \quad (3.8)$$

3.3 The Data Processing Inequality and SR

where $P_e = \frac{t-1}{2b}$ is the probability of an error. However, $b \geq t-1$, $0 < b < 1$ and $t-1 > 0$, therefore $0 < P_e \leq 0.5$. Hence $0.5 < 0.5 + P_e \leq 1$. Therefore the second term in Eqn. (3.8) is always negative and the maximum value of the third term is 0.5. Therefore $I(s, y) < 1$ always. Since $I(s, x) = 1$ then $I(s, x) > I(s, y)$ always and since s, x and y form a Markov chain $s \rightarrow x \rightarrow y$, the DPI holds for this example.

Note that since $t > 1$ that $P_e > 0$. As $t \rightarrow 1$, $P_e \rightarrow 0$ and therefore $I(s, y) \rightarrow 1$. When $t-1 = b$, $P_e = 0.5$ and $I(s, y) = 0$. It might seem strange that even though the probability of error is 0.5 that the mutual information is zero, however this corresponds to the case where, since the output is always zero, nothing can be said about the input.

3.3.3 Example 2: Binary Erasure Channel

Our second example, as considered in Chapeau-Blondeau (1999), is a form of *binary erasure channel*. This example can also be described by Fig. 3.1, where all signals are discrete time. Suppose the input signal, $s(t)$, is a discretely valued random signal that can have values $\pm s_v$. Suppose also that the output signal, $y(t)$, can have three values, $(-s_v, 0, s_v)$, which are determined by two thresholds, $\pm\theta$ such that $T[\cdot]$ is defined by

$$y = T[x] = \begin{cases} s_v & \text{if } x = s + n \geq \theta, \\ -s_v & \text{if } x = s + n \leq -\theta, \\ 0 & \text{otherwise,} \end{cases} \quad (3.9)$$

where an output value of zero indicates the complete erasure of an input value, rather than its corruption.

However, in the textbooks (Cover and Thomas 1991), the *binary erasure channel* is a channel in which there is no possibility of error, only a possibility of erasure, and hence there is no possibility of an inverted output. The channel in Chapeau-Blondeau (1999) does allow for this possibility, so it is not equivalent to the classical binary erasure channel. Hence, we will refer to the channel considered in Chapeau-Blondeau (1999) as the *CB erasure channel*.

Chapeau-Blondeau (1999) shows, for a system where $\theta > 1$ and therefore the signal is subthreshold for even noise probability densities, that channel capacity occurs when $p(s = -1) = p(s = 1) = 0.5$. As expected, this capacity is shown to have a maximum for nonzero noise, as the addition of noise allows threshold crossings to occur that otherwise would not have. Chapeau-Blondeau (1999) also claimed to show that the

capacity at the output—that is, the maximum mutual information about s contained in y —could be greater than the capacity at the input, that is, the maximum mutual information about s contained in $s + n$. Note that this possibility appears to be ruled out by the DPI. The remainder of this Section indicates the source of this apparent contradiction.

Binary symmetric channel

Suppose $\theta = 0$, and the probability of error given the input, p_e , is known, and that $p_e = p(y = s_v | s = -s_v) = p(y = -s_v | s = s_v)$. Such a channel is a *binary symmetric channel*, and the possibility of erasure no longer exists. The mutual information of this channel (Cover and Thomas 1991) is

$$I(s, y) = H(y) - H(p_e). \quad (3.10)$$

The channel capacity occurs when $p(\pm s_v) = 0.5$ (Cover and Thomas 1991) and is given by

$$C_1 = 1 + p_e \log_2(p_e) + (1 - p_e) \log_2(1 - p_e). \quad (3.11)$$

If the noise in the channel has an even PDF, then $p_e = F_n(s_v)$, where $F_n(\cdot)$ is the Cumulative Distribution Function (CDF) of the noise.

Capacity at the output of the CB erasure channel

For the classical binary erasure channel, where there is no possibility of error, only of erasure, if the probability of erasure is denoted as α , then the capacity is given by $C = 1 - \alpha$, where capacity is achieved when $p(\pm s_v) = 0.5$ (Cover and Thomas 1991).

The channel considered in Chapeau-Blondeau (1999), in which there is a possibility of error as well as erasure—so that $\theta > 0$ —is still a symmetric memoryless discrete channel, provided the noise has an even PDF. Hence, for this case, capacity is achieved when $p(\pm s_v) = 0.5$. Chapeau-Blondeau (1999) derives formulas for the mutual information for such a channel, from which channel capacity can be found by taking the input probabilities as $p(\pm s_v) = 0.5$. These formulas depend on the CDF of the noise.

For example, let the threshold values be $\theta = \pm 1.1s_v$ (subthreshold signal) and the noise be uniform with a variance of $\sigma_n^2/12$. Hence, the noise has PDF

$$p(n) = \begin{cases} \frac{1}{\sigma_n} & -\sigma_n/2 \leq n \leq \sigma_n/2, \\ 0 & \text{otherwise.} \end{cases} \quad (3.12)$$

3.3 The Data Processing Inequality and SR

Therefore, the CDF of the noise is

$$p(n < u) = F_n(u) = \begin{cases} 1 & u \geq \sigma_n/2, \\ \frac{u}{\sigma_n} + \frac{1}{2} & -\sigma_n/2 \leq u \leq \sigma_n/2, \\ 0 & u \leq -\sigma_n/2. \end{cases} \quad (3.13)$$

Using Eqn. (3.13), along with Eqns. (10)-(17) from Chapeau-Blondeau (1999), the channel capacity for the CB erasure channel—which we label as C_2 —can be calculated numerically.

Fig. 3.2 shows a plot of the capacity of the CB erasure channel, C_2 , against the rms value of the uniformly distributed noise when $\theta = \pm 1.1s_v$. Note that the rms value of the noise is given by $\sigma_n/\sqrt{12}$. It also shows the capacity of the binary symmetric channel, calculated using Eqn. (3.11) with, as in Chapeau-Blondeau (1999), $p_e = F_n(1)$, since the noise PDF is symmetric. Note that the maximum value of the capacity for the CB erasure channel corresponds to a nonzero value of noise. This shows that SR can occur in the channel capacity measure, verifying the results of Chapeau-Blondeau (1997b).

The capacity of the binary symmetric channel is a decreasing function of the noise rms amplitude, except for small values, where the capacity is one bit per sample. This makes sense, since for small values of noise, the probability of error is zero—for example, by using a threshold at zero—and for larger values of noise, the probability of error increases as the rms noise amplitude increases.

Capacity at the input to CB erasure channel

So far we have calculated and plotted the channel capacity between the input and output of two different channels. In Chapeau-Blondeau (1999), the intention was to find an analogy between channel capacity and SNR gains. To achieve this, Chapeau-Blondeau (1999) needed to define an input channel capacity and an output channel capacity. With reference to Fig. 3.1, the only way this is possible is to say the input capacity is the maximum mutual information between s and x , and the output capacity is the maximum mutual information between s and y .

However, in Chapeau-Blondeau (1999), C_1 is considered to be the maximum mutual information at the *input* of the CB erasure channel, and C_2 to be the maximum mutual information at the *output* of the CB erasure channel. When the rms noise amplitude is such that $C_2 > C_1$, this is interpreted to mean that the capacity at the output of the

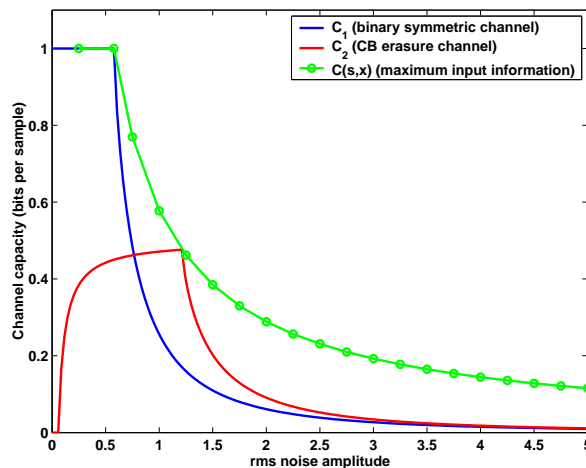


Figure 3.2. Channel capacity for several channels. Plot of channel capacity for the binary symmetric channel, C_1 , the CB erasure channel, C_2 , and the maximum input capacity, $C(s, x)$, against the rms value of uniform input noise. The maximum value of the capacity for the CB erasure channel corresponds to a nonzero value of noise. This shows that SR can occur in the channel capacity measure, verifying the results of Chapeau-Blondeau (1997b). This maximum can be interpreted to mean that a certain nonzero value of noise can minimise the information lost in the channel for a binary input signal. The capacity of the binary symmetric channel is a decreasing function of the noise rms amplitude, except for small values, where the capacity is one bit per sample. It can be seen that the capacity at the output of both the binary symmetric channel and the CB erasure channel is less than or equal to the capacity at the input, $C(s, x)$, and that hence the DPI holds. The input capacity is reached by the binary symmetric channel when the rms noise value is less than $2/\sqrt{12}$, and by the CB erasure channel when the rms noise value is about 1.215.

CB erasure channel is greater than the capacity at the input. According to the DPI, the information between s and y is always less than or equal to the information between s and x . Thus, there is a contradiction between the interpretation of Chapeau-Blondeau (1999), and the DPI.

Source of the apparent contradiction

Chapeau-Blondeau (1999) interpreted the result obtained to mean that for noise values where the ratio $C_2 > C_1$ is greater than one, the capacity at the output of the CB erasure channel is greater than the capacity at the input.

However, in contrast to Chapeau-Blondeau (1999), C_1 is the channel capacity of the binary symmetric channel, not the erasure channel, and hence cannot be considered

3.3 The Data Processing Inequality and SR

as the maximum input mutual information to the CB erasure channel. The DPI states that the mutual information at the output of the CB erasure channel is less than or equal to the maximum mutual information at the input of that channel, which is also the *input* to the binary symmetric channel. Thus Chapeau-Blondeau (1999) has shown only that for the values of rms noise where the ratio $C_2/C_1 > 1$, more information can be obtained about s by the CB erasure channel than for the binary symmetric channel. This does not show that the output of the CB erasure channel is more detectable than the input, as claimed in Chapeau-Blondeau (1999).

Using the results plotted in Fig. 3.2, we can obtain a graph for the ratio of C_2 to C_1 , as shown in Fig. 3.3 for a uniform noise distribution and for $\theta = 1.1$. Note that Fig. 3.3 is identical to the plot produced in Chapeau-Blondeau (1999) for uniform noise and $\theta = 1.1$. Where this ratio is greater than unity, the CB erasure channel has greater capacity than the binary symmetric channel.




Figure 3.3. Ratio of C_2 to C_1 . This plot shows the ratio of C_2 (capacity of CB erasure channel, with $\theta = 1.1$) to C_1 (capacity of binary symmetric channel) against the rms value of the uniformly distributed noise. For values of noise rms amplitude where the ratio is greater than unity, the CB erasure channel has greater capacity than the binary symmetric channel.

The actual capacity at the input, $C(s, x)$, can only really be taken to be the maximum of the mutual information between the input signal, s , and the input signal plus noise, $x = s + n$. As in the example of Section 3.3.2, the mutual information between x and s can be derived, this time using Eqn. (3.12) for the noise PDF. The derivation is given in

Section A.2 of Appendix A and finds that the mutual information between s and x is

$$I(s, x) = \begin{cases} 1 & \sigma_n < 2, \\ \frac{2}{\sigma_n} & \sigma_n \geq 2. \end{cases} \quad (3.14)$$

The following assumes (without proof) that due to the symmetry involved, capacity does occur when $p(\pm s_v) = 0.5$ and is given by Eqn. (3.14). Thus, $C(s, x)$ is plotted in Fig. 3.2. It can be seen that the capacity at the output of both the binary symmetric channel and CB erasure channel is less than or equal to the capacity at the input, $C(s, x)$, and that hence the DPI holds. Fig. 3.2 indicates that capacity is reached by the binary symmetric channel when the rms noise value is less than $2/\sqrt{12}$, and capacity is reached by the CB erasure channel when the rms noise value is about 1.215.

For the binary symmetric channel, where there is a noise induced maximum in the capacity, this maximum can be interpreted to mean that a certain nonzero value of noise can minimise the information lost in the channel.

3.3.4 Discussion

From this investigation, it is clear that although SNR gains may exist due to SR for periodic input signals, no information theoretical analogy exists for random noisy aperiodic signals. The simplest illustration of this is to threshold a noisy binary pulse train at its mean. For uniform noise with a maximum value less than half the pulse amplitude, the mutual information between input and output remains constant, regardless of the input SNR.

Furthermore, it has been demonstrated that since the DPI holds, the addition of more noise to a noisy signal cannot be of benefit as far as obtaining an input-output mutual information gain is concerned.

Such a result does not rule out the fact that the addition of noise at the input to a channel can maximise the mutual information at the output, in other words, the effect of SR for aperiodic signals is perfectly valid. When this occurs, an optimal value of input noise means minimising the information lost in the channel.

3.4 Cross-Spectral Measures

The aim of this Section is to examine the merits of using a frequency dependent signal-to-noise Ratio (SNR) measure for SR research. This SNR measure requires a calculation of the Cross-power Spectral Density (CSD) between the input and output signals in a system. It was first proposed as a measure for SR in Kiss (1996), but apart from a handful of other papers (Sethuraman and Kish 2003, Mingesz *et al.* 2005), has not been otherwise discussed in the SR literature.

As discussed in Chapter 2, all early work on SR considered only periodic input signals. As pointed out in DeWeese and Bialek (1995), for studies of neural systems such signals are not particularly realistic. Hence, Collins *et al* soon popularised the term Aperiodic Stochastic Resonance (ASR) in several studies, measuring its effects using correlation based measures. Independently, Kiss (1996) also recognised the need to consider systems with aperiodic inputs. However, the driving force for Kiss was the fact that it had been shown conclusively by linear response theory that for small periodic input signals the output SNR is always less than the input SNR (DeWeese and Bialek 1995). Kiss was looking for ways to obtain an SNR gain due to SR, and hence looked to systems not covered under the conditions for which linear response theory applies. This includes aperiodic signals, and signals that are strong compared to additive noise.

Hence, Kiss was motivated to consider the case of a broadband—and aperiodic, as opposed to broadband periodic—input signal. He recognised that in this case, the usual means of calculating SNR used in the SR literature—the ratio of the output signal power at the input frequency to the background noise Power Spectral Density (PSD)—is inadequate, since the output noise is usually taken as the output signal that occurs in the absence of an input signal. He noted that for nonlinear systems that the input signal can interact with the input noise to yield extra cross-modulation spectra at the output—see also Section 3.1.1. This is extra output noise that is not present in the absence of a signal. Hence, the output noise spectrum cannot be taken as that which occurs when no signal is present.

For the computation of output SNRs in such systems, it is therefore necessary to be able to separate output signal and noise components for all frequencies at which the signal exists. Hence Kiss (1996) proposed—in particular, for the context of SR research—that an appropriate means of achieving this separation is to classify the output signal as the component of the output spectrum that is the magnitude of the CSD between the

input signal and the system output, divided by the PSD of the system input, with the remainder of the output spectrum being noise.

Cross-spectral measures similar to those used in Kiss (1996) have also been used elsewhere in SR research. In particular, Neiman *et al.* (1997) uses the *coherence function* (Carter 1993, Bendat 1998) to calculate cross-correlation measures for ASR. We shall soon see that the SNR measure proposed in Kiss (1996) can be expressed as a renormalisation of the coherence function. Later, Goychuk and Hänggi (1999) and Nikitin and Stocks (2004) make use of a result that relates mutual information to correlation coefficient, by way of cross-spectral densities (Pinsker 1964).

It is shown in the remainder of this Section that the proposed separation of output signal and noise spectra has a strong basis in signal processing theory and applications. It is also shown to be related to other standard signal processing techniques and measures, including the coherence function, time delay estimation (Carter 1993) and Wiener-Kolmogorov filtering (Poor 1994). To illustrate an application in SR research, the method is applied to measure ASR in a model neural system driven by a broadband input signal. Firstly, however, the next Section gives the definition of the CSD function, and derivations of Kiss' frequency dependent SNR formula.

3.4.1 Power and Cross-Power Spectral Densities

The PSD, $S_{xx}(f)$, of an ergodic, finite power signal—such as a stationary random signal— $x(t)$, is given by the Fourier transform of its autocorrelation function (Bendat 1998), $R_{xx}(\tau)$, as

$$S_{xx}(f) = \int_{-\infty}^{\infty} R_{xx}(\tau) \exp(-j2\pi\tau) d\tau. \quad (3.15)$$

The CSD of two such signals, $x(t)$ and $y(t)$, is given by the Fourier transform of the cross correlation of x and y , $R_{xy}(\tau)$, as

$$S_{xy}(f) = \int_{-\infty}^{\infty} R_{xy}(\tau) \exp(-j2\pi\tau) d\tau. \quad (3.16)$$

Unlike the PSD function, which is always a real-valued function of frequency, the CSD is a complex-valued function of frequency. A well known inequality (Bendat 1998) relates the CSD to the PSDs of x and y ,

$$|S_{xy}(f)|^2 \leq S_{xx}(f)S_{yy}(f). \quad (3.17)$$

3.4 Cross-Spectral Measures

One of the uses for the CSD function is the coherence function, which for two signals $x(t)$ and $y(t)$ is defined as (Bendat 1998)

$$\Gamma_{xy}^2(f) = \frac{|S_{xy}(f)|^2}{S_{xx}(f)S_{yy}(f)}. \quad (3.18)$$

This function is real valued between zero and unity, and is a measure of the linearity between x and y . That is, if a perfect linear relationship exists between x and y at frequency f , then $|S_{xy}(f)|^2 = S_{xx}(f)S_{yy}(f)$ and the coherence function will be equal to unity at that frequency. Thus, the coherence function can be considered to be a correlation coefficient for the frequency domain.

The coherence function has been used previously in the context of SR, both for conventional periodic SR and aperiodic SR (Neiman *et al.* 1997, Anishchenko *et al.* 1999). The first work on ASR used the cross-correlation measure (Collins *et al.* 1995a). It was subsequently noted (Neiman *et al.* 1997) that if the coherence function can be found, the cross correlation measure and the correlation coefficient easily follow since the correlation between the input and output is

$$C_{xy} = \int_f \text{Re}[S_{xy}(f)]df. \quad (3.19)$$

The input and output autocorrelation functions, or mean square powers, are similarly obtained from the integral over all f of their PSDs. Hence, the correlation coefficient is

$$\rho_{xy} = \frac{\int_f \text{Re}[S_{xy}(f)]df}{\sqrt{\int_f S_{xx}(f)df \int_f S_{yy}(f)df}}. \quad (3.20)$$

Linear systems

CSDs can arise when two correlated signals are added. For example, if the two signals are $s(t)$ and $n(t)$, the result $x(t) = s(t) + n(t)$ has auto-correlation

$$\begin{aligned} R_{xx}(\tau) &= \langle x(t)x(t-\tau) \rangle \\ &= \langle [s(t) + n(t)][s(t-\tau) + n(t-\tau)] \rangle \\ &= R_{ss}(\tau) + R_{sn}(\tau) + R_{ns}(\tau) + R_{nn}(\tau). \end{aligned} \quad (3.21)$$

Taking the Fourier transform of both sides gives the PSD of $x(t)$ in terms of the spectral densities of $s(t)$ and $n(t)$ and the CSD of $s(t)$ and $n(t)$ as

$$S_{xx}(f) = S_{ss}(f) + 2\text{Re}[S_{sn}(f)] + S_{nn}(f), \quad (3.22)$$

where we have used the property that $S_{sn}(f) = S_{ns}^*(f)$, where $*$ denotes the complex conjugate. Note that if $s(t)$ and $n(t)$ are uncorrelated then $S_{sn}(f)$ is zero for all f and the output PSD is the sum of the PSDs of $x(t)$ and $n(t)$.

We may also be interested in the CSD of $s(t)$ and $x(t)$, i.e. $S_{sx}(f)$. This is, again, the Fourier transform of the cross correlation of $s(t)$ and $x(t)$, which is

$$\begin{aligned} R_{sx}(\tau) &= \langle s(t)x(t-\tau) \rangle \\ &= \langle s(t)[s(t-\tau) + n(t-\tau)] \rangle \\ &= R_{ss}(\tau) + R_{sn}(\tau). \end{aligned} \quad (3.23)$$

Taking the Fourier transform of both sides gives

$$S_{sx}(f) = S_{ss}(f) + S_{sn}(f). \quad (3.24)$$

The magnitude squared of S_{sx} is given by

$$\begin{aligned} |S_{sx}(f)|^2 &= S_{sx}(f)S_{sx}^*(f) \\ &= (S_{ss}(f) + S_{sn}(f))(S_{ss}(f) + S_{sn}(f))^* \\ &= S_{ss}(f)(S_{ss}(f) + 2\text{Re}[S_{sn}(f)]) + |S_{sn}(f)|^2. \end{aligned} \quad (3.25)$$

Substituting Eqn. (3.22) into Eqn. (3.25) gives

$$|S_{sx}(f)|^2 = S_{ss}(f)(S_{xx}(f) - S_{nn}(f)) + |S_{sn}(f)|^2 \quad (3.26)$$

and the coherence function is therefore

$$\Gamma_{sx}^2(f) = 1 - \frac{S_{nn}(f)}{S_{xx}(f)} + \frac{|S_{sn}(f)|^2}{S_{ss}(f)S_{xx}(f)} \quad \forall f \text{ s.t. } S_{ss}(f)S_{xx}(f) \neq 0. \quad (3.27)$$

Uncorrelated signal and noise

If the signal and noise are uncorrelated then $S_{sn}(f) = 0$ and from Eqn. (3.22) the sum of the PSDs is the output PSD. Therefore from Eqn. (3.26)

$$|S_{sx}(f)|^2 = S_{ss}^2(f), \quad (3.28)$$

and the coherence function between s and x is

$$\Gamma_{sx}^2(f) = \frac{S_{ss}(f)}{S_{ss}(f) + S_{nn}(f)}, \quad (3.29)$$

and is only unity when noise is absent. Thus, the coherence function can be reduced from unity by the presence of noise, even when the signal and noise are added and uncorrelated, simply due to the noise being nonzero at frequencies present in the signal.

Nonlinear systems

We now consider the general case of a system with input signal, $s(t)$, subject to additive noise, $n(t)$, so that the overall input is $x(t) = s(t) + n(t)$. Let the output be the result of a nonlinear transformation of $x(t)$, i.e. $y(t) = T[x(t)]$. Then the PSDs of s and y are $S_{ss}(f)$ and $S_{yy}(f)$ respectively, and $S_{sy}(f)$ is the CSD between s and y .

According to the ideas in Kiss (1996), an ideal metric for such a system is one which will give a consistent, frequency dependent, phase independent measure for an output SNR under all circumstances, including either single frequency periodic, broadband periodic, or aperiodic input signals, and regardless of whether the signal is subthreshold or not, or has a small amplitude compared to the noise or not. To this end Kiss (1996) defined the following “generalised amplification” factor, which is a complex function of frequency,

$$K(f) = \frac{S_{sy}(f)}{S_{ss}(f)}. \quad (3.30)$$

Using carets to indicate output signal and noise, as opposed to input signal and noise, we can simplify the notation used by Kiss to define the “generalised output signal” spectrum as

$$\hat{S}_{ss}(f) = S_{ss}(f)|K(f)|^2 = \frac{|S_{sy}(f)|^2}{S_{ss}(f)}. \quad (3.31)$$

Note that since the PSD is a real number, $|S_{ss}(f)|^2$ is simply $S_{ss}(f)^2$. Kiss (1996) notes that $\hat{S}_{ss}(f)$ can be considered to be the spectrum of the signal component of the output signal and that the remainder of the output spectrum to be noise, so that

$$\hat{S}_{nn}(f) = S_{yy}(f) - \hat{S}_{ss}(f). \quad (3.32)$$

Substituting Eqn. (3.31) into Eqn. (3.32) gives the output noise spectrum in terms of the signal PSD, the CSD and the overall output PSD as

$$\hat{S}_{nn}(f) = S_{yy}(f) - \frac{|S_{sy}(f)|^2}{S_{ss}(f)}. \quad (3.33)$$

Thus, the output SNR obtained with these quantities is a function of frequency given by

$$\text{SNR}(f) = \frac{\frac{|S_{sy}(f)|^2}{S_{ss}(f)}}{S_{yy}(f) - \frac{|S_{sy}(f)|^2}{S_{ss}(f)}}. \quad (3.34)$$

We now add several observations and caveats to this measure.

Firstly, in order to avoid infinities, the quantity $K(f)$, and therefore the output SNR, must be defined only for frequencies for which $S_{ss}(f)$ is nonzero and not a delta function. Delta functions will occur in the input spectrum if the input has a nonzero mean, or contains periodic components. To avoid this scenario, we will consider only non-periodic, zero mean, ergodic random signals.

Secondly, the “generalised amplification” factor can be seen to look like the same quantity as the susceptibility function, $\chi(f)$, used in the context of linear response theory, for which it is stated that

$$S_{sy}(f) = \chi(f)S_{ss}(f). \quad (3.35)$$

Linear response theory also states that for a sufficiently weak input signal, the output spectrum is

$$S_{yy}(f) = S_{yy}^0(f) + |\chi(f)|^2 S_{ss}(f), \quad (3.36)$$

where $S_{yy}^0(f)$ is the output PSD which occurs when $s(t) = 0$ (Neiman *et al.* 1997). However, for the general case we are considering, these equations are not valid, since in linear response theory, the susceptibility is taken to be constant for a system, regardless of the input, whereas $K(f)$ is defined for a combination of the system and input.

Thirdly, from Eqn. (3.18), we can rewrite Eqns. (3.31) and (3.32) in terms of the coherence function as

$$\hat{S}_{ss}(f) = \Gamma_{sy}^2(f)S_{yy}(f) \quad (3.37)$$

and

$$\hat{S}_{nn}(f) = (1 - \Gamma_{sy}^2(f))S_{yy}(f). \quad (3.38)$$

Hence, if s and y are linearly related, then $1 - \Gamma_{sy}^2(f) = 0$, the output signal is the entire output spectrum, and the output noise is zero. Hence, the SNR can be written as

$$\text{SNR}(f) = \frac{\Gamma_{sy}^2(f)}{1 - \Gamma_{sy}^2(f)}. \quad (3.39)$$

It turns out that Eqn. (3.39) is not new. This relation between frequency dependent SNR and the coherence function—as well as the same means of output signal and noise separation—can be found in Gabbiani (1996) and Borst and Theunissen (1999). In particular, it has been applied to neural systems, at least as early as 1991 by Bialek *et al.* (1991).

It can also be seen that this formulation of SNR is effectively a renormalisation of the coherence function. If the coherence function approaches unity, then the SNR approaches

3.4 Cross-Spectral Measures

infinity. If the coherence function is zero, then the SNR is zero. If the coherence function is one half, then the SNR is unity. Hence, measuring system performance by the SNR gives exactly the same information as the coherence function, and vice versa. The only possible advantage of using this SNR measure rather than the coherence function is that SNR may be more useful if comparing two systems with nearly the same performance, as the coherence function near zero or unity is effectively highly compressed when compared to the SNR.

Eqn. (3.39) can be rearranged to give the magnitude of the coherence function in terms of the SNR as

$$|\Gamma(f)| = \sqrt{\frac{\text{SNR}(f)}{1 + \text{SNR}(f)}}. \quad (3.40)$$

The very fact that the SNR can be expressed in terms of the coherence function, which is a measure of the linearity between two signals, may seem surprising since the original goal of Kiss (1996) was to derive a measure that is robust to nonlinear transformations of a signal. However, the root of this fact can be seen by re-examination of Eqns. (3.30) and (3.31). Even though $K(f)$ may not be a linear function of f , these equations *define* a linear relationship between the PSDs of $s(t)$ and $y(t)$, in the sense that $S_{sy}(f)$ is a linear function of $S_{ss}(f)$. Furthermore, this does not mean that the method is not applicable for highly nonlinear signal transfer, since $K(f)$ is not defined as being applicable for all possible input signal and noise combinations, but is defined for a *given* signal and noise combination.

Fourthly, note that electronic engineers also use Eqn. (3.30) in linear systems theory, except that $K(f)$ is usually specified as the ‘transfer function,’ $H(f)$, which is used to define a linear system. Again, the approach taken here differs from such linear systems theory, in that $K(f)$ is not defined by only the system, but also by the input signal and noise.

Applicability to linear systems

To illustrate that the frequency dependent SNR is applicable to linear systems, consider that when two signals $s(t)$ and $n(t)$ are added,

$$\begin{aligned}\hat{S}_{ss}(f) &= \frac{|S_{sx}(f)|^2}{S_{ss}(f)} \\ &= S_{xx}(f) - S_{nn}(f) + \frac{|S_{sn}(f)|^2}{S_{ss}(f)} \\ &= S_{ss}(f) + 2\text{Re}[S_{sn}(f)] + \frac{|S_{sn}(f)|^2}{S_{ss}(f)},\end{aligned}\quad (3.41)$$

and

$$\hat{S}_{nn}(f) = S_{nn}(f) - \frac{|S_{sn}(f)|^2}{S_{ss}(f)}.\quad (3.42)$$

If the signal and noise are completely uncorrelated then $S_{sn}(f)$ is zero and the output signal and noise PSDs are simply the input signal and noise PSDs. Thus the SNR of Eqn. (3.34) is simply the conventional linear systems definition,

$$\text{SNR}(f) = \frac{S_{ss}(f)}{S_{nn}(f)}.\quad (3.43)$$

Notice also that the coherence function is given as in Eqn. (3.29) and substitution of this expression into Eqn. (3.39) gives exactly Eqn. (3.43).

3.4.2 New Interpretation

The quantity labelled by Kiss (1996) as the “generalised output signal,” that is, $\hat{S}_{ss}(f)$, can be interpreted as an estimate for the input signal PSD. The process of obtaining $\hat{S}_{ss}(f)$ is implicitly a linear filtering operation, which depends on knowledge of the input PSD. Hence, this filtering is similar to matched filtering, which also depends on knowledge of the input signal. To illustrate this, consider the system as a ‘black-box,’ which has as input a known input signal, $s(t)$, and an output $y(t)$. The processes internal to the black box are the addition of the input noise, $n(t)$, to $s(t)$ and the nonlinear transformation, $T[s(t) + n(t)]$. To obtain \hat{S}_{ss} , the CSD of the output of the black-box, and the input signal, $s(t)$ is found, and divided by the PSD of the input signal. These two operations form the filtering. The SNR can then be found from Eqn. (3.34).

It is evident that such a filter will not be useful if the input signal’s PSD is unknown, since then there is no means of obtaining \hat{S}_{ss} and therefore the output SNR cannot

3.4 Cross-Spectral Measures

be determined. However, this method can provide a useful means of characterising a system, and determining performance for a known input signal, when the system transfer function is unknown. An example of such a scenario is where a known signal is applied to a physiological population of neurons, and the output is measured further down the neural pathway.

The formulation of Kiss (1996) has many similarities to the Wiener-Kolmogorov filter (Poor 1994). In this filter, the output time signal, $y(t)$, is convolved with the linear filter given by

$$K(\tau) = F^{-1} \left(\frac{S_{sy}^*(f)}{S_{yy}(f)} \right), \quad (3.44)$$

where F^{-1} is the inverse Fourier transform. This means that in the frequency domain

$$\hat{S}_{ss}(f) = \frac{S_{sy}^*(f)}{S_{yy}(f)}. \quad (3.45)$$

Such a technique is well known in signal processing, and has been applied frequently in the computational neuroscience literature in the analysis of neural spike trains (Gabbiani 1996, Manwani and Koch 1998, Manwani and Koch 1999).

Note that if the spectrum of the output signal is defined as

$$\hat{S}_{ss}(f) = \frac{|S_{sy}(f)|^2}{S_{yy}(f)}, \quad (3.46)$$

rather than as in Eqn. (3.31), then the noise spectrum at the output can be considered as the difference between the input spectrum and the output spectrum, which is

$$\hat{S}_{nn}(f) = S_{ss}(f) - \hat{S}_{ss}(f) = S_{ss}(f) - \frac{|S_{sy}(f)|^2}{S_{yy}(f)}, \quad (3.47)$$

and the SNR at the output is

$$\text{SNR}(f) = \frac{\frac{|S_{sy}(f)|^2}{S_{yy}(f)}}{S_{ss}(f) - \frac{|S_{sy}(f)|^2}{S_{yy}(f)}} = \frac{\Gamma_{sy}^2(f)}{1 - \Gamma_{sy}^2(f)}, \quad (3.48)$$

which is precisely the same formula as Eqn. (3.39).

3.4.3 Applying the Frequency Dependent SNR Formula

A simple way to illustrate the possible use of this work is to apply the generalised frequency dependent SNR formula to a model neural system driven by a broadband

and aperiodic input signal. In line with much SR research, we use a very simple model of a neuron that encapsulates its main nonlinearity—that of a threshold. The output from the neuron is a spike when the input, $s(t)$, is greater than the neuron's threshold, θ . We also assume additive Gaussian white noise, $n(t)$, at the input to the neuron, so that the output contains noise due to both the nonlinearity and the input noise. Thus, assuming an infinitesimal refractory time, the neuron's output is

$$y(t) = \begin{cases} 1 & \text{if } s(t) + n(t) > \theta, \\ 0 & \text{otherwise.} \end{cases} \quad (3.49)$$

The remainder of this Section examines the use of the frequency dependent SNR formula of Eqn. (3.39) by simulating input and output signals from this neural model.

Note that care must be taken when attempting to calculate spectral densities and SNRs by simulation. Due to the necessity of simulating signals by discrete time series, the Discrete Fourier Transform (DFT) must be used. This has not always been recognised in SR research, but Mitaim and Kosko (1998) give a clear view on how to use the DFT to calculate SNR (page 2157).

Practical implementation of the DFT is usually achieved by the Fast Fourier Transform (FFT) algorithm (Proakis and Manolakis 1996). Inherent in the FFT is a gain factor which depends on the number of samples used in the FFT. Furthermore, unlike for finite energy signals such as pulses, for random signals the magnitude of the FFT is not necessarily the correct quantity to take as the PSD. Instead, estimation techniques such as Welch's averaged, modified periodogram method are required (So *et al.* 1999).

Broadband random input signal

The simplest way to illustrate the effectiveness of the frequency dependent SNR measure is to let the input signal be a bandpass Gaussian random signal, $s(t)$, with zero mean and unity variance. Let the signal have a bandwidth of 2 kHz between 4 and 6 kHz, with a sampling frequency of 20 kHz. A 200 sample realisation of such a signal is shown in Fig. 3.4. This realisation was produced by applying an elliptic filter to a sequence of 200 samples drawn the unity variance Gaussian distribution. The PSD of such a signal—obtained using Welch's averaged, modified periodogram method (So *et al.* 1999) with 100,000 samples—is shown in Fig. 3.5.

Let the noise signal, $n(t)$, be Gaussian white noise with variance σ^2 . For a sampling frequency of 20 kHz, the noise bandwidth is therefore effectively the Nyquist frequency

3.4 Cross-Spectral Measures




Figure 3.4. Bandpass signal example. A 10 millisecond realisation of a Gaussian bandpass signal, generated by applying an elliptic filter to a sequence of 200 samples drawn from a Gaussian distribution with a unity variance. The bandwidth is 2 kHz, between 4 and 6 kHz. The sampling frequency is 20 kHz.




Figure 3.5. PSD of bandpass signal. The PSD of the bandlimited signal illustrated in Fig. 3.4. The PSD was obtained from a sequence of 100,000 samples of the signal, by using Welch's averaged, modified periodogram method. The bandpass nature of the signal is clearly evident, with a power magnitude of 0 dB between 4 kHz and 6 kHz, and a power magnitude decreasing rapidly to lower than -20 dB outside this range.

of 10 kHz. Let the neuron's threshold be $\theta = 3$ (arbitrary units). Hence, since the signal has unity variance, it will nearly always be subthreshold, and hence our simulation is exactly that which is most commonly used in threshold-based SR studies; that of a subthreshold signal to which noise is added, in order to induce threshold crossings. To illustrate the various spectra involved, Fig. 3.6 shows plots of the input and output power spectral densities, and the magnitude of the CSD obtained for 1,000,000 samples, and a noise variance of 4. It turns out that this is close to the optimal value of noise variance for this signal and threshold value. Despite this, only a very slight increase—invisible in Fig. 3.6—in the output PSD is evident in the bandpass frequencies, relative to the stop band frequencies. However, the magnitude of the CSD is relatively large when compared to non-optimal values of noise, and it is this fact which leads to the SNR containing a maximum near this point, as we shall now see.

Fig. 3.7 shows the SNR of Eqn. (3.39) plotted against frequency for a number of noise intensities. It is clear that there must be a nonzero value of noise intensity which provides the optimal output response and that therefore SR occurs. Since Fig. 3.7 indicates that the output SNR is approximately constant for all pass-band frequencies, Fig. 3.8 shows the value of the output SNR for the center frequency of 5 kHz, as a function of noise standard deviation. Inspection of Fig. 3.8 verifies that there is indeed a noise-induced maximum in the SNR at this frequency. However, also note from Figs. 3.7 and 3.8 that the SNR is very small—less than -10 dB. Fig. 3.9 shows the plot obtained by normalising the SNR to obtain the magnitude of the coherence function, as in Eqn. (3.40). It can be seen that the peak value of the coherence magnitude is about 0.25.

3.4.4 Discussion

It is clear from the results presented in this Section that although we have shown that SR occurs in the frequency dependent SNR measure in a simple model neural system, the very low maximum output SNR means that the output signal is only slightly coherent with the input signal. For other forms of input signal we might expect the output SNR to be much higher than this—see, for example, Sethuraman and Kish (2003). The main conclusion to be drawn from these simulations is that the frequency dependent output SNR formulation outlined does indeed show SR behaviour.

3.4 Cross-Spectral Measures




Figure 3.6. PSDs and CSD of input and output signals. This plot shows the PSD of both the bandpass input signal, $x(t)$ —as also shown in Figure 3.5—and the output neural pulse train, $y(t)$. It also shows the magnitude of the CSD between the input and output signals. The noise variance is 4. The PSD of the output signal is virtually flat for all frequencies. However, the CSD magnitude is relatively large for the passband compared to the stopband.




Figure 3.7. Output SNR against frequency. Output SNR against frequency for four values of noise standard deviation, σ . For both small noise and very large noise, there is no noticeable increase in SNR in the passband frequencies when compared to the stopband. However for the intermediate values of noise, there is definitely a larger SNR for the passband. This indicates that a certain range of nonzero noise levels provide a better performance than the absence of noise, or too much noise, and that therefore SR occurs as expected. However, unlike the conventional single frequency SNR measure, the measure used here works for a broadband aperiodic input signal.

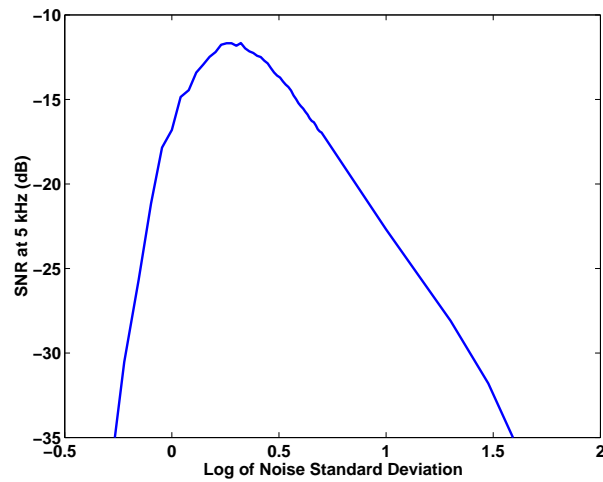


Figure 3.8. Output SNR against input noise. Output SNR at 5 kHz against log of noise standard deviation. This curve shows SR occurring in the simple neural model as expected. For smaller and larger noise standard deviations than the optimum value, either too few threshold crossings occur, or threshold crossings are almost entirely due to noise, thus causing very low SNRs.




Figure 3.9. Coherence magnitude against input noise. Magnitude of the coherence function at 5 kHz against log of noise standard deviation. This plot shows the same data as Fig. 3.8, but normalised to give the magnitude of the coherence function, as expressed in Eqn. (3.40). The peak coherence magnitude is about 0.25.

3.5 Chapter Summary

Also of significance is that the frequency dependent SNR formula has been shown to be closely related to the coherence function, and Wiener-Kolmogorov filtering. Since the latter technique has been applied extensively in studies of neural spike trains, we expect that it could be usefully applied to studies of SR in more realistic neural systems than that considered here.

3.5 Chapter Summary

The initial Section of this Chapter gave a brief discussion of the use of SNR measures in SR research, and pointed out that the SNR definition used for conventional, single frequency input signal SR work is irrelevant for broadband periodic or aperiodic signals.

The second Section discussed the SNR gain debate, and concluded that some of the reported SNR gains should not be considered to be “due to noise,” but due to the same mechanism as SR.

The third Section then analysed and discussed previous work that claimed to show an analogy to SNR gains for periodic signals in a system driven by an aperiodic signal. This analogy was shown to be false, both in abstract theory, and in the particular example studied.

Finally, the fourth Section of this Chapter discussed a frequency dependent SNR formula that does have some relevance to certain engineering problems, and showed that it is effectively the same measure as the coherence function.

3.5.1 Original Contributions for Chapter 3

This Chapter included the following original contributions:

- A demonstration that the Data Processing Inequality is not violated, despite an apparent contradiction to this in the literature.
- It is shown that the frequency dependent SNR formula first used in SR research by Kiss (1996) is equivalent to a renormalisation of the coherence function. Hence, using this formula to measure SR is equivalent to measuring it using the coherence function.

- The first application of the frequency dependent SNR formula to measure the coherence between an aperiodic bandpass signal in a simple neural model.
- It is pointed out that the frequency dependent SNR formula has previously been used in neural coding research, and is equivalent to Wiener-Kolmogorov filtering. Furthermore, this fact is used to interpret the spectra of the “generalised output signal” of Kiss (1996), as being an estimate for the input signal’s spectra.

3.5.2 Further Work

Possible future work and open questions arising from this Chapter might include:

- A more comprehensive review of engineering uses of the coherence function, and after obtaining this understanding, applying it more thoughtfully to SR research. Such a review would begin with a reading of Carter (1993).
- A more thorough investigation of the previous use of coherence function based measures in computational neuroscience. Such research could start by thoroughly understanding the application of Wiener-Kolmogorov filtering to spike-train analysis, such as the work presented in Gabbiani and Koch (1998) and Rieke *et al.* (1997).

This concludes Chapter 3. The next Chapter begins the main topic of this thesis—that of Suprathreshold Stochastic Resonance (SSR). As mentioned briefly in Chapter 2, SSR is a form of aperiodic stochastic resonance that occurs in arrays of parallel identical threshold devices.

Chapter 4

Suprathreshold Stochastic Resonance: Encoding

IN many of the systems and models in which Stochastic Resonance (SR) has been observed, the essential nonlinearity is effectively a single threshold. Usually SR occurs when an entirely subthreshold signal is subjected to additive noise, which allows threshold crossings to occur that otherwise would not have. In such systems, it is generally thought that when the input signal is suprathreshold, then the addition of noise will not have any beneficial effect on the system output.

However, the 1999 discovery of a novel form of SR in simple threshold-based systems showed that this is not the case. This phenomenon is known as Suprathreshold Stochastic Resonance (SSR), and occurs in arrays of identical threshold devices subject to independent additive threshold noise. In such arrays, SR can occur regardless of whether the signal is entirely subthreshold or not, hence the name *suprathreshold* SR.

This chapter presents a review of the original theoretical work on SSR. Novel theoretical extensions are also presented, as well as numerical analysis of previously unstudied input and noise signals, a new technique for calculating the mutual information by integration, and an investigation of a number of channel capacity questions for SSR. Finally, the chapter shows how SSR can be interpreted as a *stochastic quantisation* scheme.

4.1 Introduction

Suprathreshold Stochastic Resonance (SSR) is a form of Stochastic Resonance (SR) that occurs in arrays of identical threshold devices. A schematic model of the system is shown in Fig. 4.1, and is described in detail in Section 4.3. The discovery of SR in such a system was made by Stocks in 1999 (Luchinsky *et al.* 1999, Stocks 2000a)¹⁰. Stocks showed—using an aperiodic input signal, meaning that SSR is a form of Aperiodic Stochastic Resonance (ASR)—that SR can occur in Shannon’s average mutual information measure between the input and output of the array, under the constraint that all thresholds are set to the same value. Most importantly, SR occurs regardless of whether the input signal is entirely subthreshold or not, which is the first known occurrence of such behaviour in threshold devices. Stocks named this effect *suprathreshold stochastic resonance*, to distinguish it from the occurrence of SR in previous studies of single-threshold systems, and subsequently showed that the effect is maximised when all threshold values are set to the signal mean (Stocks 2001a). Fig. 4.1 also serves to describe the more general case when all thresholds are not identical, a case which is examined in detail in Chapter 8. In this chapter we will only consider the case of all thresholds set to the same value.

This Chapter contains original work on the topic of SSR that has been published, in part, in the open literature (McDonnell *et al.* 2001, McDonnell *et al.* 2002a, McDonnell and Abbott 2004b).

4.1.1 Chapter Structure

This chapter provides, in Section 4.2, a brief literature review that outlines the main results of all previous studies of SSR. Then in Section 4.3 we reproduce the most important results from this previous work, provide some extensions to theory, and numerically examine some hitherto unconsidered signal and noise distributions. This is followed in Section 4.4 by an investigation into a number of channel capacity questions for the SSR model. Finally, an interpretation of SSR as *stochastic quantisation* is introduced in Section 4.5. Some results relevant to this chapter are given in more detail in Appendix B.

¹⁰Note that the first description of SSR was given in a 1999 paper reviewing non-conventional forms of SR (Luchinsky *et al.* 1999). However, although this paper was published first, it references a paper submitted to *Physical Review Letters*, which was subsequently published in 2000 (Stocks 2000a).

4.2 Literature Review

4.2.1 Information Theoretic Studies of SR in Threshold Systems

As discussed in Chapter 2, although the term *aperiodic stochastic resonance* was popularised by work using correlation based measures (Collins *et al.* 1995a), the first papers to study ASR using mutual information (Levin and Miller 1996, Bulsara and Zador 1996, Heneghan *et al.* 1996) are of more relevance to SSR. These papers, along with DeWeese and Bialek (1995), paved the way for the use of information theory in SR research.

Of particular importance to SSR was the first study using mutual information to analyse a single threshold-device system (Bulsara and Zador 1996). This system is very similar to that discussed in Section 3.3.1 of Chapter 3, where the input signal to the system is a random binary signal, subject to continuously valued noise. The initial section of Bulsara and Zador (1996) calculates the input-output mutual information for this system for various threshold values, and shows that SR occurs for entirely subthreshold signals. Subsequently, Bulsara and Zador (1996) also consider mutual information in a leaky integrate-and-fire neuron model.

Also of relevance to SSR are the final paragraphs of Bulsara and Zador (1996), which discuss how future work might extend the system they considered. Specifically, Bulsara and Zador (1996) propose allowing continuously valued input signals, and an output signal which is the sum of the outputs of N binary threshold devices. It is stated that the mutual information in such a model would be of the order of $0.5 \log_2(N)$. These extensions are exactly what is carried out in the first work on SSR.

4.2.2 Original Work on SSR

The earliest work published in the open literature in which the phenomenon of SSR is presented, is in the second part of a two part review of SR in electrical circuits (Luchinsky *et al.* 1999). Although this work was published first, it references a paper submitted to *Physical Review Letters*, which was subsequently published in 2000 (Stocks 2000a). This latter paper contains the same results as Luchinsky *et al.* (1999) as well as some further analysis. In this original work of Stocks, the input signal, x , is taken to be a Gaussian random variable, and the noise on each threshold is likewise a Gaussian

4.2 Literature Review

random variable. No reduced analytical expressions are obtained for the mutual information, which is instead calculated by numerical integration and validated by digital simulation.

A later paper (Stocks 2001c) presents an analytical expression for the mutual information in the specific case of a uniformly distributed signal, and uniformly distributed threshold noise. Another longer paper (Stocks 2001a) examines the behaviour of SSR in the event of large N , as well as providing an exact expression for the mutual information in the event that both the signal and noise distributions are identical, that is, they have the same distribution and moments. Furthermore, Stocks (2001a) finds that the SSR effect is optimised in the case of all thresholds being set equal to the signal mean, as well as presenting an analysis of the output signal-to-noise ratio (SNR). Several conference and book chapters published a combination of the material presented in Stocks (2000a), Stocks (2001c) and Stocks (2001a). These include Stocks (2000c), Stocks (2000b) and Stocks (2001b). In Stocks (2001b), an interesting new measure of the performance of SSR is briefly described, that of a coding efficiency. This measure is briefly explored in Section 4.4. Furthermore, in both Stocks (2000b) and Stocks (2001b), an analysis of the optimal threshold configuration for the system in Fig. 4.1 is given. Such calculations are the subject of Chapter 8, and will not be considered here.

4.2.3 SSR in Neurons and Cochlear Implant Encoding

The original work on SSR commented briefly—see, in particular, Stocks (2000b)—that the model shown in Fig. 4.1 has similarities to ensembles of sensory neurons. This, combined with the fact that SR is well known to occur in neurons (Levin and Miller 1996), is one of the motivations for studying such a system. In order to further investigate the possibility that SSR can occur in real neurons, Stocks and Mannella (2001)—see also the book chapter, Stocks and Mannella (2000)—replaces the simple threshold device building-block that makes up the system shown in Fig. 4.1, with the FitzHugh-Nagumo neuron model. It is shown that SSR can still occur despite this change in the model. Calculations of the mutual information in such a case are more complicated than the simple model analysed in depth in this chapter, as the output is no longer discretely valued. Hence, investigation of this topic is beyond the scope of this thesis, and the reader is referred to Stocks and Mannella (2001).




Figure 4.1. Array of N noisy threshold devices. This schematic diagram shows the model in which SSR occurs. There are N identical threshold devices. The input signal is taken to be a random signal, consisting of a sequence of discrete time uncorrelated samples drawn from a continuously valued probability distribution. Each device receives the same input signal sample, x , and is subject to independent additive noise, η_i . The noise is also a random signal, and is independent from the input signal, x . The output from the n -th device, y_i , is unity if the sum of the signal and noise at its input is greater than the corresponding threshold, θ_i and zero otherwise. The overall output, y , is the sum of the individual outputs, y_i .

The only other work to consider SSR in neural coding is Hoch *et al.* (2003b), which examines how SSR can occur in leaky integrate-and-fire and Hodgkin-Huxley neural models, with an emphasis on finding the optimal noise level for a large number of neurons.

Cochlear implants¹¹ are prosthetic devices (Dorman and Wilson 2004) that enable profoundly deaf people to hear. The operation of cochlear implants requires direct electrical stimulation of the cochlear nerve. Such stimulation requires sophisticated methods of signal encoding, and although hearing can be restored, patients still have difficulty perceiving speech in a noisy room, or music. An interesting and potentially very important proposed application of both SR in general, and the SSR effect, is to incorporate its effects in the design of cochlear implant encoding. The idea is based

¹¹Historical note: the world's first multi-channel cochlear implant was invented in Australia by Graeme Clarke—the first patient to receive the implant was in Melbourne, 1978 (Clark 1986).

4.2 Literature Review

on the fact that people requiring cochlear implants are missing the natural sensory hair cells that a functioning inner ear uses to encode sound in the cochlear nerve. It is known that these hair cells undergo significant Brownian motion—i.e. randomness (Jaramillo and Wiesenfeld 1998, Bennett *et al.* 2004). Hence, the hypothesis is that employing the principles of SSR to re-introduce this natural randomness to the encoding of sound could improve speech comprehension in patients fitted with cochlear implants (Morse *et al.* 2002, Stocks *et al.* 2002, Allingham *et al.* 2003, Allingham *et al.* 2004, Morse and Stocks 2005). Investigation and further discussion of this work is beyond the scope of this thesis, and the reader is referred to the stated references, as well as to other references proposing the use of SR to improve signal encoding in cochlear implants (Morse and Roper 2000, Morse and Meyer 2000, Hohn 2001, Chatterjee and Robert 2001, Rubinstein and Hong 2003).

4.2.4 Work on SSR by Other Authors

Several other authors besides Stocks and colleagues have now published results on SSR. On the theoretical side, Hoch *et al.* (2003b)—identical work is given in Hoch *et al.* (2003b) and Wenning (2004)—derives an approximation for the mutual information through the system of Fig. 4.1 that applies in the case of large N . Using this result, it is shown that for the case of a Gaussian signal and independent Gaussian noise—as studied in Stocks (2000a) and Stocks (2001a)—the value of noise intensity at which the mutual information reaches its maximum converges to a fixed value as N increases. As mentioned, Hoch *et al.* (2003b) also examine SSR in a neural coding context, including an investigation into the effects of SSR under energy efficiency coding constraints. We consider such a question from a different perspective in Chapter 9.

An analysis of the SSR model in terms of Fisher information is given in Rousseau *et al.* (2003). This paper also discusses for the first time the case of deterministic time-varying input signals, rather than random input signals. Such an input signal allows analysis of the performance of the SSR model in terms of SNR in the conventional SR sense. A detailed investigation of exactly this topic is reported in Rousseau and Chapeau-Blondeau (2004). Subsequently, Wang and Wu (2005) also use Fisher information to measure the performance of the SSR model, by examining how the ‘thickness’ of the ‘tail’ of the noise distribution affects the output response.

The SSR model shown in Fig. 4.1 is mentioned in Sato *et al.* (2004), however this work considers the input signal to always be entirely subthreshold. In such a case SR can still occur, however in this chapter, unlike Sato *et al.* (2004), we will not restrict attention to entirely subthreshold input signals.

More practically oriented work includes an investigation of the potential use of SSR in motion detection systems (Harmer 2001, Harmer and Abbott 2001, McDonnell and Abbott 2002). There has also been a proposal to use the entire SSR model as the comparator component of a sigma-delta modulator (Oliaei 2003), which is a specific type of analog-to-digital converter that makes use of feedback.

4.2.5 Similarity to SSR in Other Work

As already mentioned, the model shown in Fig. 4.1 can also serve as a model for a population of neurons. Furthermore, we will see in Chapters 6-8 that it also serves as a model for quantisation and analog-to-digital conversion (ADC). However, in quantisers and ADCs it is certainly not conventional to assume that all thresholds have the same value, as is the case in this chapter. There are however, some other systems in which all thresholds are identical, and hence, the SSR model shares similarities with work previously considered in the literature. The remainder of this subsection briefly points out references to such work. Finally, it discusses how the results on SSR first given in Stocks (2000a) differs from these works.

DIMUS sonar arrays

DIMUS (Digital Multibeam Steering) sonar arrays (Rudnick 1960) are arrays of spatially distributed hydrophones that each ‘clip’ their inputs to provide a digital binary signal. If there are N hydrophones, then the overall output is the sum of the N individual binary signals. Such a method for acquiring acoustic sonar signals was first used prior to the 1960s, and provided the introduction of digital signal processing to the field. Due to the spatially distributed nature of the hydrophones, it is usually assumed that noise at the input to each hydrophone is mutually independent of every other hydrophone, while the signal at each remains identical. As noted in Stocks (2001a), such a situation is virtually identical to the SSR model. However, unlike the SSR model, DIMUS has never been analysed from an information theoretic perspective, and no results have been published indicating the performance of the DIMUS sonar array as

4.2 Literature Review

the noise intensity increases from zero to input SNRs smaller than zero dB (Anderson 1960, Rudnick 1960, Wolff *et al.* 1962, Remley 1966, Kanefsky 1966, Berndt 1968, Fitelson 1970, Wang 1972, Bershada and Feintuch 1974, Wang 1976, Anderson 1980, Tuteur and Presley 1981). The interested reader is referred to the cited references.

Semi-continuous channels

Although it is not immediately apparent on a first reading of a paper by Chang and Davisson (1988), which considers algorithms for calculating channel capacity for “an infinite input and finite output channel,” it actually is highly relevant to the SSR scenario. Both the channel considered by Chang and Davisson (1988)—see also Davisson and Leon-Garcia (1980)—and the SSR model of Fig. 4.1 have an input signal consisting of a continuously valued random signal, and an output signal that consists of a finite number of states. As will be discussed in Section 4.4, calculations of channel capacity values for various N given in Chang and Davisson (1988), are in fact equivalent to channel capacity values for SSR, under certain assumptions on the noise distribution.

Detection scenarios

A summation of independent binary values is also used in Kay (2000) in a detection scenario. This work describes the occurrence of SR in a suboptimal detector. However, unlike SSR, the input signal is only ever a single value, and is not considered to be a random variable.

4.2.6 Discussion

All of the situations above—apart from Chang and Davisson (1988)—are considered from the viewpoint of signal detection, and the SNR of the output signal. As in Chang and Davisson (1988), another way of describing such systems is in terms of information theory. In each case, the system in which a signal is propagated to its detected output state can be described as a *channel*. An important branch of information theory is concerned with the transmission of information through channels. Often, it is assumed that the signal is discrete-time in nature. Often, the basis for this assumption is that the input signal can be described as a random signal, and has a certain finite bandwidth. Under such an assumption, the sampling theorem guarantees that a continuous time signal can be sampled at the Nyquist rate. Subsequently, the original continuous time

signal can be perfectly reconstructed from those samples. In such a situation, or indeed any case where a signal is composed of *iid*—independent and identically distributed—samples from some probability distribution, Shannon’s mutual information measure provides a measure of the *average* number of bits per sample that can be transmitted by a single sample in some channel. When combined with a sample rate, measured in samples per second—often called bandwidth—an *information rate* through a channel can be stated, by multiplying the sample rate by the mutual information, to obtain the number of bits per second that are transmitted through the channel.

The remainder of this Chapter studies the system in Fig. 4.1 from this information theoretic viewpoint. All input signals are assumed to be discrete-time sequences of samples drawn from some stationary probability distribution. Thus, this work differs from the detection scenario of Kay (2000), which considers a constant signal. Such a signal does not convey new information with an increasing number of samples, and cannot be considered from an information theoretic viewpoint.

Finally, note that another branch of information theory is concerned with the problem of *lossy source coding*. This includes studies of quantisation theory, where a continuously valued signal is compressed to a discretely valued encoding. Although combined source-channel coding research is increasing, information theory traditionally separates source coding and channel coding into two separate independent components of a communications system. An interesting feature of the SSR model is that it can be interpreted both as a channel model and as a source coding model. This is because the channel is semi-continuous; it has a continuously valued input signal, but a discretely valued output signal. This viewpoint leads to a natural interpretation of the SSR model as a quantisation scheme. This subject is further studied in Chapters 6 and 7. On the other hand, since there is channel noise in the system, and the output is a noisy version of the input signal, the system can also naturally be seen as a channel. The remainder of this chapter looks at this topic, and analyses SSR from a channel coding viewpoint.

4.3 Suprathreshold Stochastic Resonance

4.3.1 The SSR Model

Fig. 4.1 shows a schematic diagram of the model system in which SSR occurs. This system consists of N threshold devices—which we will also refer to as comparators—which all receive the same sample of a random input signal, x . This random signal is assumed to consist of a sequence of independent samples drawn from a distribution with Probability Density Function (PDF), $P(x)$. For such a situation to apply in a real system, the independence of each sample usually requires that a random continuous time signal is bandlimited and sampled at the Nyquist rate (Proakis and Salehi 1994), prior to being input to this system as a sequence of discrete-time samples.

The i -th device in the model is subject to continuously valued *iid*—independent and identically distributed—additive noise, η_i ($i = 1, \dots, N$), drawn from a probability distribution with PDF $R_\eta(\eta)$. Each noise signal is required to also be independent of the signal, x . For each individual comparator, the output signal, y_i , is unity if the input signal, x , plus the noise on that comparators thresholds, η_i , is greater than the threshold value, θ_i . The output signal is zero otherwise. The outputs from each comparator, y_i , are summed to give the overall system output signal, y . Hence, y is a discrete signal which can take integer values between 0 and N .

The output of device i is then given by

$$y_i = \begin{cases} 1 & \text{if } x + \eta_i \geq \theta_i, \\ 0 & \text{otherwise.} \end{cases} \quad (4.1)$$

The overall output of the array of comparators is $y = \sum_{i=1}^N y_i$. This can be expressed as a function of x in terms of the signum (sign) function as

$$y(x) = \frac{1}{2} \sum_{i=1}^N \text{sign}[x + \eta_i - \theta_i] + \frac{N}{2}. \quad (4.2)$$

As mentioned in Section 4.2.5, Kay (2000) has previously published work describing the occurrence of SR in a suboptimal detector. This work uses an expression similar to Eqn. (4.2) in terms of the signum function, indicating some similarities between SSR and detection problems. However, the problem of deciding whether a constant signal has been detected, as carried out in Kay (2000), is beyond the scope of this thesis; the main task considered here is that of signal transmission and quantisation, rather than detection.

Note that Eqn. (4.2) completely describes the *transfer function* for the model system of Fig. 4.1. However, the output of the array of comparators is non-deterministic—except in the complete absence of noise—and represents a lossy encoding of the input signal. Hence, the transfer function is not deterministic, since for a given input value, x , the output, $y(x)$, depends on the set of random variables, $\{\eta_i\}$, $i = 1, \dots, N$.

This chapter considers only the case of all threshold values being identical. Hence, we let $\theta_i = \theta$ for all i and, without loss of generality, in this Chapter the subscript, i , is dropped from all references to threshold values. It is also dropped for references to noise signals, η , but this does not imply that the noise on each threshold is no longer independent.

Since the encoding given by the transfer function of Eqn. (4.2) is not deterministic, probabilistic measures are required for mathematical analysis of the system. The key function required is therefore the joint PDF between the input and output signals, $P(x, y)$. Denoting the probability mass function of the output signal as $P_y(n)$ and making use of Bayes' theorem (Shiryaev 1996, Yates and Goodman 2005) gives the joint PDF as

$$P(x, y) = P(y = n|x)P(x) \quad (4.3)$$

$$= P(x|y = n)P_y(n). \quad (4.4)$$

We will describe the conditional distribution of the output given the input, denoted by $P(y = n|x)$, as the *transition probabilities*, since $P(y = n|x)$ gives the probability that the encoding for a given input value, x , is encoded by output state n . From here on we will abbreviate the notation $P(y = n|x)$ to $P(n|x)$. The transition probabilities can be used to obtain $P_y(n)$ as

$$P_y(n) = \int_{-\infty}^{\infty} P(n|x)P(x)dx \quad n \in 0, \dots, N. \quad (4.5)$$

In this Chapter we will always assume that the PDF, $P(x)$, is known, except briefly in Section 4.4, where channel capacity is considered. To progress further requires a method for calculating the transition probabilities. Using the notation of Stocks (2000a), let $P_{1|x}$ be the probability of any comparator being 'on'—that is, the sum of the signal, x , and noise, η , exceeds the threshold value, θ , for a single comparator—given that the input signal value, x , is known. Thus

$$P_{1|x} = P(x + \eta > \theta|x). \quad (4.6)$$

4.3 Suprathreshold Stochastic Resonance

This probability depends on the noise PDF, $R(\eta)$, as

$$P_{1|x} = \int_{\theta-x}^{\infty} R(\eta) d\eta = 1 - F_R(\theta - x), \quad (4.7)$$

where $F_R(\cdot)$ is the Cumulative Distribution Function (CDF) of the noise. If $R(\eta)$ is an even function of η then

$$P_{1|x} = F_R(x - \theta). \quad (4.8)$$

Since all thresholds have the same value, then—as noted in Stocks (2000a)—the transition probabilities are given by the binomial distribution (Shiryaev 1996) as

$$P(n|x) = \binom{N}{n} P_{1|x}^n (1 - P_{1|x})^{N-n} \quad n \in 0, \dots, N. \quad (4.9)$$

In many applications of the binomial distribution, the probability of a single event is a constant. Instead, here $P_{1|x}$ is a function of x and the transition probabilities are binomially distributed for a given value of x .

However, for a given value of n , the nature of the transition probabilities for a given n as a function of x is not as easily characterised. We can however easily find the value of x at which the maximum value, or peaks, of $P(n|x)$ for a given n occurs, by differentiating $P(n|x)$ with respect to x , setting to zero and solving for x . The details of this procedure are given in Section B.1 of Appendix B, which shows that the peak occurs when

$$P_{1|x} = \begin{cases} 1 & \text{for } n = 0, \\ \frac{n}{N} & \text{for } n \in 1, \dots, N - 1, \\ 0 & \text{for } n = N. \end{cases} \quad (4.10)$$

In all cases we examine, the CDF will be sigmoidal rather than peaked for $n = 0$ and $n = N$, which means that the peaks occur at $\pm\infty$. Hence, for the context of determining the peaks, we will ignore these values of n in this discussion. Substituting for Eqn. (4.7) in Eqn. (4.10) and solving for x gives

$$x = \theta + F_R^{-1} \left(1 - \frac{n}{N} \right) \quad \text{for } n \in 1, \dots, N - 1, \quad (4.11)$$

where $F_R^{-1}(\cdot)$ is the Inverse Cumulative Distribution Function (ICDF) of the noise distribution. For even noise PDFs,

$$x = \theta + F_R^{-1} \left(\frac{n}{N} \right) \quad \text{for } n \in 1, \dots, N - 1. \quad (4.12)$$

These values of x at which the maximum value of $P(n|x)$ occurs are known as the *mode* of the conditioned random variable, y given x (Yates and Goodman 2005).

4.3.2 Measuring SSR with Mutual Information

We will refer to the model described in the above subsection as the *SSR model*.

The mutual information—see Section 2.3.1 in Chapter 2 for a definition—between the input and output signals of the SSR model can be expressed as

$$\begin{aligned} I(x, y) &= H(y) - H(y|x) \\ &= - \sum_{n=0}^N P_y(n) \log_2 P_y(n) - \left(- \int_{-\infty}^{\infty} P(x) \sum_{n=0}^N P(n|x) \log_2 P(n|x) dx \right). \end{aligned} \quad (4.13)$$

As noted in Stocks (2001a), since the input signal is continuously valued and the output is discretely valued, the mutual information is that of a semi-continuous channel. Since $P_y(n)$ is a function of $P(x)$ and $P(n|x)$ —see Eqn. (4.5)—the mutual information can be expressed in terms of only $P(x)$ and $P(n|x)$. The transition probabilities, $P(n|x)$, are given by Eqn. (4.9), which for a given N depends only on $P_{1|x}$ and therefore only on the noise PDF, $R(\eta)$, and the threshold values, θ .

Stocks (2000a) simplifies this expression for the mutual information by use of Eqn. (4.9). In particular, consider the entropy of the output for a given value of the input,

$$\hat{H}(y|x) = - \sum_{n=0}^N P(n|x) \log_2 P(n|x). \quad (4.14)$$

Note that the quantity, \hat{H} , is labelled with a caret to make explicit that this formula is not the *average* conditional entropy, $H(y|x)$. Substituting Eqn. (4.9) into Eqn. (4.14) and simplifying gives

$$\hat{H}(y|x) = -N \left(P_{1|x} \log_2 P_{1|x} + (1 - P_{1|x}) \log_2 (1 - P_{1|x}) \right) - \sum_{n=0}^N P(n|x) \log_2 \binom{N}{n}. \quad (4.15)$$

This simplification uses the fact that $\sum_{n=0}^N P(n|x) = 1$ and that the expected value of the binomial distribution is $\sum_{n=0}^N nP(n|x) = NP_{1|x}$. Eqn. (4.15) is identical to one expressed by Davisson and Leon-Garcia (1980) and Chang and Davisson (1990), although their use of this equation is under far more specific conditions.

The average conditional entropy, or equivocation, is

$$\begin{aligned} H(y|x) &= \int_x P(x) \hat{H}(y|x) dx \\ &= - \sum_{n=0}^N P_y(n) \log_2 \binom{N}{n} \\ &\quad - N \int_x P(x) \left(P_{1|x} \log_2 P_{1|x} + (1 - P_{1|x}) \log_2 (1 - P_{1|x}) \right) dx, \end{aligned} \quad (4.16)$$

4.3 Suprathreshold Stochastic Resonance

where the simplification comes from use of Eqn. (4.5). Notice that the integral on the RHS is N times the conditional entropy of the output of a single comparator. Hence, since $\sum_{n=0}^N P_y(n) \log_2 \binom{N}{n}$ is always smaller than zero, the overall average conditional entropy is smaller than the sum of the average conditional entropy of its parts, by this factor. This means that the average uncertainty about the input signal, x , is reduced by observing the added outputs of more than one single threshold device.

Note that in the absence of noise, the system is deterministic, and $P_{1|x}$ can only be zero or unity. This means that the only possible overall output states are 0 or N , and the conditional output entropy is always zero, since there is never any uncertainty about the output given the input. Thus, the mutual information is always exactly one bit per sample, provided the threshold is set to the signal mean, so that $P_y(0) = P_y(N) = 0.5$.

Substituting Eqn. (4.16) into Eqn. (4.13) gives the mutual information as

$$I(x, y) = - \sum_{n=0}^N P_y(n) \log_2 P^*(n) + N \int_x P(x) \left(P_{1|x} \log_2 P_{1|x} + (1 - P_{1|x}) \log_2 (1 - P_{1|x}) \right) dx, \quad (4.17)$$

where $P^*(n) = P_y(n) / \binom{N}{n}$ and thus $P^*(n) = \int_x P(x) P_{1|x}^n (1 - P_{1|x})^{N-n} dx$. The above derivation is given in Stocks (2001a).

Change of probability measure

We now introduce a mathematically convenient change of probability measure. The function $P_{1|x}$ is a function of x . The range of x is the range of allowable values of the input, and for infinite support input PDFs, $x \in [-\infty, \infty]$ —an example of this is the Gaussian distribution. We will see that certain equations can be simplified if integrations are performed over the interval $[0, 1]$ rather than $[-\infty, \infty]$. This can be achieved by a change of probability measure as follows.

Let $\tau = P_{1|x}$ so that from Eqn. (4.7), $\tau = 1 - F_R(\theta - x)$. Then, differentiating τ with respect to x —assuming τ is differentiable for all x —gives

$$\frac{d\tau}{dx} = R(\theta - x). \quad (4.18)$$

Multiplying both sides by $P(x)$ and rearranging leaves

$$P(x)dx = \frac{P(x)}{R(\theta - x)} d\tau. \quad (4.19)$$

Substituting for $x = \theta - F_R^{-1}(1 - \tau)$ gives

$$P(x)dx = \frac{P(\theta - F_R^{-1}(1 - \tau))}{R(F_R^{-1}(1 - \tau))}d\tau. \quad (4.20)$$

The LHS of the above expression is the infinitesimally small probability that the input is x . The RHS is $d\tau$ times the ratio of the signal PDF to the noise PDF, calculated at the values x and $(\theta - x)$ respectively. Let this RHS expression be $Q(\tau)d\tau$ so that

$$Q(\tau) = \frac{P(x)}{R(\theta - x)} \Big|_{x=\theta - F_R^{-1}(1 - \tau)}. \quad (4.21)$$

Thus, provided the support of $P(x)$ is contained in the support of $R(\theta - x)$, since otherwise division by zero occurs, then $Q(\tau)$ is a PDF defined on the interval $[0, 1]$. We will see later in this Chapter that the case of uniform signal and uniform noise, with the signal variance larger than the noise variance, is an example where $Q(\tau)$ is not a PDF.

Thus, making a change of variable in Eqn. (4.16) gives

$$\begin{aligned} H(y|x) &= - \sum_{n=0}^N P_y(n) \log_2 \binom{N}{n} \\ &\quad - N \int_{\tau=0}^{\tau=1} Q(\tau) (\tau \log_2 \tau d\tau + (1 - \tau) \log_2 (1 - \tau)) d\tau, \end{aligned} \quad (4.22)$$

and in Eqn. (4.17) gives

$$\begin{aligned} I(x, y) &= - \sum_{n=0}^N P_y(n) \log_2 P^*(n) \\ &\quad + N \int_{\tau=0}^{\tau=1} Q(\tau) (\tau \log_2 \tau d\tau + (1 - \tau) \log_2 (1 - \tau)) d\tau \end{aligned} \quad (4.23)$$

where

$$P^*(n) = \int_{\tau=0}^{\tau=1} Q(\tau) \tau^n (1 - \tau)^{N-n} d\tau. \quad (4.24)$$

Input and noise PDFs that are even functions about identical means

Further simplification can be made in the case where both the signal and noise PDFs are even valued functions, with identical means, and all thresholds are set equal to that mean. Without loss of generality, assume the mean is zero. Hence if $R(\eta)$ is even and $\theta = 0$, from Eqn. (4.8), $P_{1|x} = F_R(x) = 1 - P_{1|-x}$. Considering the integral

$$A = \int_{x=-\infty}^{x=\infty} P(x)(1 - P_{1|x}) \log_2 (1 - P_{1|x}) dx, \quad (4.25)$$

4.3 Suprathreshold Stochastic Resonance

and putting $s = -x$, then making a change of variable and letting $P(x)$ be even so that $P(x) = P(-x)$, gives

$$\begin{aligned} A &= - \int_{s=-\infty}^{s=\infty} P(-s)(1 - P_{1|-s}) \log_2 (1 - P_{1|-s}) ds \\ &= \int_{s=-\infty}^{s=\infty} P(s)P_{1|s} \log_2 P_{1|s} ds. \end{aligned} \quad (4.26)$$

Thus Eqn. (4.16) reduces to

$$H(y|x) = - \sum_{n=0}^N P_y(n) \log_2 \binom{N}{n} - 2N \int_x P(x)P_{1|x} \log_2 P_{1|x} dx, \quad (4.27)$$

and Eqn. (4.17) reduces to

$$I(x, y) = - \sum_{n=0}^N P_y(n) \log_2 P^*(n) + 2N \int_x P(x)P_{1|x} \log_2 P_{1|x} dx. \quad (4.28)$$

Also, $Q(\tau)$ simplifies when $\theta = 0$ and $P(x)$ and $R(\eta)$ are even, to yield

$$Q(\tau) = \frac{P(F_R^{-1}(1 - \tau))}{R(F_R^{-1}(1 - \tau))} = \frac{P(F_R^{-1}(\tau))}{R(F_R^{-1}(\tau))}. \quad (4.29)$$

It turns out that besides being mathematically convenient, the function $Q(\tau)$ has further significance, and assists greatly in intuitively understanding the behaviour of SSR. We shall return to this point in Section 4.3.3 and Chapter 5.

The average conditional entropy can now be written in terms of the function $Q(\tau)$ as

$$H(y|x) = - \sum_{n=0}^N P_y(n) \log_2 \binom{N}{n} - 2N \int_{\tau=0}^{\tau=1} Q(\tau)\tau \log_2 \tau d\tau, \quad (4.30)$$

and the mutual information as

$$I(x, y) = - \sum_{n=0}^N P_y(n) \log_2 P^*(n) + 2N \int_{\tau=0}^{\tau=1} Q(\tau)\tau \log_2 \tau d\tau. \quad (4.31)$$

Identical signal and noise distributions

Further analytical simplification of Eqn. (4.13) is possible in the case where the signal's PDF is 'matched' with the noise PDF, so that $P(x) = R(\theta - x) \forall x$. This means that the signal and noise must have identical even moments.

Note that for some signal and noise distribution pairs, it is not possible for $P(x) = R(\theta - x) \forall x$. Examples include (i) if $P(x)$ and $R(\eta)$ are uniform, and $E[x] \neq \theta - E[\eta]$

and (ii) if both $P(x)$ and $R(\eta)$ have only positive support, with infinite tails, that is, $P(x) > 0 \forall x \in [0, \infty]$.

We begin with a result given in Stocks (2001a) showing that all output states are equiprobable in this situation, that is

$$P_y(n) = \frac{1}{N+1} \quad n = 0, \dots, N. \quad (4.32)$$

This result can be proved as follows. Firstly, consider $Q(\tau)$ as given by Eqn. (4.21). If $P(x) = R(\theta - x) \forall x$ then it is clear that $Q(\tau) = 1 \forall \tau$. Hence, from Eqn. (4.24),

$$P^*(n) = \int_{\tau=0}^{\tau=1} \tau^n (1-\tau)^{N-n} d\tau. \quad (4.33)$$

This is simply a Beta function (Spiegel and Liu 1999), which evaluates to

$$P^*(n) = \frac{n!(N-n)!}{(N+1)!}, \quad (4.34)$$

and thus,

$$P_y(n) = \binom{N}{n} P^*(n) = \frac{1}{N+1}. \quad (4.35)$$

Interestingly, this means that the output entropy is maximised, since, as noted in Stocks (2001a), it is well known that maximum entropy occurs when all states are equally likely (Cover and Thomas 1991). Thus, the output entropy is

$$H(y) = - \sum_{n=0}^N \frac{1}{N+1} \log_2 \left(\frac{1}{N+1} \right) = \log_2 (N+1). \quad (4.36)$$

This does not however mean that mutual information is maximised, since it depends also on the average conditional entropy, $H(y|x)$. It is possible for a discrete signal like y to have a uniform distribution, yet still have a very large average uncertainty about y given x , whereas a smaller output entropy might provide a far smaller $H(y|x)$.

Stocks (2001a) calculates the integral on the RHS of Eqn. (4.16), to obtain an exact expression for the average conditional entropy. This integral can be solved equivalently from Eqn. (4.22), by using the fact that $Q(\tau) = 1$ to get

$$\int_{\tau=0}^{\tau=1} Q(\tau) (\tau \log_2 \tau d\tau + (1-\tau) \log_2 (1-\tau)) d\tau = 2 \int_{\tau=0}^{\tau=1} \tau \log_2 \tau d\tau = -\frac{1}{2 \ln 2}. \quad (4.37)$$

In order to further simplify the average conditional entropy, Stocks (2001a) uses the identity (see Section B.2 in Appendix B for a proof):

$$- \sum_{n=0}^N \log_2 \binom{N}{n} = \sum_{n=1}^N (N+1-2n) \log_2 n. \quad (4.38)$$

4.3 Suprathreshold Stochastic Resonance

When Eqns. (4.37) and (4.38) are substituted into Eqn. (4.22) we get

$$H(y|x) = \frac{N}{2 \ln 2} + \frac{1}{N+1} \sum_{n=1}^N (N+1-2n) \log_2 n. \quad (4.39)$$

Hence, the mutual information is

$$I(x, y) = \log_2 (N+1) - \frac{N}{2 \ln 2} - \frac{1}{N+1} \sum_{n=2}^N (N+1-2n) \log_2 n. \quad (4.40)$$

This equation is plotted against increasing N in Fig. 4.2. What is quite remarkable about this result is that it is independent of the nature of the PDFs of the signal and noise, other than that $P(x) = R(\theta - x)$ for all valid values of x , so that both PDFs have the same shape, but may possibly have different means, and be mutually reversed along the x -axis about their means.




Figure 4.2. Mutual information for the case of $P(x) = R(\theta - x)$. This plot shows the mutual information against an increasing number of array elements, $N = 0, \dots, 511$, for the case of $P(x) = R(\theta - x)$, calculated from Eqn. (4.40). Notice that the mutual information increases with increasing N , although the rate of increase decreases with increasing N .

Further simplification of Eqn. (4.40) is insightful. Firstly, the second term of Eqn. (4.39) can be expressed as

$$\begin{aligned} \frac{1}{N+1} \sum_{n=1}^N (N+1-2n) \log_2 n &= \log_2 (N!) - \frac{2}{N+1} \sum_{n=1}^N n \log_2 n \\ &= \log_2 (N!) - N \log_2 (N+1) \\ &\quad - 2 \sum_{n=1}^N \frac{n}{N+1} \log_2 \left(\frac{n}{N+1} \right). \end{aligned} \quad (4.41)$$

Thus, the average conditional entropy can be written as

$$H(y|x) = \log_2(N!) - N \log_2(N+1) + \frac{N}{2 \ln 2} - 2 \sum_{n=1}^N \frac{n}{N+1} \log_2\left(\frac{n}{N+1}\right). \quad (4.42)$$

Note that the first two terms scale with $N \log N$, since $\log_2(N!)$ approaches $N \log N$ for very large N . Such large N behaviour is discussed in Chapter 5. The term inside the summation in the final term is always between 0 and 0.5, so therefore the summation is always smaller than N . Hence, the average conditional entropy consists of two terms of the order of $N \log N$ and two terms of the order of N .

Stocks (2001a) also shows that as N becomes large, Eqn. (4.42) reduces to $I(x, y) \simeq 0.5 \log_2(N+1)$, and that this implies that a maximum must occur in the mutual information for a nonzero noise intensity for any $N > 1$. The large N behaviour of SSR is further examined in Chapter 5. Additional analytical progress on calculating the mutual information for small N can only be made if the signal and/or noise distributions are specified. This is the focus of the next subsection.

4.3.3 SSR for Specific Signal and Noise Distributions

This subsection examines whether SSR can occur for a range of different signal and noise distributions. Before doing this, we firstly describe seven different probability distributions, and also make some comments on the function, $Q(\tau)$.

Seven different probability distributions

Table 4.1 gives expressions for the PDF, CDF and ICDF of a number of different probability distributions. Section B.3 in Appendix B gives specific details of employing each of these distributions as the signal or noise distribution. The PDF, CDF and ICDF of each distribution are plotted in Fig. 4.3, where the variance is $\sigma = 1$ for all distributions except Cauchy, and the parameter $\lambda = 1$ for the Cauchy distribution. For the CDF, z is limited to the values of support of the corresponding PDF, and for the ICDF $w \in [0, 1]$. Note that $P_{1|x}$ can be obtained from the CDF by use of Eqn. (4.7) and the mode of $P(n|x)$ for each n by use of Eqn. (4.11).

The first five PDFs can be seen to be even functions about their means. The Gaussian and logistic PDFs are almost identical in shape, and have infinite support, whereas the uniform PDF has finite support, and is flat. The Laplacian PDF has infinite support, but

4.3 Suprathreshold Stochastic Resonance

is more sharply peaked than either the Gaussian or logistic distributions. The Cauchy distribution—also known as the Lorentzian distribution—is unusual, in that it has an infinite variance, yet is still a PDF. Instead of characterising its width by its standard deviation, a property known as its Full Width at Half Maximum (FWHM) is usually used, as explained in Section B.3 of Appendix B. The FWHM is the distance between points on the curve at which the function reaches half its maximum value. This is given for the Cauchy distribution by the parameter, λ_x . The Rayleigh and exponential distributions exist only for positive support. This fact ensures that no PDF, $Q(\tau)$, can be found for these cases for all τ , since there will always exist points of support for which $P(x)$ is nonzero, yet $R(\theta - x) = 0$.

For the CDFs, in each case there is a substantial range of z for which the CDF is approximately linearly increasing. The support of the CDF is the same as the support for the corresponding PDF. The ICDF can be seen to have support between zero and unity. This reflects the fact that probabilities must be between zero and unity.

Observations about $Q(\tau)$

It will prove to be convenient to parameterise the forthcoming results in terms of the ratio of noise standard deviation to signal standard deviation, $\sigma = \sigma_\eta / \sigma_x$. Likewise, for the Cauchy distribution, let $\sigma_\lambda = \lambda_\eta / \lambda_x$. Note from Appendix B.3 that the variance of the noise is always a function of σ_η , and the variance of the signal is a function of σ_x . The reason that this parameterisation is convenient is that the mutual information turns out—in the cases examined here, where θ is equal to the signal mean—to be a function of the ratio, σ , so that it is invariant to a change in σ_x provided σ_η changes by the same proportion.

This result can be seen by deriving the function $Q(\tau)$ for specific signal and noise distributions cases. For the case of signal and noise sharing the same distribution—but with not necessarily equal variances—the function $Q(\tau)$ is listed in Table 4.2 for $\theta = 0$. In all cases, $Q(\tau)$ is a function of the ratio σ . Hence, by inspection of Eqns. (4.24), (4.22) and (4.23) it is clear that $P_y(n)$, $H(y)$, $H(y|x)$, and $I(x, y)$ will also be functions of σ , and not σ_η or σ_x in isolation. Note that if $\theta \neq 0$ then $Q(\tau)$ will depend on the ratio $\frac{\theta}{\sigma_x}$, as well as θ and the ratio, σ , and therefore so will the mutual information.

Table 4.1. List of probability distributions and corresponding PDFs.

Distribution	support	mean	variance	PDF, $P(x)$	CDF, $F_R(z)$	ICDF, $F_R^{-1}(w)$
Gaussian	$[-\infty, \infty]$	0	σ_x^2	$\frac{1}{\sqrt{2\pi\sigma_x^2}} \exp\left(-\frac{x^2}{2\sigma_x^2}\right)$	$0.5 + 0.5\text{erf}\left(\frac{z}{\sqrt{2}\sigma_\eta}\right)$	$\sqrt{2}\sigma_\eta\text{erf}^{-1}(2w - 1)$
Uniform	$[-\frac{\sigma_x}{2}, \frac{\sigma_x}{2}]$	0	$\frac{\sigma_x^2}{12}$	$\begin{cases} \frac{1}{\sigma_x}, & x \in [-\frac{\sigma_x}{2}, \frac{\sigma_x}{2}], \\ 0, & \text{otherwise} \end{cases}$	$\begin{cases} 0, & z < -\frac{\sigma_\eta}{2}, \\ \frac{z}{\sigma_\eta} + \frac{1}{2}, & z \in [-\frac{\sigma_\eta}{2}, \frac{\sigma_\eta}{2}], \\ 1, & z > \frac{\sigma_\eta}{2}. \end{cases}$	$\sigma_\eta(w - \frac{1}{2})$
Laplacian	$[-\infty, \infty]$	0	σ_x^2	$\frac{1}{\sqrt{2}\sigma_x} \exp\left(-\frac{\sqrt{2} x }{\sigma_x}\right)$	$\begin{cases} \frac{1}{2} \exp\left(\frac{\sqrt{2}z}{\sigma_\eta}\right), & z \leq 0, \\ 1 - \frac{1}{2} \exp\left(\frac{-\sqrt{2}z}{\sigma_\eta}\right), & z \geq 0. \end{cases}$	$\begin{cases} \frac{\sigma_\eta}{\sqrt{2}} \ln(2w), & w \in [0, \frac{1}{2}], \\ \frac{-\sigma_\eta}{\sqrt{2}} \ln(2(1-w)), & w \in [\frac{1}{2}, 1]. \end{cases}$
Logistic	$[-\infty, \infty]$	0	$\sigma_x^2 = \frac{\pi^2 b_x^2}{3}$	$\frac{\exp\left(-\frac{x}{b_x}\right)}{b_x(1+\exp\left(-\frac{x}{b_x}\right))^2}$	$\frac{1}{1+\exp\left(-\frac{z}{b_\eta}\right)}$	$b_\eta \ln\left(\frac{w}{1-w}\right)$
Cauchy	$[-\infty, \infty]$	0	∞	$\frac{\lambda_x}{\pi} \frac{1}{\lambda_x^2 + x^2}$	$\frac{1}{2} + \frac{1}{\pi} \arctan\left(\frac{z}{\lambda_\eta}\right)$	$\lambda_\eta \tan\left(\pi(w - \frac{1}{2})\right)$
Exponential	$[0, \infty]$	σ_x	σ_x^2	$\frac{1}{\sigma_x} \exp\left(-\frac{x}{\sigma_x}\right)$	$1 - \exp\left(-z/\sigma_\eta\right)$	$\sigma_\eta \ln\left(\frac{1}{1-w}\right)$
Rayleigh	$[0, \infty]$	$\sigma_x \sqrt{\frac{\pi}{2}}$	$(2 - \frac{\pi}{2})\sigma_x^2$	$\frac{x}{\sigma_x^2} \exp\left(-\frac{x^2}{2\sigma_x^2}\right)$	$1 - \exp\left(-\frac{z^2}{2\sigma_\eta^2}\right)$	$\sigma_\eta \sqrt{2 \ln\left(\frac{1}{1-w}\right)}$

4.3 Suprathreshold Stochastic Resonance

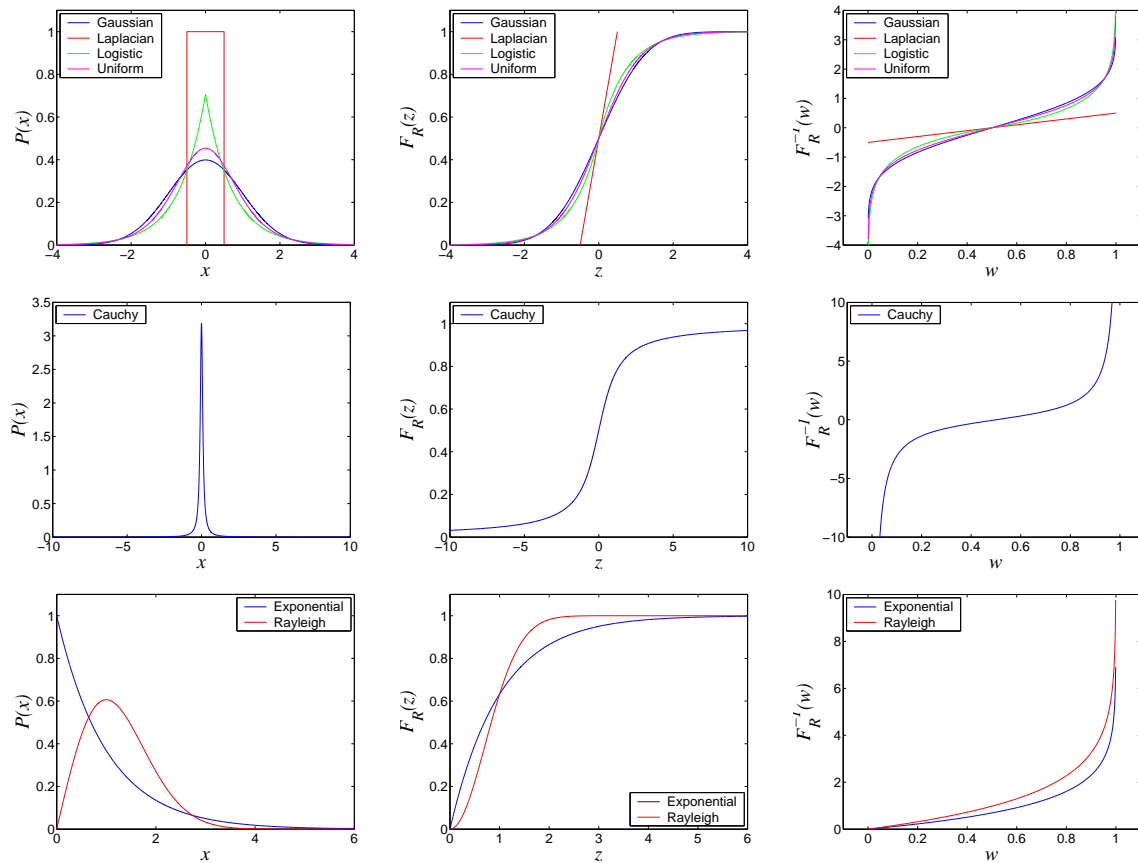


Figure 4.3. Various probability distributions. Plot showing the PDF, $P(x)$, the CDF, $F_R(z)$, and the ICDF, $F_R^{-1}(w)$, for seven different distributions. The left hand panels show the PDFs, the middle panels show the CDFs and the right hand panels the ICDFs. The first row shows the four double-sided, zero-mean, finite variance distributions, the second row shows the infinite variance Cauchy distribution, and the third row shows the two finite variance one-sided PDFs. All distributions have been plotted with a variance of $\sigma = 1$, except for the Cauchy distribution, which has been plotted with a FWHM of $\lambda = 1$.

Table 4.2. List of ratios of signal PDF to noise PDF, $Q(\tau)$, for $\theta = 0$. This table also shows $H(\tau)$ —that is, the entropy of $Q(\tau)$. The label ‘NAS’ indicates that there is no analytical solution for the entropy.

Distribution	$Q(\tau)$	$H(\tau)$
Gaussian	$\sigma \exp \left((1 - \sigma^2) \left(\text{erf}^{-1}(2\tau - 1) \right)^2 \right)$	$-\log_2(\sigma) - \frac{1}{2 \ln 2} \left(\frac{1}{\sigma^2} - 1 \right)$
Uniform, $\sigma \geq 1$	$\begin{cases} \sigma, & -\frac{1}{2\sigma} + 0.5 \leq \tau \leq \frac{1}{2\sigma} + 0.5, \\ 0, & \text{otherwise.} \end{cases}$	$\log_2 \sigma$
Laplacian	$\begin{cases} \sigma(2\tau)^{(\sigma-1)} & \text{for } 0 \leq \tau \leq 0.5, \\ \sigma(2(1-\tau))^{(\sigma-1)} & \text{for } 0.5 \leq \tau \leq 1. \end{cases}$	NAS
Logistic	$\sigma \frac{(\tau(1-\tau))^{(\sigma-1)}}{(\tau^\sigma + (1-\tau)^\sigma)^2}$	NAS
Cauchy	$\sigma \lambda \frac{1 + \tan^2(\pi(\tau-0.5))}{(1 + \sigma^2 \tan^2(\pi(\tau-0.5)))}$	NAS

The PDFs, $Q(\tau)$, are also plotted for various values of σ in Fig. 4.4 for $\theta = 0$. Table 4.2 also lists the entropy of τ —which is equivalent to the negative of the relative entropy between $P(x)$ and $R(x)$ —for two cases. In this table, NAS indicates that No Analytical Solution can be found for the entropy.

We are now in a position to comment further on interpreting the significance of the function $Q(\tau)$. Recall that $P(x)dx = Q(\tau)d\tau$. Assume that $Q(\tau)$ is a PDF—which will be the case if the support of $P(x)$ is contained in the support of $R(\theta - x)$, and is given by Eqn. (4.21). Thus, the entropy of the random variable τ is

$$H(\tau) = - \int_{\tau=0}^{\tau=1} Q(\tau) \log_2(Q(\tau)) d\tau = - \int_{x=-\infty}^{x=\infty} P(x) \log_2 \left(\frac{P(x)}{R(\theta - x)} \right) dx, \quad (4.43)$$

which is the negative of the *relative entropy* between $P(x)$ and $R(x - \theta)$, that is $H(\tau) = -D(P(x)||R(\theta - x))$ —see Section 2.3.1 in Chapter 2 for the definition of relative entropy. If $R(\eta)$ is even about a mean of zero, and θ is equal to the signal mean of zero, then the entropy of τ is the negative of the relative entropy, $D(P(x)||R(x))$. Since relative entropy is always positive, this means that the entropy of $Q(\tau)$ is always negative or zero. Since τ is a continuously valued random variable, its entropy is differential entropy, and negative entropy is allowable.

The relative entropy for the four cases of matched Gaussian, logistic, Laplacian and Cauchy distributions is shown for $\theta = 0$ in Fig. 4.5, against increasing σ . The Gaussian case was calculated from the exact formula given in Table 4.2. The other cases were found numerically. Clearly, the relative entropy is zero at $\sigma = 1$, which is as expected, since at $\sigma = 1$, $P(x) = R(x) \forall x$. For $\sigma < 1$, the relative entropy becomes very large as

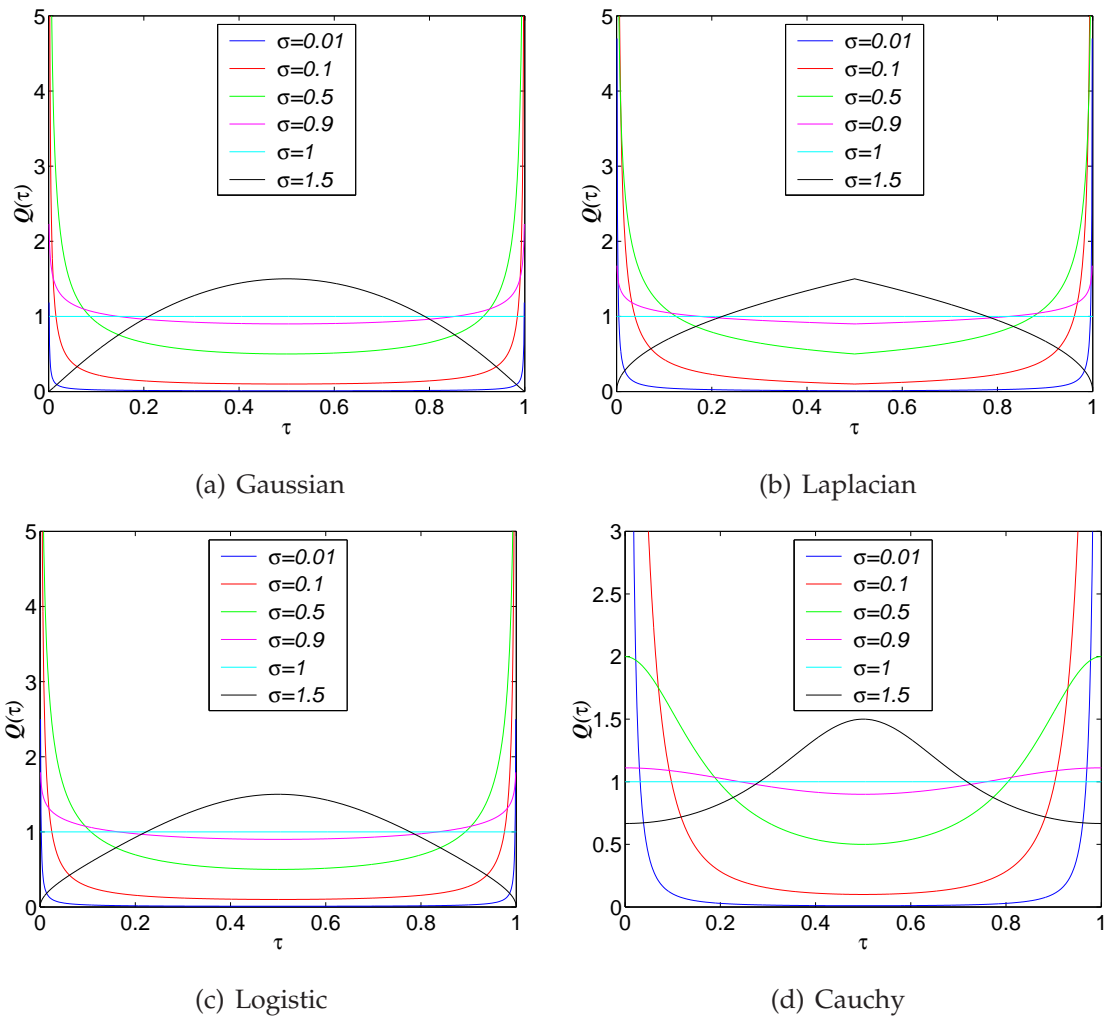


Figure 4.4. $Q(\tau)$ for various even PDFs and $\theta = 0$. Plot showing $Q(\tau)$ for various values of σ for four even, infinite support PDFs, where both $R(\cdot)$ and $P(\cdot)$ have identical PDFs and $\theta = 0$. For $\sigma < 1$, $Q(\tau)$ is convex, for $\sigma = 1$, $Q(\tau)$ is uniform, and for $\sigma > 1$, $Q(\tau)$ is concave. For the Gaussian, Laplacian and logistic cases, $Q(\tau)$ is infinite at $\tau = 0$ and $\tau = 1$ for $\sigma < 1$, and is equal to zero at $\tau = 0$ and $\tau = 1$ for $\sigma > 1$. For the Cauchy case, $Q(0) = Q(1) = \frac{1}{\sigma\lambda} \forall \sigma$.

σ gets smaller, and for $\sigma > 1$ gets larger, but more slowly with σ . The Gaussian case always gives the largest relative entropy for the same value of σ .

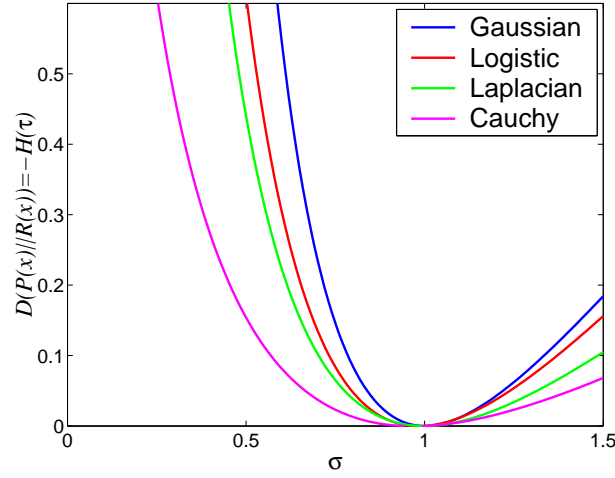


Figure 4.5. Relative entropy between $P(x)$ and $R(x)$. This figure shows the relative entropy—which is equivalent to the negative of the entropy of τ —plotted against σ for four cases of matched signal and noise distributions, for $\theta = 0$. The Gaussian case was calculated from the exact formula given in Table 4.2. The other cases were found numerically. Clearly, the relative entropy is zero at $\sigma = 1$, which is as expected, since at $\sigma = 1$, $P(x) = R(x) \forall x$. For $\sigma < 1$, the relative entropy becomes large as σ gets smaller, and for $\sigma > 1$ also becomes larger, but more slowly with σ . The Gaussian case always gives the largest relative entropy for the same value of σ .

Also, the first two moments of the random variable, τ , are

$$E[\tau] = \int_{\tau=0}^{\tau=1} \tau Q(\tau) d\tau, \quad (4.44)$$

$$E[\tau^2] = \int_{\tau=0}^{\tau=1} \tau^2 Q(\tau) d\tau. \quad (4.45)$$

Notice also that the term that appears on the RHS in Eqn. (4.30) and Eqn. (4.31) is $E[\tau \log_2(\tau)]$.

For the specific distributions listed in this Section that have infinite support, $Q(\tau)$ has support $\tau \in [0, 1]$. For the uniform case, when $\sigma > 1$, τ has support $[-\frac{1}{2\sigma} + 0.5, \frac{1}{2\sigma} + 0.5]$. Note that for uniform signal and noise, $Q(\tau)$ is only valid for $\sigma \geq 1$, since for $\sigma \leq 1$, $\frac{d\tau}{dx} = 0$ for $x < -\frac{\sigma\eta}{2}$ and $x > \frac{\sigma\eta}{2}$. In all the cases in Table 4.2, $Q(\tau)$ is a PDF, as one would expect from Eqn. (4.20). This is proven in Section B.4 in Appendix B for the cases of Gaussian and Laplacian signal and noise.

Note also from Table 4.2 and Fig. 4.4 that, apart from the Cauchy case, when $\sigma = 1$, $Q(\tau) = 1 \forall \tau$, that is, $Q(\tau)$ is uniform. Furthermore, for $\sigma > 1$, $Q(0) = Q(1) = 0$, so

4.3 Suprathreshold Stochastic Resonance

that $Q(x)$ is concave, and for $\sigma < 1$, $Q(x)$ is convex and $Q(0) = Q(1) = \infty$. For the Cauchy case, a finite limit exists at $\tau = 0$ and $\tau = 1$, such that $Q(0) = Q(1) = \frac{1}{\sigma_\lambda} \forall \sigma$.

That a PDF may have values that approach infinity may at first seem surprising. However, to be a PDF, a function only needs to be nonnegative for its entire support, and have a total area of unity when integrated with respect to its support. Infinite values do not preclude this, no more than the fact that a Gaussian PDF is never equal to zero, even for infinite values of its support. An example of a relatively well known PDF with infinite values is the PDF of a sine wave with a random phase and amplitude A (Damper 1995, Wannamaker *et al.* 2000a),

$$P(x) = \begin{cases} \frac{1}{\pi\sqrt{A^2-x^2}}, & x \in [-A, A], \\ 0 & \text{otherwise.} \end{cases} \quad (4.46)$$

This PDF is plotted in Fig. 4.6 for $A = 0.5$.




Figure 4.6. PDF of a random phase sine wave. This figure shows the PDF of a random phase sine wave that has an amplitude of $A = 0.5$. Notice that the PDF approaches infinity at $x = \pm A$, just as the PDF $Q(\tau)$ does for small values of σ .

Interpreting $Q(\tau)$

For uniform signal and noise with $\sigma < 1$, define $Q'(\tau) = Q(\tau)$ to be valid for $x \in [-\frac{\sigma_\eta}{2}, \frac{\sigma_\eta}{2}]$. Hence, $Q'(\tau)$ is not a PDF for $\sigma < 1$, since $\int_{\tau=0}^{\tau=1} \sigma d\tau = \sigma$. Therefore any integrations over τ in this case require extra consideration of $x < -\frac{\sigma_\eta}{2}$ and $x > \frac{\sigma_\eta}{2}$. This will not be the case for the integral on the RHS of Eqn. (4.28) since in this range of x , $P_{1|x}$ is either zero or unity and therefore $P_{1|x} \log_2 P_{1|x} = 0$. Likewise, it will not be the

case in calculating $P_y(n)$, since $P_{1|x}(1 - P_{1|x}) = 0$, except when $n = 0$ or $N = 1$. A simple integration in these cases finds that

$$\begin{aligned} P_y(0) = P_y(N) &= \int_{\tau=0}^{\tau=1} Q'(\tau)\tau^N d\tau + \int_{x=\sigma_\eta/2}^{x=\sigma_x/2} \frac{1}{\sigma_x} 1^N dx \\ &= \frac{\sigma}{N+1} + \frac{1}{2} - \frac{\sigma}{2}. \end{aligned} \quad (4.47)$$

Note that the integral in the region $x \in [\sigma_\eta/2, \sigma_x/2]$ is the region of the support of $P(x)$ for which all thresholds are always ‘on.’ This occurs since for these values of x , the noise is never negative enough to cause any threshold devices to be ‘off,’ that is, we have $x + \eta > \theta \forall x \in [\sigma_\eta/2, \sigma_x/2]$.

Thus, a possible interpretation of the significance of $Q(\tau)$ is that it is related to the output distribution, $P_y(n)$. As previously noted, for $\sigma < 1$, $Q(\tau)$ is concave, and for $\sigma > 1$, $Q(\tau)$ is convex. For $\sigma = 1$, $Q(\tau)$ is uniform. Recall also that $\sigma = 1$ corresponds to the case of optimal output entropy, so that $P_y(n) = \frac{1}{N+1}$, that is, the output probability mass function is uniform. Furthermore, for $\sigma = 0$, only output states 0 and N are available, each with probability one half. Also, for uniform signal and noise and $\sigma < 1$ where $Q'(\tau)$ is not a PDF, there is a contribution to $P_y(0)$ and $P_y(N)$ that cannot be explained by this interpretation of $Q'(\tau)$.

The obvious possibility is to suggest that

$$P_y(n) = \int_{\tau=\frac{n}{N+1}}^{\tau=\frac{n+1}{N+1}} Q(\tau) d\tau. \quad (4.48)$$

However, this does not agree with Eqn. (4.24). In the case of large N however, it will be seen in Chapter 5 that the output distribution does indeed approach $Q(\tau)$.

Hence, given that $Q(\tau)$ is a PDF, an intuitive explanation of its significance is that it is the PDF of the average fraction of thresholds that are ‘on’, given x , that is, $Q(\tau)d\tau$ is the probability that 100τ is the average percentage of times that $x + \eta > \theta$. Hence, since there are N thresholds, the average number of thresholds crossed is $N\tau$, and $Q(\tau)$ is the probability that $N\tau$ thresholds are crossed. We discuss this in more detail in Section 5.4.3 in Chapter 5.

By contrast with the probability mass function, $P_y(n)$, of the output states of the discrete random variable, y , the random variable of the average number of thresholds crossed is a continuous random variable. The reason that the average value of y is a continuous random variable, is that it is being conditioned on x , which is also a continuous random variable. Hence, the average value of y given x —that is, $E[y|x]$ —is

4.3 Suprathreshold Stochastic Resonance

described by a PDF, which we are interpreting as $Q(\tau)$. Thus, the CDF of $Q(\tau)$ evaluated at x is the probability that $E[y|x]/N$ is less than x .

Furthermore, the expected value of τ , as given by Eqn. (4.44), is then the actual expected value of the output y (not conditioned on x) and the expected value of τ^2 , as given by Eqn. (4.45), is the mean square value of the output (not conditioned on x).

This subsection has presented the distributions we will study for the signal and/or noise, as well as commenting on the function $Q(\tau)$, and giving an interpretation for its significance in the SSR model. The next subsection presents an exact result for the mutual information for the specific case of signal and noise both being uniformly distributed, and $\sigma \leq 1$.

4.3.4 Exact Result for Uniform Signal and Noise, and $\sigma \leq 1$

Exact analytical results for the SSR model are rare, and in most cases quantities need to be found numerically. However, Stocks (2001c) gives an exact expression for the mutual information in the case of a uniform signal and uniform noise when $\sigma \leq 1$ and $\theta = 0$.

In this case, $P_y(0)$ and $P_y(N)$ have already been derived, as given by Eqn. (4.47). For other n , from Eqn. (4.24) for $x \in [-\frac{\sigma\eta}{2}, \frac{\sigma\eta}{2}]$,

$$P^*(n) = \sigma \int_0^1 \tau^n (1 - \tau)^{N-n} d\tau. \quad (4.49)$$

For x outside this range, there is no additional contribution to $P^*(n)$, as already discussed. Eqn. (4.49) is the same expression as Eqn. (6) in Stocks (2001c). As noted in Stocks (2001a), this integral is simply a Beta function. The solution to such a Beta function (Spiegel and Liu 1999) is

$$P^*(n) = \sigma \frac{\Gamma(n+1)\Gamma(N-n+1)}{\Gamma(N+2)}, \quad (4.50)$$

where $\Gamma(\cdot)$ is the Gamma function (Spiegel and Liu 1999). For integer k , $\Gamma(k+1) = k!$, and hence, since n and N are integers,

$$P^*(n) = \sigma \frac{n!(N-n)!}{(N+1)!} = \frac{\sigma}{(N+1)\binom{N}{n}}. \quad (4.51)$$

Therefore

$$P_y(n) = \binom{N}{n} P^*(n) = \frac{\sigma}{N+1}. \quad (4.52)$$

Thus, the complete output probability mass function for uniform signal and noise and $\sigma \leq 1$ is

$$P_y(n) = \begin{cases} \frac{\sigma}{N+1} + \frac{1}{2} - \frac{\sigma}{2} & \text{for } n = 0, N \\ \frac{\sigma}{N+1} & \text{for } n = 1, \dots, N-1. \end{cases} \quad (4.53)$$

It is clear that when $\sigma = 0$, $P_y(0) = P_y(N) = 0.5$, which is expected, as in the absence of noise either all comparators are switched on, or all are off. It is also clear that for $\sigma = 1$ that $P_y(n) = \frac{1}{N+1} \forall n$, which is verified by Eqn. (4.32).

Thus, the output entropy is

$$\begin{aligned} H(y) &= -2P_y(0) \log_2 P_y(0) - (N-1)P_y(1) \log_2 P_y(1) \\ &= -\left(\frac{2\sigma}{N+1} + 1 - \sigma\right) \log_2 \left(\frac{\sigma}{N+1} + \frac{1}{2} - \frac{\sigma}{2}\right) - \frac{(N-1)\sigma}{N+1} \log_2 \left(\frac{\sigma}{N+1}\right). \end{aligned} \quad (4.54)$$

The conditional output entropy can be evaluated from Eqn. (4.27). First,

$$\int_{\tau=0}^{\tau=1} Q'(\tau) \tau \log_2 \tau d\tau = \frac{-\sigma}{4 \ln 2}, \quad (4.55)$$

which is valid for all x , as previously discussed. Secondly,

$$\begin{aligned} -\sum_{n=0}^N \log_2 \binom{N}{n} P_y(n) &= -\sum_{n=1}^{N-1} \log_2 \binom{N}{n} P_y(n) \\ &= -\frac{\sigma}{N+1} \sum_{n=1}^{N-1} \log_2 \binom{N}{n} - P_y(0) \log_2 1 - P_y(N) \log_2 1 \\ &= -\frac{\sigma}{N+1} \sum_{n=0}^N \log_2 \binom{N}{n} \\ &= \frac{\sigma}{N+1} \sum_{n=2}^N (N+1-2n) \log_2 n, \end{aligned} \quad (4.56)$$

since, as before, $-\sum_{n=0}^N \log_2 \binom{N}{n} = \sum_{n=1}^N (N+1-2n) \log_2 n$. Thus,

$$H(y|x) = \sigma \left(\frac{N}{2 \ln 2} + \frac{1}{N+1} \sum_{n=2}^N (N+1-2n) \log_2 n \right). \quad (4.57)$$

It is clear that this is exactly σ times Eqn. (4.39), the average conditional entropy for identical signal and noise distributions at $\sigma = 1$, and that therefore $H(y|x)$ is a linearly increasing function of σ for this specific case. This is due to $P_y(n)$ being uniform at $\sigma = 1$, and $P_y(n)$ being uniform—that is, $P_y(n) = \sigma/(N+1)$ for all n except $n = 0$ and $n = N$. The $n = 0$ and $n = N$ terms do not appear in the conditional entropy term due to being multiplied by the logarithm of unity.

4.3 Suprathreshold Stochastic Resonance

Therefore, the complete exact expression for the mutual information for uniform signal and noise and $\sigma \leq 1$ is

$$I(x, y) = - \left(\frac{2\sigma}{N+1} + 1 - \sigma \right) \log_2 \left(\frac{\sigma}{N+1} + \frac{1}{2} - \frac{\sigma}{2} \right) - \frac{(N-1)\sigma}{N+1} \log_2 \left(\frac{\sigma}{N+1} \right) - \frac{\sigma}{N+1} \sum_{n=2}^N (N+1-2n) \log_2 n - \frac{N\sigma}{2 \ln 2}. \quad (4.58)$$

This expression is that derived in Stocks (2001a) and stated as Eqn. (7) in Stocks (2001c). Note that if $\sigma = 0$ in Eqn. (4.58) then $I(x, y) = 1$ and if $\sigma = 1$, Eqn. (4.58) becomes identical to Eqn. (4.40).

This exact expression for the mutual information is plotted against σ for various values of N in Fig. 4.7, as well as the average conditional output entropy and the output entropy. Apart from $N = 1$, it is clear that there is a maximum in the mutual information for some value of σ between zero and unity. As N gets larger, the value of σ at which the maximum occurs gets closer to unity. It is also clear that both the output entropy, and the average conditional output entropy are strictly increasing with increasing σ . Note that, as indicated by Eqn. (4.57), the average conditional output entropy is a linear function of σ , whereas the output entropy increases with a larger slope for smaller σ .

The maximum in the mutual information occurs for the value of σ for which the slope of the output entropy is equal to the slope of the average conditional output entropy. Label this optimal value as σ_o . This can be found by differentiating $I(x, y)$ with respect to σ and setting the result to zero. Carrying this out, and after some algebra, the following expression gives σ_o as a function of N ,

$$\sigma_o = \frac{N+1}{N-1+2^{B+1}}, \quad (4.59)$$

where $B = \frac{N(N+1)}{2 \ln 2(N-1)} + \frac{1}{N-1} \sum_{n=2}^N (N+1-2n) \log_2 n$. The mutual information at this optimal value is indicated in Fig. 4.7(a) with crosses, for $N = 2, \dots, 1000$.

Before discussing these observations further, the next subsection presents numerical results for the other signal and noise distributions given in Table 4.1.

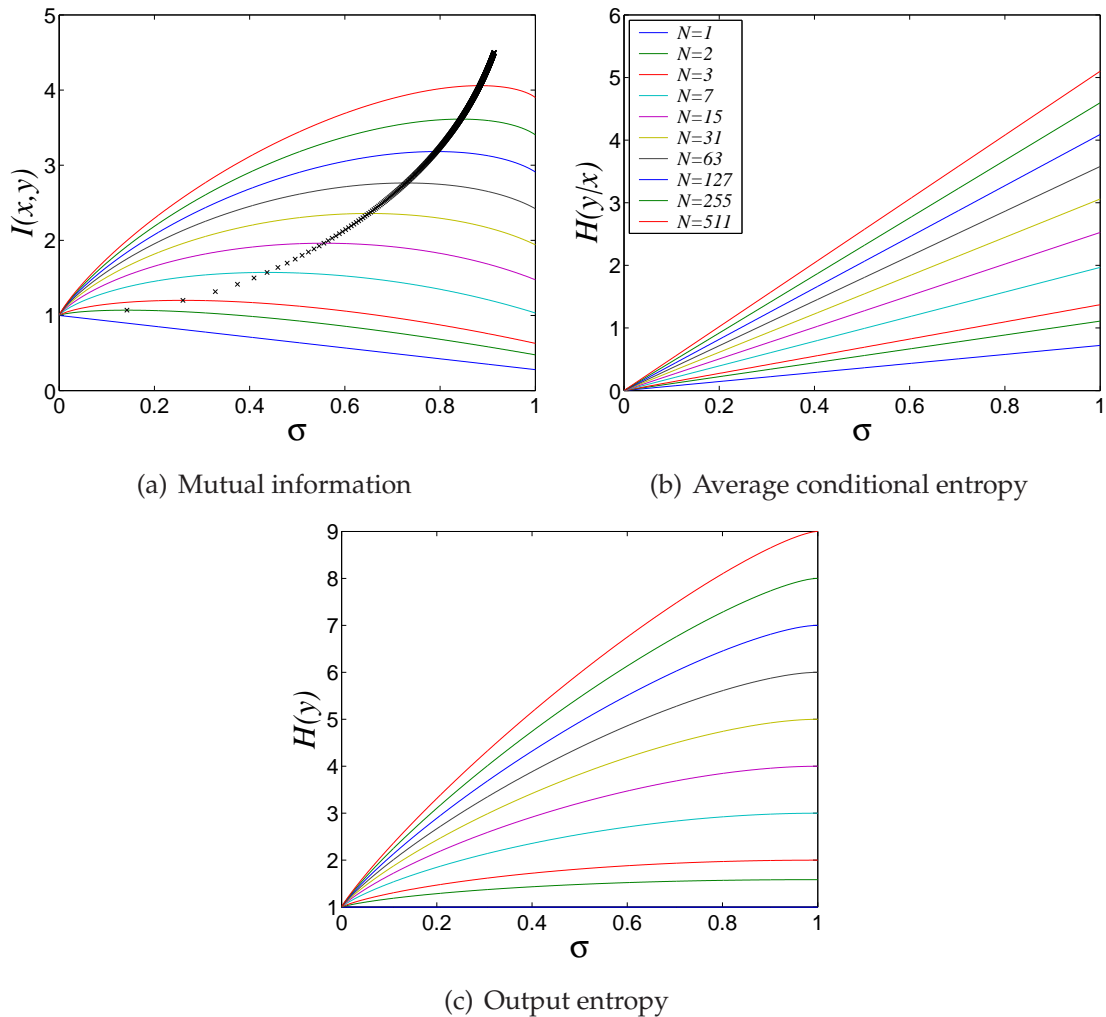


Figure 4.7. Exact $I(x, y)$ for uniform signal and noise and $\sigma \leq 1$. Fig. 4.7(a) shows the exact mutual information, Fig. 4.7(b) shows the average conditional output entropy, and Fig. 4.7(c) shows the output entropy, in the case of uniform signal and noise, and $\sigma \leq 1$, for various values of N . The threshold value is $\theta = 0$. The increasing value of the curves up the y -axis correspond to larger values of N . The average conditional entropy can be seen to increase linearly with increasing σ , as indicated in Eqn. (4.57). The output entropy is also an increasing function of σ , but appears to increase faster for smaller σ than for larger σ . The mutual information shows a noise induced maximum, and the value of σ at which this occurs increases with increasing N . The optimal value of mutual information is shown with a cross in Fig. 4.7(a) for $N = 2, \dots, 1000$.

4.3.5 Numerical Results

In this Section we show that SSR is not a phenomenon that occurs only for the specific cases of uniform or Gaussian signal and noise distributions. We begin by considering various cases of ‘matched’ signal and noise, so that the signal and noise both have the same distribution, although with different variances. This is the same situation as studied in Stocks (2000a), Stocks (2001c) and Stocks (2001a). The facts that an exact solution for the mutual information holds in such a case when $\sigma = 1$, and the mutual information is a function of the ratio, σ , and not both the signal and the noise variance independently, are the reasons that this situation is the natural situation to examine first. However, there is no other reason why the signal and noise need to have the same distribution.

Therefore, we also briefly look at some examples where the noise is Gaussian, and the signal has some other distribution, to illustrate that the qualitative behaviour of SSR does not depend on the actual signal and noise distribution to any large extent.

The mutual information between the input and output signals of the array of threshold devices can be obtained by numerical integration for any given signal and noise distribution. Some technical details on how the numerical integrations have been carried out are presented in Section B.5 of Appendix B.

Matched signal and noise

We now present figures showing the mutual information, average conditional entropy, and output entropy for five different matched signal and noise distributions, for a range of σ and N , and $\theta = 0$. Each of these five distributions have PDFs which are even functions about a mean of zero.

Gaussian Fig. 4.8 shows the mutual information, average conditional entropy, and output entropy for Gaussian signal and noise.

Uniform Fig. 4.9 shows the mutual information, average conditional entropy, and output entropy for uniform signal and noise. Fig. 4.9 also shows with circles the exact mutual information, average conditional entropy and output entropy already plotted in Fig. 4.7, for a number of values of σ . Clearly, the numerically calculated results agree with the known exact values in all cases.

Laplacian Fig. 4.10 shows the mutual information, average conditional entropy, and output entropy for Laplacian signal and noise.

Logistic Fig. 4.11 shows the mutual information, average conditional entropy, and output entropy for logistic signal and noise.

Cauchy Fig. 4.12 shows the mutual information, average conditional entropy, and output entropy for Cauchy signal and noise.

In each case, the mutual information shows a noise induced maximum, and the value of σ at which this occurs increases with increasing N , and corresponds to the point where the slope of the output entropy is equal to the slope of the average conditional entropy. Clearly, the maximum in $I(x, y)$ must occur for $\sigma \leq 1$, since for $\sigma > 1$, the output entropy decreases while the average conditional entropy increases. Hence, the slope of each can never be the same, and therefore no stationary point in the mutual information exists for $\sigma > 1$. For $\sigma \leq 1$, both the average conditional entropy and the output entropy can be seen to increase with increasing σ . For $\sigma \geq 1$, the average conditional entropy continues to increase, but the output entropy decreases. This makes sense, since from Eqn. (4.36), the output entropy reaches its maximum value at $\log_2(N + 1)$ at $\sigma = 1$.

Each figure shows with circles the exact mutual information, average conditional entropy, and output entropy at $\sigma = 1$, calculated from Eqns. (4.40) and (4.36). Also shown with circles is the fact that the mutual information and output entropy is exactly unity at $\sigma = 0$, and that the average conditional entropy is zero at $\sigma = 0$. Clearly, the numerically calculated results agree with the known exact values in all cases.

Given the very similar qualitative behaviour for these five signal and noise pairs, it is of interest to examine more closely the differences between them. To facilitate this, Fig. 4.13 shows the mutual information, average output entropy, and output entropy for all five cases superimposed, for $N = 127$. As of course should be the case, given the previously presented theory, for $\sigma = 1$ the mutual information, average output entropy and output entropy is the same for all signal/noise pairs. A close inspection shows that the Gaussian case gives the largest maximum mutual information out of all the plotted cases. A further observation is that the maximum mutual information for the Gaussian, Laplacian and logistic cases occurs for nearly the same values of σ —a value that is much smaller than the maximising σ for the uniform case. The Cauchy

4.3 Suprathreshold Stochastic Resonance

case has a slightly smaller maximum, but also occurs for a much smaller σ than the uniform case.




Figure 4.8. $I(x,y)$, $H(y|x)$ and $H(y)$ for Gaussian signal and noise. Fig. 4.8(a) shows the mutual information, Fig. 4.8(b) shows the average conditional output entropy, and Fig. 4.8(c) shows the output entropy, in the case of Gaussian signal and noise for various values of σ and N . The threshold value is $\theta = 0$. The increasing value of the curves up the y -axis correspond to larger values of N . The circles at $\sigma = 0$ —for which the mutual information and output entropy is always exactly unity—and $\sigma = 1$, indicate exact values for each quantity, calculated using Eqns. (4.40) and (4.36). For $\sigma \leq 1$, both the average conditional entropy and the output entropy can be seen to increase with increasing σ . For $\sigma \geq 1$, the average conditional entropy continues to increase, but the output entropy decreases. This makes sense, as the output entropy reaches its maximum value at $\sigma = 1$ of $\log_2(N + 1)$. The mutual information shows a noise induced maximum, and the value of σ at which this occurs increases with increasing N , and corresponds to the point where the slope of the output entropy is equal to the slope of the average conditional entropy.




Figure 4.9. $I(x,y)$, $H(y|x)$ and $H(y)$ for uniform signal and noise. Plot showing the mutual information, average conditional output entropy, and output entropy, against σ for various values of N and uniform signal and noise. The solid line indicates numerical integration and the circles indicate the exact results that appear in Fig. 4.7. The increasing value of the curves up the y -axis correspond to larger values of N .




Figure 4.10. $I(x,y)$, $H(y|x)$ and $H(y)$ for Laplacian signal and noise. Plot showing the mutual information, average conditional output entropy, and output entropy against σ for various values of N and Laplacian signal and noise. The increasing value of the curves up the y -axis correspond to larger values of N . The solid lines were obtained by numerical integration and the circles at $\sigma = 0$ —for which the mutual information and output entropy is always exactly unity—and $\sigma = 1$, indicate exact values for each quantity, calculated using Eqns. (4.40) and (4.36).

4.3 Suprathreshold Stochastic Resonance




Figure 4.11. $I(x,y)$, $H(y|x)$ and $H(y)$ for logistic signal and noise. Plot showing the mutual information, average conditional output entropy, and output entropy against σ for various values of N and logistic signal and noise. The increasing value of the curves up the y -axis correspond to larger values of N . The solid lines were obtained by numerical integration and the circles at $\sigma = 0$ —for which the mutual information and output entropy is always exactly unity—and $\sigma = 1$, indicate exact values for each quantity, calculated using Eqns. (4.40) and (4.36).




Figure 4.12. $I(x,y)$, $H(y|x)$ and $H(y)$ for Cauchy signal and noise. Plot showing the mutual information, average conditional output entropy, and output entropy against σ for various values of N and Cauchy signal and noise. The increasing value of the curves up the y -axis correspond to larger values of N . The solid lines were obtained by numerical integration and the circles at $\sigma = 0$ —for which the mutual information and output entropy is always exactly unity—and $\sigma = 1$, indicate exact values for each quantity, calculated using Eqns. (4.40) and (4.36).

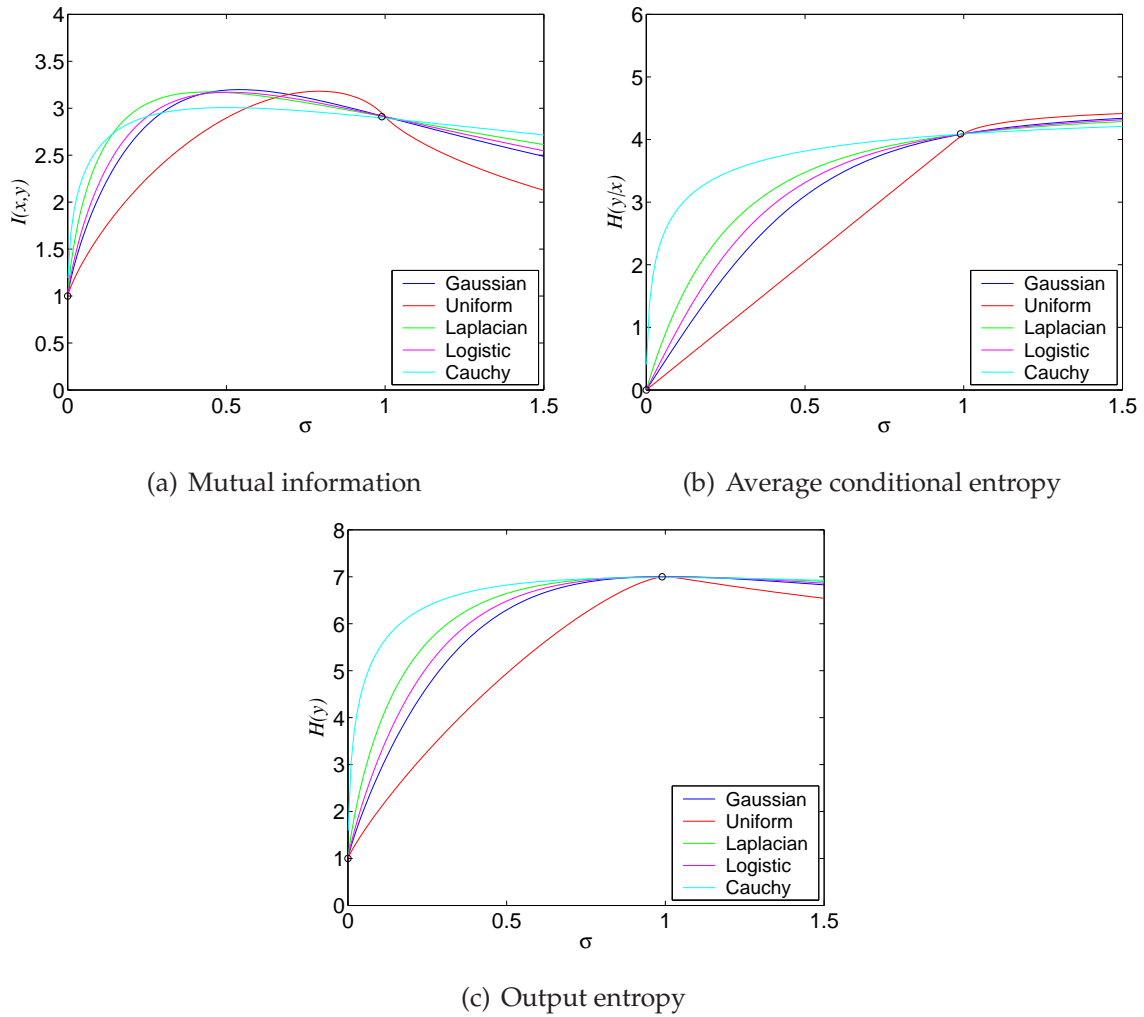


Figure 4.13. $I(x, y)$, $H(y|x)$ and $H(y)$ for $N = 127$ and five signal and noise pairs with even PDFs. Plot showing the mutual information, average conditional output entropy, and output entropy against σ for $N = 127$ and each of five signal and noise pairs. The solid lines were obtained by numerical integration and the circles at $\sigma = 0$ —for which the mutual information and output entropy is always exactly unity—and $\sigma = 1$, indicate exact values for each quantity, calculated using Eqns. (4.40) and (4.36). Notice that in all cases the mutual information, output entropy and average conditional output entropy have the same value at $\sigma = 1$, as it should be, given the known exact expressions at this point. The maximum mutual information for the Gaussian, Laplacian and logistic cases occurs for nearly the same values of σ , a value which is much smaller than the maximising σ for the uniform case. For the Cauchy case, the maximum σ is slightly smaller, but also occurs for a much smaller value of σ than the uniform case.

4.3 Suprathreshold Stochastic Resonance

We now consider two cases of distributions with single sided PDFs, that is, the PDF of random variables which are only defined for positive values.

Fig. 4.14 shows the mutual information, average conditional entropy, and output entropy for Rayleigh signal and noise, for a range of σ and N , and two values of the threshold. The first case, shown with solid lines, is for the threshold set to the signal mean, which is $\theta = \sigma_x \sqrt{\frac{\pi}{2}}$. This ensures that the mutual information at $\sigma = 0$ is equal to unity. However, due to the one-sided nature of the noise PDF, this also means that for all values of x greater than the mean, the output will always be N , since the noise can never help these values of the signal cross below any thresholds. This causes the maximum value of the mutual information to be much smaller than for the same N in the Gaussian case. The second case shown in Fig. 4.14 is for the threshold value set to twice the signal mean. Although this decreases the mutual information for small σ , it substantially increases it for large σ , since there is now far fewer values of x for which the output will always be zero.

Fig. 4.15 shows the mutual information, average conditional entropy, and output entropy for exponential signal and noise, for a range of σ and N . The threshold value is set to the signal mean, which is $\theta = \sigma_x$. This ensures that the mutual information at $\sigma = 0$ is equal to unity. Although the results shown in Fig. 4.15 show the same qualitative behaviour as in the Gaussian signal and noise case, the maximum value of the mutual information for each N is much smaller. As with the Rayleigh case, this is because with the threshold set to the signal mean, there is a large range of x for which the output is always zero.

To compare the mutual information in the Rayleigh and exponential cases to the even PDF cases, Fig. 4.16 shows the mutual information, average output entropy, and output entropy for the both these distributions for $N = 127$, as well as the case of Gaussian signal and noise. The Rayleigh case is shown for both threshold values plotted in Fig. 4.14. It is clear that the mutual information is far smaller for the Rayleigh and exponential cases than for the Gaussian case. However in the case of the Rayleigh distribution, where the threshold is set to twice the signal mean, the mutual information for large σ is very close to that obtained for the Gaussian case. Indeed, at $\sigma = 1$, it is only a small fraction smaller, indicating that even though the exact formula for the mutual information at $\sigma = 1$ given in Eqn. (4.40) does not apply for one-sided signal and noise PDFs, it can give a close upper bound, for an optimally set threshold.

In the next subsection we briefly consider mixed signal and noise distributions.




Figure 4.14. $I(x,y)$, $H(y|x)$ and $H(y)$ for Rayleigh signal and noise. Plot showing the mutual information, average conditional output entropy, and output entropy against σ for various values of N and Rayleigh signal and noise. The solid lines are the results for the threshold set to the signal mean, so that the mutual information is exactly unity in the absence of noise. The dashed lines are the results for the threshold set to twice the signal mean, showing that in this case a much larger maximum mutual information occurs. This happens due to the one-sided nature of the noise PDF; for all signal values above the mean, the noise cannot cause thresholds to be crossed, and the overall output will always be N . For a larger mean, this is less likely to occur, and for sufficiently large σ , the noise is better matched to the signal. The increasing value of the curves up the y -axis correspond to larger values of N .




Figure 4.15. $I(x,y)$, $H(y|x)$ and $H(y)$ for exponential signal and noise. Plot showing the mutual information, average conditional output entropy, and output entropy against σ for various values of N and exponential signal and noise. The threshold is set to the signal mean, $\theta = \sigma_x$, so that the mutual information is exactly unity in the absence of noise. As with the Rayleigh case, the maximum mutual information is far smaller than for the even PDFs, due to the one-sided nature of the noise PDFs. The increasing value of the curves up the y -axis correspond to larger values of N .

4.3 Suprathreshold Stochastic Resonance

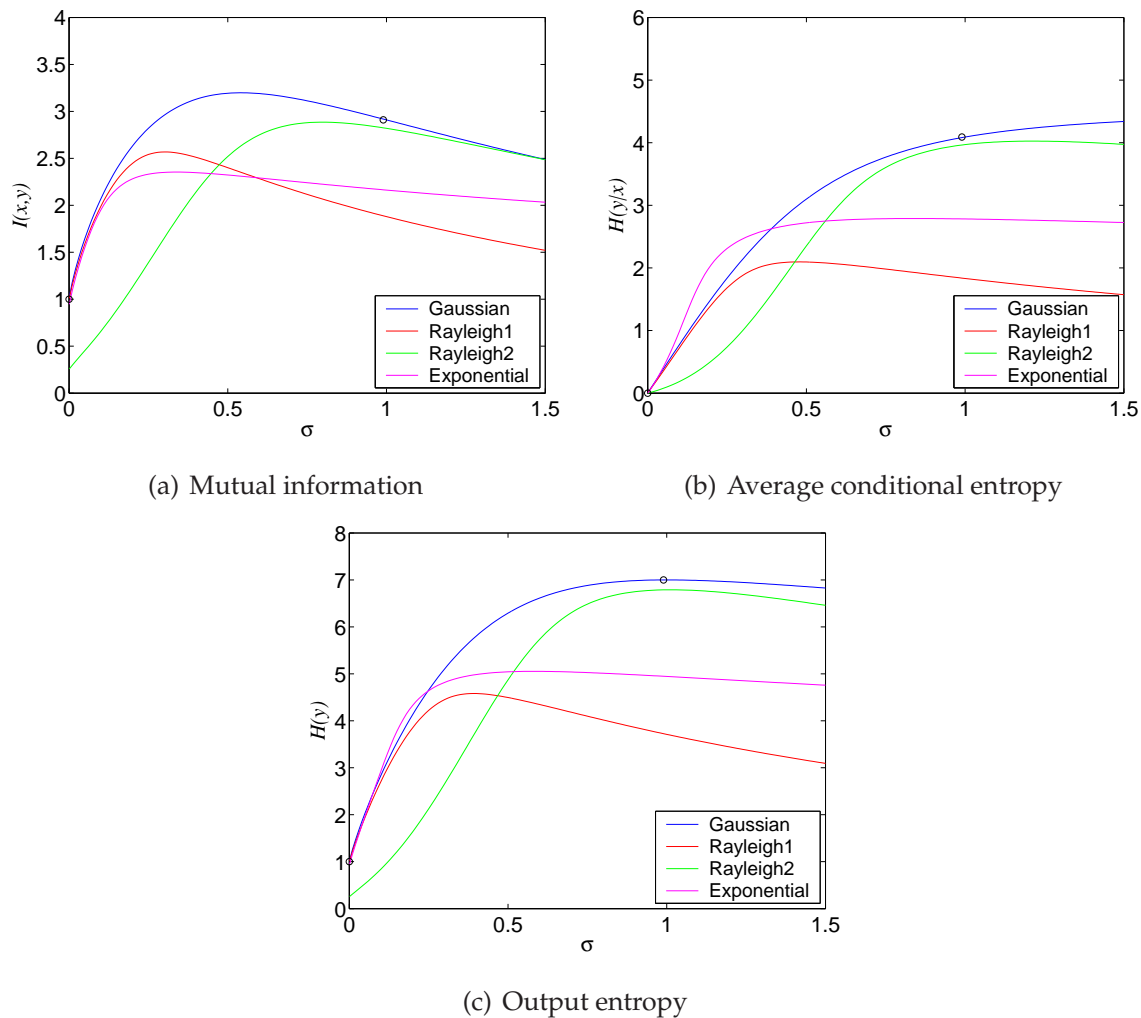


Figure 4.16. $I(x,y)$, $H(y|x)$ and $H(y)$ for $N = 127$ and the Rayleigh and exponential cases.

Plot showing the mutual information, average conditional output entropy, and output entropy against σ for $N = 127$ and the cases of Rayleigh signal and noise, and exponential signal and noise. The Rayleigh case is shown for two different threshold values, firstly the threshold set to the signal mean (Rayleigh1), and secondly the threshold set to twice the signal mean (Rayleigh2). The Gaussian case is shown for comparison. The circles indicate known exact results at $\sigma = 0$ and $\sigma = 1$ that apply for the Gaussian case. Note that there is no exact result for the Rayleigh or exponential cases, due to the one-sided nature of their PDFs. It is clear that the maximum mutual information is far smaller for the Rayleigh and exponential cases than for the Gaussian case. This is due to the one-sided nature of the noise PDF. However, for a threshold set to allow more threshold crossings, the mutual information in the Rayleigh case approaches that of the Gaussian case for large σ near $\sigma = 1$.

Mixed signal and noise

Unlike in the matched cases above, in general the mutual information will be a function of both the signal and noise variances independently. Hence, our results for the mutual information in this Section are plotted as a function of the ratio of noise standard deviation to signal standard deviation, and the signal standard deviation will be set to unity.

We consider the case of Gaussian noise. This is usually the first assumption to be made when the distribution of a noise source is unknown. Partly this is because Gaussian distributions are ubiquitous in nature, and partly this is because for all continuously valued distributions with a known variance, the Gaussian distribution has the largest entropy.

Figs. 4.17, 4.18, 4.19 4.20 and 4.21 show with dashed lines the numerical results for the cases of uniform, Laplacian, logistic, Cauchy and Rayleigh signals, with Gaussian noise. The solid lines show the previously plotted results for noise having the same distribution as the signal. In all cases except for a uniform or Rayleigh signal, Gaussian noise provides a slightly larger maximum mutual information than for the matched case. In the Rayleigh case, a much larger mutual information results from Gaussian noise. Fig. 4.22 shows the mutual information for all cases with Gaussian noise for $N = 127$. Very similar qualitative behaviour is seen for each. The maximum mutual information is between 3 and 3.2 bits per sample for all cases, and occurs for a value of σ close to that for matched Gaussian signal and noise. Fig. 4.22 also shows that the mutual information for a Rayleigh signal and Gaussian noise, is very close to that of Gaussian signal and Gaussian noise, indicating that Gaussian noise is a far better match for a Rayleigh signal than Rayleigh noise.

Such results lead into the focus of the next section, that of *channel capacity*. Channel capacity is defined as the *maximum* possible mutual information, subject to certain constraints. The results presented here lead us to ask questions like “what noise distribution provides the maximum mutual information in the SSR model for a given signal?” and “what signal distribution provides the maximum mutual information for a given noise distribution?”

4.3 Suprathreshold Stochastic Resonance




Figure 4.17. $I(x, y)$, $H(y|x)$ and $H(y)$ for a uniform signal and Gaussian noise. Plot showing the mutual information, average conditional output entropy, and output entropy against noise intensity, σ , for various values of array size, N . The solid lines show uniform signal and uniform noise, and the dashed lines show uniform signal and Gaussian noise, with the signal standard deviation set to $\sigma_x = \sqrt{12}$ in both cases. The increasing value of the curves up the y -axis correspond to larger values of N .




Figure 4.18. $I(x, y)$, $H(y|x)$ and $H(y)$ for a Laplacian signal and Gaussian noise. Plot showing the mutual information, average conditional output entropy, and output entropy against σ for various values of N . The solid lines show Laplacian signal and Laplacian noise, and the dashed lines show Laplacian signal and Gaussian noise, with $\sigma_x = 1$ in both cases. The increasing value of the curves up the y -axis correspond to larger values of N .




Figure 4.19. $I(x, y)$, $H(y|x)$ and $H(y)$ for a logistic signal and Gaussian noise. Plot showing the mutual information, average conditional output entropy, and output entropy against σ for various values of N . The solid lines show logistic signal and logistic noise, and the dashed lines show logistic signal and Gaussian noise, with $\sigma_x = 1$ in both cases. The increasing value of the curves up the y -axis correspond to larger values of N .




Figure 4.20. $I(x, y)$, $H(y|x)$ and $H(y)$ for a Cauchy signal and Gaussian noise. Plot showing the mutual information, average conditional output entropy, and output entropy against noise intensity, σ , for various values of N . The solid lines show Cauchy signal and Cauchy noise, and the dashed lines show Cauchy signal and Gaussian noise. In both cases, for the Cauchy signal, the FWHM for the signal is $\lambda_x = 1$ and the FWHM for the noise is $\lambda_\eta = \sigma/5$. The increasing value of the curves up the y -axis correspond to larger values of N .

4.3 Suprathreshold Stochastic Resonance




Figure 4.21. $I(x,y)$, $H(y|x)$ and $H(y)$ for a Rayleigh signal and Gaussian noise. Plot showing the mutual information, average conditional output entropy, and output entropy against σ for various values of N . The solid lines show Rayleigh signal and Rayleigh noise, and the dashed lines show Rayleigh signal and Gaussian noise. In both cases, $\sigma_x = 1$ and the threshold value, θ , is set to the signal mean so that $\theta = \sqrt{2 - \frac{\pi}{2}}\sigma_x$. The increasing value of the curves up the y -axis correspond to larger values of N . Clearly, Gaussian noise provides a far larger maximum mutual information than for matched signal and noise in this case.

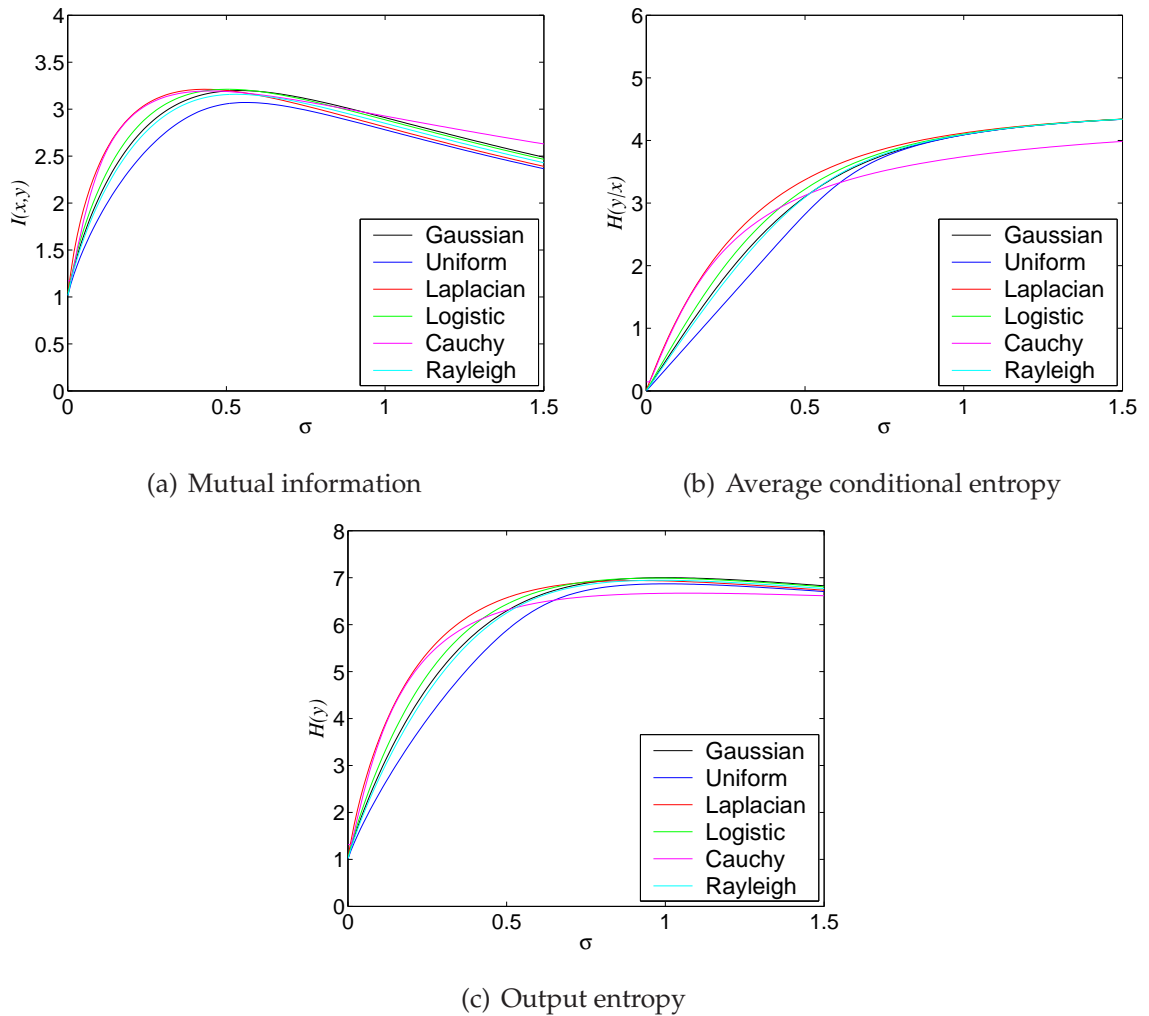


Figure 4.22. $I(x,y)$, $H(y|x)$ and $H(y)$ for $N = 127$, various signals and Gaussian noise.

Plot showing the mutual information, average conditional output entropy, and output entropy against σ for $N = 127$ and various signal distributions with Gaussian noise. The maximum mutual information is between 3 and 3.2 bits per sample, and occurs for nearly the same values of σ for all cases.

4.4 Channel Capacity for SSR

In information theory, the term channel capacity is usually defined as being the maximum possible (on average) mutual information, per sample, that can be transmitted through a given channel. Usually, the channel is fixed and the maximisation performed over all possible source probability distributions. Thus the problem of finding channel capacity, $C(x, y)$, can be expressed as the optimisation problem,

$$\text{Find: } C = \max_{\{P(x)\}} I(x, y). \quad (4.60)$$

Often there are prescribed constraints on the source distribution, $P(x)$, such as a fixed average power, or a finite alphabet, that is, $P(x)$ is discretely valued, with a fixed, finite number of states. Relevant attributes of the given channel are the noise power and distribution, and bandwidth. The second most well known example of a channel capacity formula is that for a binary symmetric channel—see Section 3.3.3 of Chapter 3.

The most well known example of a channel capacity formula is Shannon's famous, but often misused, formula for the channel capacity of a bandlimited, additive white Gaussian noise channel (Shannon 1948),

$$C = B \log_2 \left(1 + \frac{P_s}{P_n} \right) \quad \text{bits per second}, \quad (4.61)$$

where B is the channel bandwidth, P_s is the prescribed maximum mean square signal power and P_n is the mean square noise power (within the channel bandwidth). Note that this formula has many names—see Section 3.3.1 in Chapter 3.

The channel capacity of Eqn. (4.61) can be rewritten in its most well known form, in terms of the input SNR, as

$$C = B \log_2 (1 + \text{SNR}) \quad \text{bits per second}. \quad (4.62)$$

The main reason that this formula is often misapplied—see Berger and Gibson (1998) for a discussion—is that it only applies when the channel noise has a Gaussian distribution that is independent of the signal and is additive in nature. It also does not apply if the noise is not white in the channel bandwidth—that is, if the power spectral density of the noise is not constant for all frequencies in the passband of the channel—or if there are constraints other than the power constraint on the input signal. Furthermore, capacity can only be achieved if the input signal has a Gaussian distribution.

In channel capacity research, conventionally the channel is prescribed, and subject to certain imposed input signal constraints, the optimal signal distribution found. Previously in this chapter, we have seen that the mutual information in the SSR model is maximised by a nonzero value of noise. If we determine the optimal ratio of noise standard deviation to signal standard deviation, σ , we have effectively found channel capacity. At first glance, when compared to conventional studies of information channels, varying the noise level for a given signal seems as if it is equivalent to modifying the channel, rather than looking for an optimal source distribution.

However, recall that in the case of even PDFs, the mutual information for matched signal and noise is a function of σ . Thus, adjusting σ to find the maximum mutual information for given signal and noise pairs is also equivalent to modifying the source distribution, by increasing or decreasing its variance, for fixed noise.

Thus, the questions of “what noise variance maximises the mutual information, for given matched signal and noise distributions and a fixed signal variance?” and “what signal variance maximises the mutual information, for given matched signal and noise distributions and a fixed noise variance?” are equivalent. The following subsection looks at this question, and in light of the above discussion, is no different to conventional channel capacity questions, since the channel is fixed—that is, a noise distribution and variance is given—and the source distribution is to be found, subject to the constraint that it has the same PDF as the noise, but can have any variance.

4.4.1 Matched and Fixed Signal and Noise PDFs

In light of the proceeding discussion, consider the case of a channel being prescribed—that is, the number of comparators, N , and the noise PDF, $R(\eta)$, and variance, σ_η^2 , will remain fixed—and the constraint on the signal that its PDF is the same as the noise, other than the variance.

Thus, in this situation, finding the SSR channel capacity means finding the optimal signal variance—or, equivalently, power—for a given noise variance. Thus, we can express this problem in the following manner,

$$\begin{array}{ll}
 \text{Find:} & C(x, y) = \max_{\text{var}[x]} I(x, y) \\
 \text{subject to:} & R(\eta) \text{ and } \text{var}[\eta] \text{ fixed} \\
 \text{and} & P(x) \text{ matched with } R(\eta).
 \end{array} \tag{4.63}$$

4.4 Channel Capacity for SSR

The result of solving this optimisation problem is the optimal value of the variance of the signal. Since this value will hold for a specified variance of the noise, the result will give the optimal value of σ .

Although solving this problem analytically is intractable—apart from the uniform signal and noise case—numerical solution is straightforward. The result of such an optimisation is shown in Fig. 4.24(a), which shows the channel capacity for the cases of matched Gaussian, Laplacian and logistic signal and noise, for increasing N , and Fig. 4.23, which shows the values of σ at which channel capacity is achieved. Also shown in these figures is the capacity for uniform signal and noise, where σ_0 is calculated from Eqn. (4.59) and plotted in Fig. 4.23, and the mutual information at this σ_0 , plotted in Fig. 4.24(a). Fig. 4.24(a) also shows the mutual information obtained at $\sigma = 1$ using the exact formula of Eqn. (4.40), which holds in all cases.




Figure 4.23. Optimal σ for Gaussian, Laplacian and logistic signal and noise. This figure shows σ_0 —that is, the value of σ that achieves channel capacity—against increasing N . The Gaussian, Laplacian and logistic cases were calculated numerically, and the uniform case was calculated from the exact formula of Eqn. (4.59). Although it is not completely clear from this figure with the scale limited to $N \leq 1000$, for larger N , σ_0 asymptotically approaches a fixed value.

From Fig. 4.24(a) the difference between the capacity of each of the three infinite support distributions is very small, and indeed is hard to resolve on the scale shown. Hence, the difference between each case is plotted in Fig. 4.24(b). This plot shows that the difference in channel capacity between each case appears to be decreasing as N increases, apart from an initial increase in difference for small N .

Fig. 4.24(c) shows the difference between channel capacity, and the mutual information at $\sigma = 1$. The capacity is, of course, always greater than the exact mutual information at $\sigma = 1$, although the difference between the capacity and this value appears to also approach a constant slightly greater than 0.2 bits per sample for large N , apart from the uniform case. In the uniform case, the optimising σ approaches unity for large N . Thus, as previously noted here, and in Stocks (2001a), the maximum in the mutual information appears to always be between $\sigma = 0$ and $\sigma = 1$, since at $\sigma = 0$, the mutual information is always 1 bit per sample.

It is clear from Fig. 4.24(a) that as N increases, the channel capacity also increases. However, from Fig. 4.23, the optimising value of σ appears to converge asymptotically to a constant value, σ_0 , as N increases. For Gaussian signal and noise, this large N value appears to be $\sigma_0 \simeq 0.607$, for Laplacian signal and noise, $\sigma_0 \simeq 0.5$ and for logistic signal and noise $\sigma_0 \simeq 0.564$. It is not visible on the axes of Fig. 4.23, but for uniform signal and noise, the large N value of $\sigma_0 \simeq 1$. This is shown to be the case in Chapter 5.

Furthermore, the difference between channel capacities for the various distributions appears as if it might approach zero for very large N . These results indicate the possibility that proofs of this behaviour may be possible in a large N limit. Chapter 5 is devoted to the investigation of expressions for the mutual information and channel capacity that hold for large N .

4.4.2 Channel Capacity in General

We saw for the case of mixed signal and noise distributions that Gaussian noise can provide a larger maximum mutual information than noise matched with the signal. Hence, the natural question to ask is “what signal distribution maximises the mutual information for a given noise distribution?” In contrast to the previous subsection, where we imposed the constraint that the signal distribution must be matched with the source distribution, and therefore the shape of the signal PDF is known, to solve this problem the optimal source PDF must be found. This problem can be expressed as

$$\begin{aligned} \text{Find:} & & C(x, y) &= \max_{P(x)} I(x, y) \\ \text{subject to:} & & R(\eta) \text{ and } \text{var}[\eta] &\text{ fixed.} \end{aligned} \quad (4.64)$$

The constraint that specifies that the noise distribution is fixed can be dropped from the expression of the problem, as it simply specifies the nature of the channel. Thus,

4.4 Channel Capacity for SSR

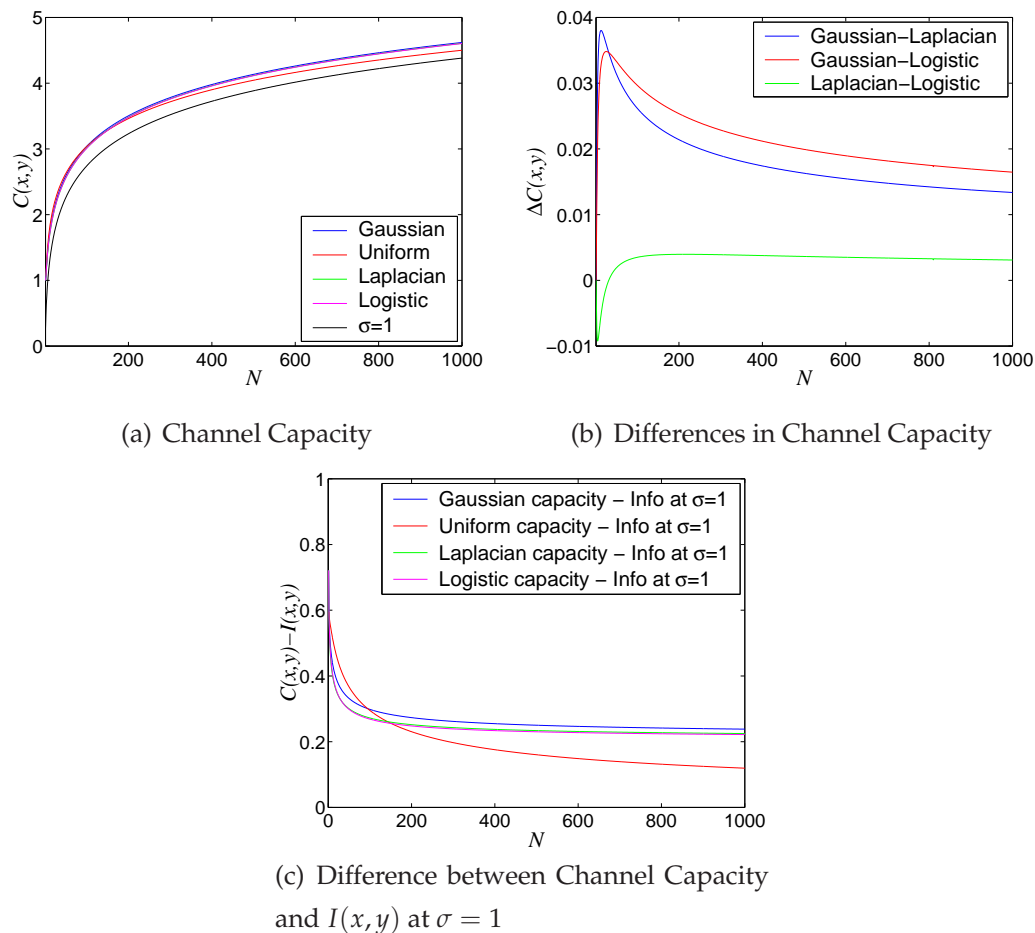


Figure 4.24. Channel capacity for matched signal and noise. Fig. 4.24(a) shows the channel capacity calculated numerically for Gaussian, Laplacian and logistic signal and noise pairs against increasing N . It also shows the exact channel capacity for uniform signal and noise, calculated from Eqns. (4.59) and (4.58), and the mutual information that holds for all cases where $P(x) = R(\theta - x)$ so that $\sigma = 1$, calculated from Eqn. (4.40). Clearly, the channel capacity is almost identical for the Gaussian, Laplacian and logistic cases, and is larger than the mutual information at $\sigma = 1$. Fig. 4.24(b) shows the difference between the channel capacity each of these three pairs is very small, and the difference gets smaller with larger N . Fig. 4.24(c) shows the difference between the channel capacity and the mutual information at $\sigma = 1$ for each case. Apart from the uniform case, this difference appears to approach a constant value, again indicating that the maximum mutual information occurs for $\sigma < 1$ for these cases for all N .

the optimisation problem to be solved is

$$\text{Find:} \quad C(x, y) = \max_{P(x)} I(x, y). \quad (4.65)$$

For channels with continuously valued inputs and outputs, usually a constraint such as a peak power or fixed average power constraint are specified. This is to give some realism to a signal's distribution, since otherwise the optimal $P(x)$ could have infinite power. Such channel capacity problems can be solved by use of the Arimoto-Blahut algorithm (Arimoto 1972, Blahut 1972, Cover and Thomas 1991).

This iterative algorithm begins with an initial guess for $P(x)$. It then uses this $P(x)$ along with the channel characteristics to calculate the joint probabilities of all pairs of input and output values, $P(x, n)$. The current mutual information can be calculated from these distributions. The algorithm then calculates a new $P(x)$, one that is optimal for the current joint probabilities. This process is repeated, until the algorithm converges to a solution such that the difference between successive values of the mutual information is smaller than a specified error tolerance, and therefore the channel capacity has been found.

The Arimoto-Blahut algorithm is fairly general, and can be easily extended to incorporate various constraints on the input signal, or specified costs for using various values of the input (Blahut 1972, Hoch *et al.* 2003b). It can also be used to calculate rate-distortion functions—see Chapter 9.

Applying the Arimoto-Blahut algorithm to calculating channel capacity for the SSR model, as specified in Problem (4.65) does indeed find solutions. Specific values for channel capacity for each N are slightly larger than in the case of matched signal and noise. In fact, although not once thought of as being a case of a noise optimised signal, such a calculation has been previously performed in Chang and Davisson (1988). The main purpose of Chang and Davisson (1988) is to specify two algorithms based on the Arimoto-Blahut algorithm for finding channel capacity in an infinite-input, finite-output channel. The SSR model is exactly such a channel.

Chang and Davisson (1988) use an example situation to test each of these algorithms, where the input signal is a continuously valued variable, x' , between zero and unity, and the output signal is integer valued between 0 and L . The channel transition probabilities are specified to be given by the binomial distribution as a function of x' as

$$P(k|x') = \binom{L}{k} (x')^k (1 - x')^{(L-k)} \quad k = 0, \dots, L. \quad (4.66)$$

4.4 Channel Capacity for SSR

If $P_{1|x}$ is substituted for x' , Eqn. (4.66) is recognisable as being equivalent to Eqn. (4.9). This substitution is allowable, since $P_{1|x} \in [0, 1]$ and $x' \in [0, 1]$. With reference to Table 4.1, a situation where $P_{1|x}$ increases linearly with x for all valid x , occurs when the noise is uniform, and the signal PDF is limited to the support of the uniform noise. Therefore, the channel of Chang and Davisson (1988) is exactly the same as the SSR model when subjected to uniform noise, provided the signal and noise have the same support.

However it turns out after replicating the results in Chang and Davisson (1988)—which does not give the actual optimal $P(x)$, only the capacity—the input PDF that achieves the channel capacity is quite often unrealistic, in that in most cases $P(x)$ has more than one local maximum, and has small regions of highly probable values located symmetrically a long distance from its mean, particularly if non-uniform noise is allowed, and the signal support allowed to be very large. Even in such situations, the capacity found, although larger than for matched signal and noise, is only larger by a few percent.

Thus, for the case of SSR, a more realistic channel capacity question is one that specifies a set of input PDF constraints that might include a maximum variance, combined with a minimum entropy. Furthermore, constraints on the number of local maxima of the PDF might be included. Such questions are left for future work.

Alternatively, constraints can be set on the output distribution. Such a situation can arise naturally in a neural coding context, where coding of information might be subject to strict energy constraints. This scenario has been studied in Hoch *et al.* (2003b), which uses the Arimoto-Blahut algorithm to derive optimal input PDFs under an output energy constraint, such as might occur in real neurons.

4.4.3 Coding Efficiency

Recall that the channel capacity was given for the binary symmetric channel, in terms of the probability of error, by Eqn. (3.11) in Section 3.3.3 of Chapter 3. For this channel, there is a well defined upper bound on the channel capacity of one bit per sample, since the output can only take two states. This is verified by Eqn. (3.11), which shows that the channel capacity is always less than or equal to unity, since terms of the form $p \log p$ are always less than or equal to zero when p is a probability.

Likewise, since the SSR model has a discretely valued output signal, it also has a well defined limit to its channel capacity. This is $\log_2(N + 1)$ bits per sample, since the

output has $N + 1$ states. This contrasts with the additive Gaussian white noise channel considered as the channel for Shannon's channel capacity formula, for which the output is a continuously valued signal. For such systems, there is no upper bound on the channel capacity, which can theoretically be infinite.

Note that it is possible to achieve this maximum channel capacity in a situation where not all thresholds are equal to the signal mean. For example, in the absence of noise, this maximum channel capacity is achieved by setting the thresholds such that all output states are equally likely. Such a scenario however is beyond the scope of this chapter, and will be considered in Chapter 8. Here, we are only considering the case of all thresholds identical.

The existence of such an upper bound on channel capacity allows the definition of a *coding efficiency* measure as

$$\Gamma(N) = \frac{C(x, y)}{\log_2(N + 1)}. \quad (4.67)$$

This measure gives the fraction of the maximum possible mutual information provided by given signal and noise distributions in the SSR model. If Γ is multiplied by 100, then the coding efficiency can be measured as a percentage. This idea is expressed in Stocks (2001b). An alternative way of measuring the same thing is to measure the difference between $C(x, y)$ and $\log_2(N + 1)$. However, since this is an absolute measure, and Γ is a relative measure, we will focus on Γ .

The results for channel capacity in the previous parts of this section can be re-plotted in terms of coding efficiency. Fig. 4.25 shows the coding efficiency for matched signal and noise for the Gaussian, uniform, Laplacian and logistic cases, as well as the coding efficiency when $P(x) = R(\theta - x)$, so that $\sigma = 1$. For large N , the coding efficiency appears to asymptotically approach a value slightly less than 0.5. As noted in Stocks (2001b), this means that SSR for matched signal and noise is less than 50 percent of the maximum noiseless channel capacity. However, given that in most systems, noise is unavoidable, it may be the case that a coding efficiency of about 50 percent is quite acceptable.

4.5 SSR as Stochastic Quantisation

We have seen in this chapter how the mutual information in the SSR model has a maximum for a nonzero noise intensity. We have also seen how the maximum mutual




Figure 4.25. Coding efficiency for matched signal and noise. This figure shows the coding efficiency, Γ , as a function of N , for matched signal and noise for the Gaussian, uniform, Laplacian and logistic cases. Also shown is the coding efficiency when $P(x) = R(\theta - x)$, so that $\sigma = 1$. For large N , the coding efficiency appears to asymptotically approach a value slightly less than 0.5.

information in the absence of noise is exactly one bit per sample, since all thresholds have the same value and the output can only ever be zero or N . Why is it that the presence of noise allows a larger information rate per sample? The explanation is that the presence of *independent* noise on each threshold has the effect of distributing all N thresholds to different values, rather than the same value. Instead of modelling the noise as being added to the signal, it is equivalent to model the noise as being threshold noise, so that the threshold is a random variable. Hence, if the threshold is θ , then the signal, x , is thresholded N times by N independent samples taken from the random variable, η , described by the PDF, $R(\theta - \eta)$.

Thus, when noise is present, effectively all threshold values will be unique. This allows the output to take values other than zero and N , and thus allow the output to become a $\log_2(N + 1)$ bit *stochastic quantisation* of the input signal. For noise with a small variance compared to the signal, most of the output states will occur with a very small—or zero, for noise with finite support—probability. Hence, very little extra information can be gained, since only a fraction of the $N + 1$ output states are utilised. However, when the noise variance is such that each output state is occupied with a probability that reflects the shape of the input PDF, far more information can be gained, although as we have seen, this appears to be at most about fifty percent of the maximum possible channel capacity. For a noise variance larger than optimal, the dependence between

the input and output signals decreases, and the mutual information starts to decrease again.

Such a description of SSR as stochastic quantisation is given in Stocks (2000b), which states that

“At any instant of time, finite noise results in a distribution of thresholds that, in turn, leads to the signal being ‘sampled’ at N randomly spaced points across the signal space.”

As discussed, the realisation that nonzero noise causes a distribution of thresholds is the key to understanding why a nonzero noise intensity gives best performance, the end result being a non-deterministic—or stochastic—quantisation of the signal.

Example random distributions of thresholds across the support of the Gaussian PDF is shown in Fig. 4.26. In the subfigures, Fig. 4.26(a) shows the optimal noiseless threshold values for this Gaussian source, while Figs. 4.26(b), 4.26(c) and 4.26(d) show examples of the random threshold distribution in the SSR model, for three different values of Gaussian noise variance. Clearly, although the example of the random threshold distribution does not give the optimal threshold distribution, as we have seen in this chapter it does give a distribution that allows the *average* mutual information of the SSR model to approach half the maximum noiseless channel capacity.

It is now clear why channel capacity occurs at $\sigma_o \rightarrow 1$ in the case of uniform signal and noise, but approaches a value much less than unity in the cases of infinite support PDFs. In the uniform case, when $\sigma > 1$ the random distribution of thresholds for any given input sample sometimes places thresholds outside the dynamic range of the signal. This can only reduce the mutual information. Furthermore, since for the uniform PDF, all values of x are equally probable, it is desirable that the random threshold distribution is just as likely to consist of a threshold near the mean of the signal as it is close to its maximum and minimum value. For large N , this can only occur if the noise PDF has a variance close to that of the signal.

By contrast, for the infinite support signal and noise cases, the most probable values of the signal occur within several standard deviations of the mean. Therefore, the noise distribution does not need to be as wide as the signal distribution for the optimal random threshold distribution to occur. Hence, channel capacity occurs for $\sigma_o \ll 1$ for these cases.

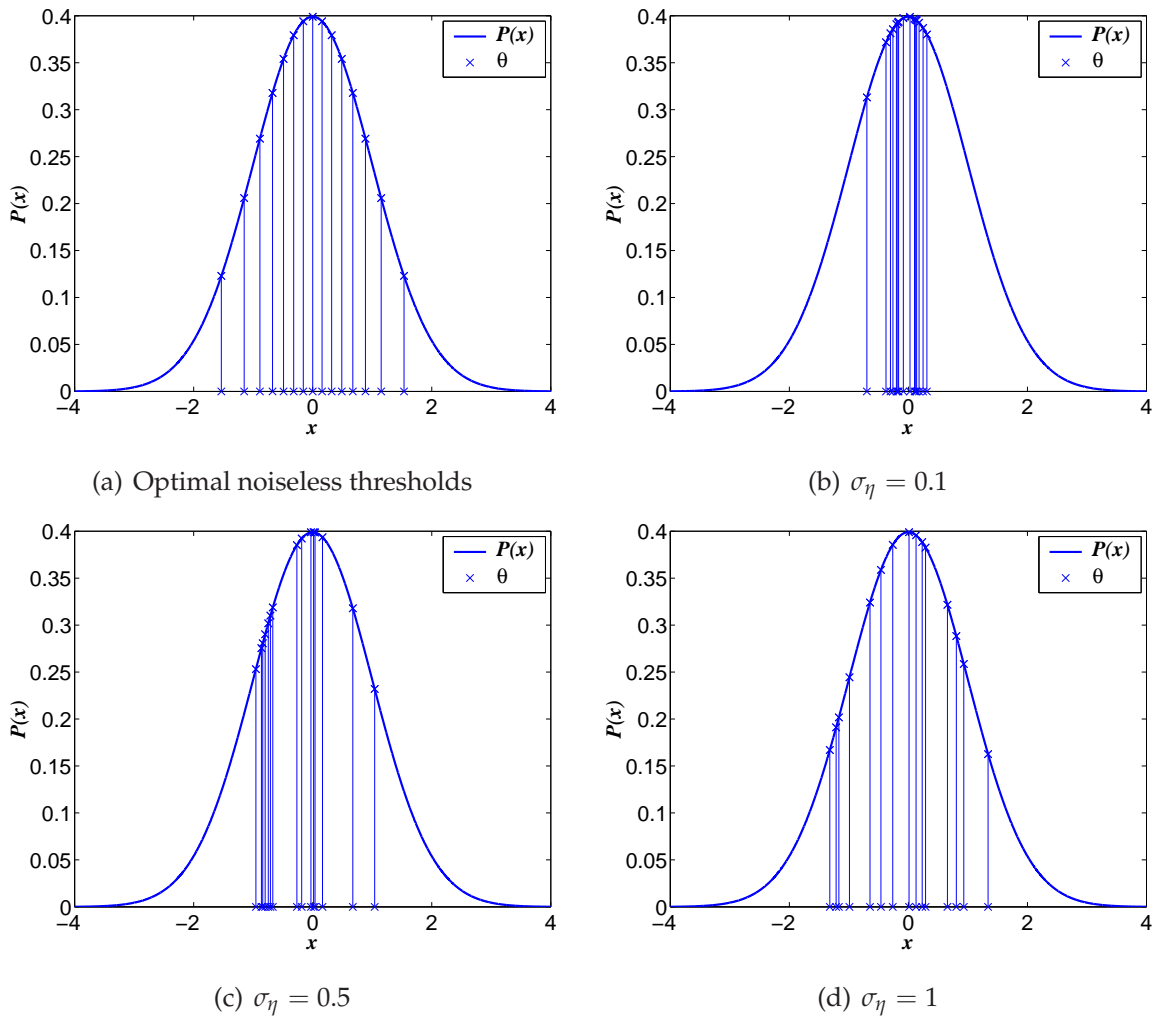


Figure 4.26. Random threshold distribution for Gaussian signal and noise. Fig. 4.26(a) shows the optimal—that is, maximum mutual information—threshold values, θ , for $N = 15$ and a unity variance Gaussian source. The stem plot indicates with crosses the values of x for each threshold, and the value of $P(x)$ at which each threshold occurs. Figs. 4.26(b), 4.26(c) and 4.26(d) show particular instances of the random SSR threshold distribution for three different values of noise variance. Notice that for small noise variance, most thresholds are likely to end up close to the signal mean. For intermediate and larger variances, more thresholds are likely to be located closer to the tails of the source.

4.5.1 Encoder Transfer Function

Given this stochastic quantisation interpretation of the SSR model, it is natural to wonder how the SSR output compares with a standard quantisation of the input. We will examine this question more closely in Chapter 6, however here we illustrate the affect of independent threshold noise in this context by comparing the transfer function of a conventional deterministic quantiser's encoder, with the mean transfer function of the SSR model's encoding. When conventional quantisation is introduced in standard textbooks, it is usual to graphically illustrate the transfer function by plotting the output values that correspond to certain input values.

For example, suppose a quantiser's output has 3 bits, so that $N = 7$ thresholds are required, and there are $N + 1 = 8$ output states. For the case where the thresholds uniformly quantise a signal between ± 1 , the transfer function is shown with a thick black line in Fig. 4.27(a). Note that this transfer function is deterministic, so therefore there is no uncertainty about which output value is attained for every input value.

In contrast, the average transfer function for SSR is simply the *expected* value of the output, y , given the input, x . Recall from Eqn. (4.9) that the conditional probability distribution of the output given the input, is given by the binomial distribution. Here, the expected value of the binomial distribution is given by

$$E[y|x] = NP_{1|x}, \quad (4.68)$$

which gives the average transfer function for SSR. However, unlike conventional quantisers, the output is not-deterministic. One way of characterising this uncertainty is with the conditional variance, $\text{var}[y|x]$, which for the binomial distribution is given by

$$\text{var}[y|x] = NP_{1|x}(1 - P_{1|x}). \quad (4.69)$$

Fig. 4.27(a) shows the *average* transfer function for the SSR model, with $N = 7$, for various values of noise standard deviation, σ_η , and Gaussian noise, as a function of x . Fig. 4.27(b) shows the variance corresponding to each value of σ_η as a function of x . We see that for the intermediate values of σ_η , the average transfer function appears to more closely approximate the deterministic transfer function, whereas for $\sigma_\eta = 0.3$ and $\sigma_\eta = 1.2$, the average transfer function appears to have too sharp, or too broad a slope to give a good approximation. The relationship between the ideal and the average transfer function is broadly measured by what is known as *bias*.

4.5 SSR as Stochastic Quantisation

However, the performance of the SSR model as a quantiser cannot be measured only by how well its average transfer function compares with an ideal transfer function. From Fig. 4.27(b), we see that the variance of the output for a given value of the input always increases with increasing σ_η . Hence, there is a tradeoff required between the bias and the variance for a given x . We will see quantitatively in Chapter 6 how the performance of the SSR model depends on both of these quantities, and also on the input distribution.

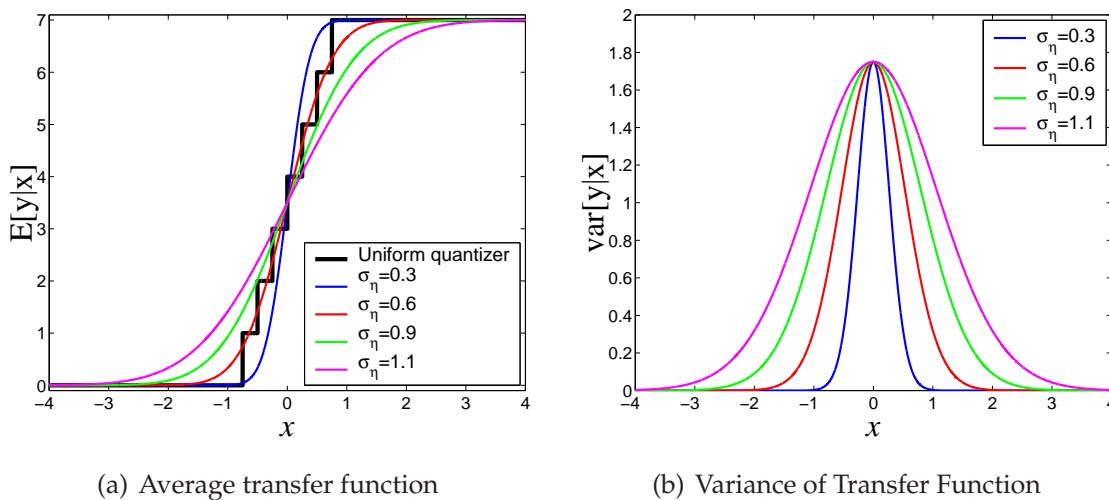


Figure 4.27. Average transfer function, and its variance for SSR. Fig. 4.27(a) shows with a thick black line the output of a conventional deterministic 3-bit uniform quantiser as a function of the input, x . It also shows the *average* transfer function for the SSR model, with $N = 7$, for various values of noise standard deviation, σ_η , and Gaussian noise. Fig. 4.27(b) shows the variance corresponding to each value of σ_η as a function of x .

4.5.2 Decoding the SSR Output

This description of SSR as a stochastic quantisation of the signal leads naturally to ask whether SSR can be described in terms of conventional quantisation. The answer is yes. Such theory usually considers a quantiser as having two stages—the encoding stage and the decoding stage. The encoding stage can be described in terms of mutual information, for randomly noisy encoding. Thus, as well as describing the interpretation of SSR as an information channel, the current Chapter serves as a description for the *encoding* of a signal in a quantisation scheme. Analysis of SSR in terms of the *decoder* component of a quantiser is left for Chapter 6.

Furthermore, the realisation that SSR can be described as a stochastic quantisation of a signal also leads to the subject of optimal quantisation. Usually a quantisation scheme is considered to be deterministic, and optimising the quantisation requires the optimal selection of threshold values. There is no reason why such optimal threshold selection cannot be considered for the case of noisy thresholds. This is the focus of Chapters 8 and 9.

4.6 Chapter Summary

This Chapter commences with a historical literature review of the initial discovery of suprathreshold stochastic resonance, and all subsequent substantial work published on the topic. We also briefly compare and contrast SSR with other similar work, and discuss why the SSR model can be analysed from an information theoretic perspective.

Next, Section 4.3 comprehensively defines, and discusses the SSR model, and replicates the approach used by Stocks to measure the mutual information between its input and output signals. We also introduce a generic change of probability measure and show how this approach can be used to calculate the mutual information. This change of probability measure results in a new PDF, $Q(\tau)$, which we interpret as being the PDF of the average transfer function of the SSR model, and is used to show that the mutual information is a function of σ for specific matched signal and noise. Exact results for identical signal and noise distributions, and the case of uniform signal and noise, as derived by Stocks are re-derived. Finally, Section 4.3 presents numerical results for seven different matched signal and noise distributions, as well as for various signal distributions subject to Gaussian noise.

Section 4.4 then discusses finding channel capacity for the SSR model. We first present numerical results showing the capacity obtained, and the value of σ which attains capacity, for four specified matched signal and noise cases. We then discuss briefly how the Arimoto-Blahut algorithm can be applied to find an input PDF which achieves capacity, when the input PDF is not specified. The last subsection in Section 4.4 gives a short discussion of how a coding efficiency measure can be defined for SSR, since there is a known finite upper bound on the mutual information.

Finally, Section 4.5 gives an interpretation of the SSR model as being a *stochastic quantiser*, since the independent noise on each threshold acts to randomise the effective threshold values. This point is illustrated by plots showing examples of the values of

such random thresholds, and plots showing the average transfer function for the SSR model, when compared to a conventional quantiser.

4.6.1 Original Contributions for Chapter 4

This chapter includes the following original contributions:

- The first literature review of all previous work on the SSR model.
- The introduction of a generic change of probability measure in the equations used to determine the mutual information through the SSR model, and the realisation that this change of probability measure results in a PDF that describes the average transfer function of the SSR model. This PDF is derived for several cases, for which it is proved that the mutual information is a function of the ratio, σ , rather than a function of the noise variance, and signal variance independently.
- The entropy of $Q(\tau)$ is shown to be equal to the relative entropy between the signal and noise PDFs.
- Numerical calculations of the output entropy, average conditional entropy and mutual information, for a number of signal and noise distributions not previously considered are plotted. These include the Laplacian, logistic, Cauchy, Rayleigh and exponential distributions. SSR is shown to occur in all cases.
- Numerical calculations of the mutual information for mixed signal and noise distributions are plotted. SSR is shown to occur in all cases.
- Channel capacity is found as a function of N for the Gaussian, Laplacian, uniform and logistic cases, as is the value of σ that achieves capacity. The difference between channel capacity and the mutual information at $\sigma = 1$ is calculated, and found to converge towards a constant value with increasing N .
- A comparison of coding efficiency for each signal and noise pair.
- Illustrations of why the signal encoding provided by the SSR model can be interpreted as a *stochastic quantisation*.

4.6.2 Further Work

Possible future work and open questions arising from this chapter might include:

- The use of the Arimoto-Blahut algorithm to find input PDFs that attain channel capacity in the SSR model, subject to the specification of constraints. Such constraints might include a specified variance, combined with a unimodal constraint.
- Investigation of multi-peaked or discretely valued signal and noise distributions.
- More thorough investigation of cases of mixed signal and noise distributions.

This concludes Chapter 4, which introduces the SSR model, and presents results that apply for any value of N . Chapter 5 now examines approximations for the information transmission in the SSR model, under the assumption of a large number of threshold devices.

Chapter 5

Suprathreshold Stochastic Resonance: Large N Encoding

THIS chapter discusses the behaviour of the mutual information and channel capacity in the suprathreshold stochastic resonance model as the number of threshold elements becomes large or approaches infinity. The results in Chapter 4 indicate that the mutual information and channel capacity might converge to simple expressions of N in the case of large N . The current chapter finds that accurate approximations do indeed exist in the large N limit.

5.1 Introduction

Section 4.4 of Chapter 4 presents results for the mutual information and channel capacity through the Suprathreshold Stochastic Resonance (SSR) model shown in Fig. 4.1. Recall that σ is the ratio of the noise standard deviation to the signal standard deviation. For the case of matched signal and noise distributions and a large number of threshold devices, N , the optimal value of σ —that is, the value of σ that maximises the mutual information and achieves channel capacity—appears to asymptotically approach a constant value with increasing N . This indicates that analytical expressions might exist in the case of large N for the optimal noise intensity and channel capacity. We also saw that the channel capacity appeared as if it might become independent of the signal and noise distribution for large N . We now explore these questions and find accurate large N approximations in some specific cases. Firstly however we briefly summarise relevant results given in Chapter 4 and then describe work already completed on this topic by other authors.

5.1.1 Mutual Information

As described in Chapter 4, for N threshold devices, the output of the SSR model is a discretely valued signal, y , that takes on integer values between 0 and N . As explained, y can be considered to be a stochastic encoding of the continuously valued input signal x .

Regardless of the fact that N is allowed to approach infinity in the present chapter, the output can never be considered to be a continuously valued signal, even if it is decoded to a finite range. This fact has its roots in the notion of *countable* and *uncountable* infinities—the set of integers is called *countably infinite*, whereas the set of real numbers is called *uncountably infinite* (Courant and Robbins 1996). An illustration of this notion is to consider the question “how many real numbers are located between the integers m and $m + 1$?” The answer is that there are infinitely many real numbers between m and $m + 1$. By extension, since there are infinitely many integers, there are infinitely more real numbers than there are integers. Furthermore, there are infinitely many real numbers between any given pair of distinct real numbers, regardless of how small the difference between them. By contrast, the distance between any two consecutive integers is finite. Thus, the real numbers are said to be uncountably infinite, whereas the integers are countably infinite in size.

This fact is also the basis for the difference between the entropy of a discrete random variable, which is always non-negative, and the differential entropy of a continuously valued random variable, which may be positive, negative or zero.

Thus, theoretically, quantisation of a continuously valued signal always gives a lossy encoding, even if the *quantisation noise*—that is, the noise introduced into an encoding of a number by quantisation, or truncation—can be made arbitrarily small, or to asymptotically approach zero, by allowing N to approach infinity. This is manifested by a term in the mutual information of the order of $\log_2 N$. This term approaches infinity for large N , and accounts for the fact that the quantisation noise can decrease inversely with N . We will return to this fact in Chapters 6 and 7.

In light of this discussion, it is important that large N approximations to the mutual information between the input and the output of the SSR model do not assume that the output is continuously valued, without rigorous justification.

Before proceeding to examining in detail previous results on this topic in Stocks (2001c) and Hoch *et al.* (2003a), we firstly summarise the relevant results from Chapter 4. As in Chapter 4, this chapter only considers the case of all threshold values being identical, so that all threshold values are equal to θ . We assume throughout this Chapter, apart from Section 5.5.1, that the Probability Density Functions (PDFs) of the signal, $P(x)$, and the noise, $R(\eta)$, are known, share the same distribution, and are even functions about a mean of zero. Thus, we will always set the threshold value to be $\theta = 0$.

Review of key results from Chapter 4

We define $P_y(n)$, $n = 0, \dots, N$ to be the probability mass function of the output signal, y , and $P(n|x)$, $n = 0, \dots, N$ to be the set of transition probabilities giving the probability that the output is $y = n$ given input signal value, x .

The mutual information between the input signal, x , and the output, y , of the SSR model can be written as

$$\begin{aligned} I(x, y) &= H(y) - H(y|x) \\ &= - \sum_{n=0}^N P_y(n) \log_2 P_y(n) - \left(- \int_{-\infty}^{\infty} P(x) \sum_{n=0}^N P(n|x) \log_2 P(n|x) dx \right). \end{aligned} \quad (5.1)$$

5.1 Introduction

The entropy of the output for a given value of the input can be expressed in terms of the probability that any given device is 'on', given x , $P_{1|x}$ as

$$\hat{H}(y|x) = N \left(P_{1|x} \log_2 P_{1|x} + (1 - P_{1|x}) \log_2 (1 - P_{1|x}) \right) + \sum_{n=0}^N P(n|x) \log_2 \binom{N}{n}. \quad (5.2)$$

For even signal and noise PDFs, and all thresholds equal to zero we have the average conditional output entropy as

$$H(y|x) = - \sum_{n=0}^N P_y(n) \log_2 \binom{N}{n} - 2N \int_x P(x) P_{1|x} \log_2 P_{1|x} dx, \quad (5.3)$$

and Eqn. (5.1) reduces to

$$I(x, y) = - \sum_{n=0}^N P_y(n) \log_2 P^*(n) + 2N \int_x P(x) P_{1|x} \log_2 P_{1|x} dx, \quad (5.4)$$

where $P^*(n) = P_y(n) / \binom{N}{n}$ so that $P^*(n) = \int_x P(x) P_{1|x}^n (1 - P_{1|x})^{N-n} dx$.

Recall also that integrations over the signal PDF's support variable, x , can be simplified to expressions defined in terms of integrations over the variable, τ , which exists only between zero and unity, via the transformation $\tau = P_{1|x}$. The resultant expressions are functions of the PDF, $Q(\tau)$, where

$$Q(\tau) = \frac{P(x)}{R(x)} \Big|_{x=F_R^{-1}(\tau)}. \quad (5.5)$$

5.1.2 Literature Review

Stocks' seminal paper (Stocks 2000a) on SSR briefly discusses the scaling of SSR with N , for large N . It states that an exact expression for mutual information can be found at $\sigma = 1$ for identical signal and noise distributions. From this expression, it can be shown that the mutual information scales with $0.5 \log_2(N)$ for large N . Stocks notes that this means that the channel capacity for large N is about half the maximum noiseless channel capacity, since in the absence of noise, the maximum mutual information is the maximum entropy of the output signal, which is $\log_2(N + 1)$. The mathematical details of this result are not presented in Stocks (2000a), but left for Stocks (2001c) and Stocks (2001a).

Stocks (2001c) presents an exact result for the mutual information in the SSR model for the case of uniform signal and noise and $\sigma \leq 1$. He also derives an approximation

to this exact expression that holds in the case of large N , and shows that the mutual information scales with $0.5\sigma \log_2(N)$. Using this expression, Stocks (2001c) derives formulas for the channel capacity for a given N , and the value of σ at which capacity occurs, σ_o , finding that as $N \rightarrow \infty$, $\sigma_o \rightarrow 1$. Thus the main result is that for uniform signal and noise and large N , capacity occurs at $\sigma_o = 1$ and scales with $0.5 \log_2(N)$.

It is shown in Section 4.3.2 of Chapter 4 that the mutual information at $\sigma = 1$ is independent of the signal and noise distribution, provided $P(x) = R(\theta - x)$ for all valid x . This means that the result of Stocks (2001c) showing that the mutual information at $\sigma = 1$ scales with $0.5 \log_2(N)$ holds for any matched signal and noise distributions where $P(x) = R(\theta - x)$. This is also demonstrated in Stocks (2001a).

The only other authors to consider SSR in the large N regime find that, in contrast to uniform signal and noise—as studied in Stocks (2001c)—the channel capacity for Gaussian signal and noise occurs for a value of $\sigma_o \simeq \sqrt{1 - 2/\pi} \simeq 0.603$ (Hoch *et al.* 2003a, Hoch *et al.* 2003b). We have already seen in Section 4.4 that for Gaussian signal and noise the channel capacity appears to occur for σ_o approaching this value for $N = 1000$.

In contrast to Stocks (2001c)—which makes use of an exact expression for the mutual information, and derives a large N approximation by approximating a summation with an integral—Hoch *et al.* (2003a) begins by using a Fisher information based approximation to mutual information. The optimal value of σ is found by an analytical approximation to the stationary point of the resultant expression. Hoch *et al.* (2003b) and Wenning (2004) also give the results of Hoch *et al.* (2003a), but provide more details. In addition, Hoch *et al.* (2003b) gives an extensive investigation into quantifying SSR in populations of spiking neuron models, including a section on SSR and energy constraints. We also study the behaviour of SSR under constraints on the energy available in the model in Chapter 9.

5.1.3 Chapter Structure

The remainder of this chapter is organised as follows. Section 5.2 discusses Stocks' large N approximation to the mutual information at $\sigma = 1$ and Section 5.3 discusses his approximation to the mutual information for uniform signal and noise and $\sigma \leq 1$. Hoch *et al.*'s Fisher information approach to approximating the mutual information for large N and all σ is discussed in Section 5.4, which also provides an alternative and

more accurate derivation of their results. Finding the noise intensity at which channel capacity occurs is the subject of Section 5.5. Several results relevant to this chapter are given in more detail in Appendix C.

5.2 SSR for Large N and $\sigma = 1$

Recall from Section 4.3.2 in Chapter 4 that for the specific case of signal and noise PDFs, and θ , such that $P(x) = R(\theta - x)$, the following exact results are obtained,

$$H(y) = \log_2(N + 1), \quad (5.6)$$

$$H(y|x) = \frac{N}{2 \ln 2} + \frac{1}{N + 1} \sum_{n=1}^N (N + 1 - 2n) \log_2 n, \quad (5.7)$$

and

$$I(x, y) = \log_2(N + 1) - \frac{N}{2 \ln 2} - \frac{1}{N + 1} \sum_{n=2}^N (N + 1 - 2n) \log_2 n. \quad (5.8)$$

If both $P(x)$ and $R(x)$ have the same distribution, then we have $\sigma = 1$. Although, we consider only this situation in this Section, the above results hold for any case of θ , $P(x)$ and $R(\eta)$ such that $P(x) = R(\theta - x)$ for all valid x .

It is shown in Stocks (2001a) that as N approaches infinity, Eqn. (5.8) reduces to the approximation, $I(x, y) \simeq 0.5 \log_2(N + 1)$, and that this implies that a maximum must occur in the mutual information for a nonzero noise distribution (Stocks 2001a). This section examines the derivation of this result in detail, and finds a slightly more accurate expression for the large N mutual information at $\sigma = 1$.

5.2.1 Average Conditional Entropy for Large N and $\sigma = 1$

We begin by finding a simplified version of Eqn. (5.7) that holds for large N . Rather than beginning with Eqn. (5.7) though, we start with Eqn. (5.3). Firstly, it is convenient to express the first term in Eqn. (5.3) without the combinatorial term. Note that for $\sigma = 1$, $P_y(n) = \frac{1}{N+1} \forall n$. This is proven in Section 4.3.2 in Chapter 4—see Eqn. (4.35). After some algebra—see Section C.1 of Appendix C—the following identity is proven,

$$-\sum_{n=0}^N P_y(n) \log_2 \binom{N}{n} = \log_2(N!) - \frac{2}{N+1} \sum_{n=1}^N n \log_2 n. \quad (5.9)$$

Note that Eqn. (5.9) can also be expressed as in the third term of Eqn. (5.8) as

$$-\sum_{n=0}^N P_y(n) \log_2 \binom{N}{n} = \frac{1}{N+1} \sum_{n=2}^N (N+1-2n) \log_2 n, \quad (5.10)$$

which is the expression used in Stocks (2001a). However, here it will prove more convenient to use Eqn. (5.9). Note firstly that this expression still holds for any value of N , not just large N .

We will now see that both terms of Eqn. (5.9) can be simplified by approximations that hold for large N . Firstly, for the $\log_2(N!)$ term, we can make use of Stirling's formula (Spiegel and Liu 1999), which is valid for large N ,

$$N! \sim \sqrt{(2\pi N)} N^N \exp(-N). \quad (5.11)$$

This approximation is particularly accurate if the log is taken of both sides,

$$\log_2(N!) \sim N \log_2 N + 0.5 \log_2 N - \frac{N}{\ln 2} + 0.5 \log_2(2\pi), \quad (5.12)$$

where, for $N = 100$, the absolute error is approximately 1.2×10^{-3} , dropping to about 1.2×10^{-4} for $N = 1000$. Thus, the absolute error is $O(N^{-1})$. This is to be expected, since Stirling's formula results from an asymptotic expansion of the Gamma function, with a second term of $1/(12x)$. The percentage errors are of course far smaller. Hence, this approximation to $\log_2(N!)$ can be used for moderately small N .

Secondly, the sum in the second term of Eqn. (5.9) can be simplified by way of the Euler-Maclaurin summation formula (Spiegel and Liu 1999). Section C.2 of Appendix C shows that

$$\frac{2}{N+1} \sum_{n=1}^N n \log_2 n \simeq N \log_2(N+1) - \frac{N(N+2)}{2 \ln 2(N+1)} + O\left(\frac{\log N}{N}\right). \quad (5.13)$$

Even though Eqn. (5.13) on its own can be reduced further for large N , we are looking for a large N approximation to Eqn. (5.9). Thus, it is safer to firstly subtract Eqn. (5.13) from Eqn. (5.12), before letting N become large. Carrying this out gives

$$\begin{aligned} -\sum_{n=0}^N P_y(n) \log_2 \binom{N}{n} &\simeq 0.5 \log_2 N - N \log_2 \left(1 + \frac{1}{N}\right) - \frac{N}{2 \ln 2} \left(2 - \frac{N+2}{N+1}\right) \\ &\quad + 0.5 \log_2(2\pi) - O\left(\frac{\log N}{N}\right). \end{aligned} \quad (5.14)$$

5.2 SSR for Large N and $\sigma = 1$

Noting that for $|x| < 1$, $\ln(1+x) = x - x^2/2 + x^3/3 - x^4/4 \dots$ (Spiegel and Liu 1999), we have

$$\begin{aligned} N \log_2 \left(1 + \frac{1}{N}\right) &= \frac{1}{\ln 2} \left(1 - 1/(2N) + 1/(3N^2) - 1/(4N^3) \dots\right) \\ &= \frac{1}{\ln 2} + O\left(\frac{1}{N}\right). \end{aligned} \quad (5.15)$$

Using Eqn. (5.15) allows a simplification of Eqn. (5.14), which when combined with Eqn. (4.37), and substituted into Eqn. (5.3), gives the average conditional entropy as

$$\begin{aligned} H(y|x) &= \frac{N}{2 \ln 2} - \sum_{n=0}^N P_y(n) \log_2 \binom{N}{n} \\ &\simeq 0.5 \log_2 N - \frac{1}{\ln 2} + \frac{1}{2 \ln 2} \frac{N}{N+1} + 0.5 \log_2 (2\pi) - O\left(\frac{\log N}{N}\right) \\ &= 0.5 \log_2 N + 0.5 \left(\frac{N}{N+1} - 2\right) \log_2 (e) + 0.5 \log_2 (2\pi) - O\left(\frac{\log N}{N}\right). \end{aligned} \quad (5.16)$$

For large N , the average conditional output entropy of Eqn. (5.16) can be approximated as

$$H(y|x) \simeq 0.5 \log_2 \left(\frac{2\pi N}{e}\right), \quad (5.17)$$

which scales with $0.5 \log_2 N$.

5.2.2 Mutual Information for Large N and $\sigma = 1$

It is demonstrated in Section 4.3.2 of Chapter 4 that for the conditions of this section, $H(y) = \log_2(N+1)$, as given in Eqn. (5.6). Thus, subtracting Eqn. (5.16) from Eqn. (5.6) gives the mutual information as

$$\begin{aligned} I(x, y) &\simeq \log_2(N+1) - 0.5 \log_2 N - 0.5 \left(\frac{N}{N+1} - 2\right) \log_2 (e) \\ &\quad - 0.5 \log_2 (2\pi) + O\left(\frac{\log N}{N}\right) \\ &= 0.5 \log_2 \left(N + 2 + \frac{1}{N}\right) - 0.5 \left(\frac{N}{N+1} - 2\right) \log_2 (e) \\ &\quad - 0.5 \log_2 (2\pi) + O\left(\frac{\log N}{N}\right). \end{aligned} \quad (5.18)$$

Letting N approach infinity in Eqn. (5.18) gives an approximation to the large N mutual information for $\sigma = 1$ as

$$I(x, y) \simeq 0.5 \log_2 \left(\frac{(N+2)e}{2\pi}\right). \quad (5.19)$$

Eqn. (5.19) differs slightly from that stated in Eqn. (7) in Stocks (2001a), which can be written as

$$I(x, y) = 0.5 \log_2 \left(\frac{N+1}{e} \right). \quad (5.20)$$

The explanation of the discrepancy is that Stocks (2001a) uses the Euler-Maclaurin summation formula to implicitly calculate $\log_2(N!)$ in the large N approximation to the average conditional entropy, under the assumption that the remainder terms are not of consequence. It turns out that these terms are of consequence, the reason being that the Bernoulli numbers, B_p , do not decrease with p , other than for the first few terms—see Section C.2 of Appendix C. After the third term, the Bernoulli numbers increase with p . Stirling's approximation for $N!$ used here takes this into account (Abramowitz and Stegun 1972), and therefore gives a more accurate approximation than Stocks (2001a).

The increased accuracy of Eqn. (5.19) can be confirmed by comparing both Eqn. (5.19) and Eqn. (5.20) with the exact expression for $I(x, y)$ of Eqn. (5.8), as N increases. Fig. 5.1 shows the mutual information for $\sigma = 1$ obtained by the exact expression given by Eqn. (5.8) compared with the approximations of Eqn. (5.18)—with the $O\left(\frac{\log N}{N}\right)$ term ignored—Eqn. (5.19) and Eqn. (5.20). Fig. 5.2 shows that the error between the exact expression and the first two approximations approaches zero as N increases, whereas the error between Eqn. (5.8) and Eqn. (5.20) approaches a nonzero constant for large N , of about $0.5 \log_2 \left(\frac{e^2}{2\pi} \right) \simeq 0.117$.

If one was however truly interested in a very large N approximation, the percentage error resulting from this term of course becomes very small. This can be seen as follows. For very large N , Eqn. (5.19) can be simplified to

$$I(x, y) \simeq 0.5 \log_2 \left(\frac{Ne}{2\pi} \right). \quad (5.21)$$

This expression—also shown in Fig. 5.1—is obtained in Section 5.4 by an alternative derivation of an approximation to the mutual information at $\sigma = 1$ for large N —see Eqn. (5.67).

Since Eqn. (5.21) can be written as $I \simeq 0.5 \log_2(N) - 0.6044$, for very large N the constant term becomes insignificant. However due to the very slowly increasing nature of the logarithm function, N needs to be very large for 0.6 bits per sample to be considered very small. For example, if $N = 10^{10}$, the relative error is 3.77 percent. It takes $N \simeq 10^{37}$ for the error to fall beneath one percent!

Part of our interest in a large N approximation to the mutual information is due to the underlying motivation of studying SSR, which is its possible relationship to neural coding. However, although there are approximately $N = 10^{11}$ neurons in the whole human brain, it is unlikely that the number of neurons devoted to a specific task is any more than 1000 (Hoch *et al.* 2003b), in which case the error incurred by dropping the constant term from Eqn. (5.21) is about 14 percent. Since this is quite significant, we suggest that it is appropriate to consider the constant term in Eqn. (5.21) if a quantitative approximation to the mutual information in the SSR model is required. However, the main conclusion from this analysis, as first noted in Stocks (2001a), is that the mutual information scales with $0.5 \log_2(N)$ as N increases—this conclusion remains unchanged whether the constant term is included or not.

From Fig. 5.2 it appears that surprisingly—given that Eqn. (5.19) was derived from Eqn. (5.18)—Eqn. (5.19) gives a better approximation than Eqn. (5.18) for small N . This can be seen more clearly in Fig. 5.3, which plots the difference between the absolute errors of these two equations. This discrepancy is due to the discarded terms that make these equations valid for large N . Two of the discarded terms from Eqn. (5.18) cancel each other out for small N , which makes Eqn. (5.19) misleadingly more accurate. For $N > 6$ the true nature of the approximations starts to show, with Eqn. (5.18) being more accurate than Eqn. (5.19), as indicated by the negative value of the plot. As expected though, Fig. 5.3 shows that the difference between each equation approaches zero as N gets larger.

In summary, this section has derived a large N expression for the mutual information at $\sigma = 1$, given by Eqn. (5.19), which is very accurate—an error of less than 0.01 bits per sample, even for N as low as 20. As N gets larger, the error converges towards zero. This is an improvement on the approximation given in Stocks (2001a), for which the error converges towards $0.5 \log_2\left(\frac{e^2}{2\pi}\right) \simeq 0.117$ bits per sample for large N .

For larger N , Eqn. (5.19) can be simplified to Eqn. (5.21). For example, if N is of the order of 1000, the error in Eqn. (5.21) is of the order of 10^{-3} .

Having now derived the specific case of an accurate asymptotic expression for the mutual information for $\sigma = 1$, as carried out in Stocks (2001a), the next logical step is to try to obtain an asymptotic expression for arbitrary σ . This is the focus of the next two sections. Firstly, in Section 5.3 a result is given for the specific case of uniform signal and noise, and then Section 5.4 gives results for arbitrary matched signal and noise distributions.

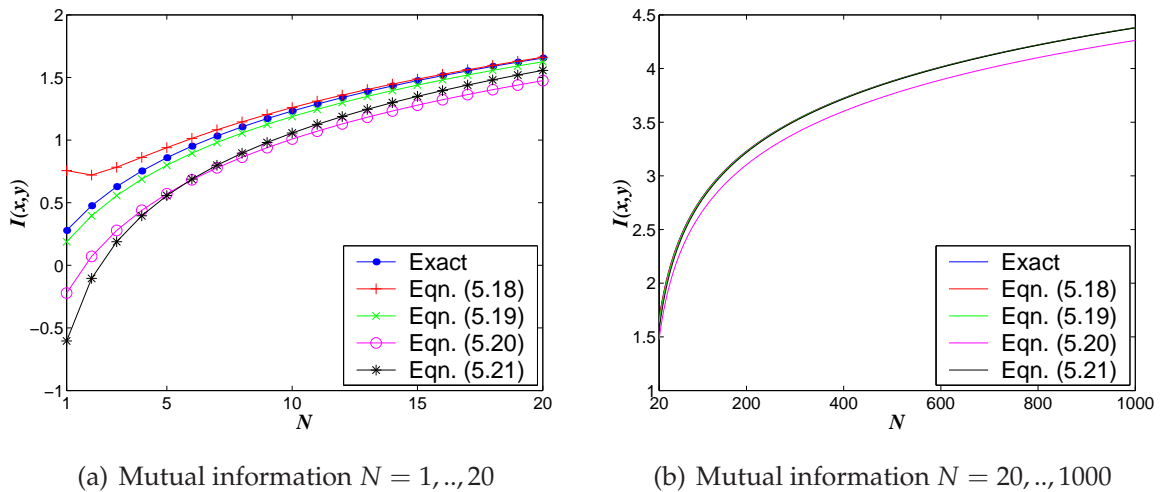


Figure 5.1. Mutual information for $\sigma = 1$ and large N . This figure shows the exact mutual information for noise intensity, $\sigma = 1$, given by Eqn. (5.8), against increasing array size, N , as well as the four approximations given by Eqns. (5.18), (5.19), (5.20) and (5.21). The plot has been divided into two scales to indicate how well the approximations work for small N and large N . In Fig. 5.1(a), the line plots are used as an aid to the eye, with the marker showing the mutual information at the integer values of N between 1 and 20. It appears from this plot that Eqn. (5.19) gives the best approximation for small N . Fig. 5.1(b) shows that for larger N , the approximations of Eqns. (5.18), (5.19), (5.21) are indistinguishable by eye, whereas the approximation of Eqn. (5.20) clearly gives a larger error. The actual errors in each approximation are shown in Fig. 5.2.

5.2 SSR for Large N and $\sigma = 1$

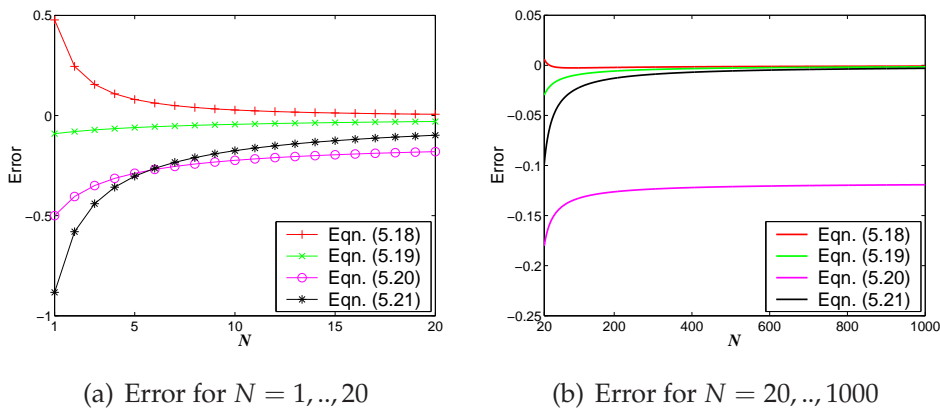


Figure 5.2. Error in large N approximations to $I(x, y)$ for $\sigma = 1$. Plot showing the error between the exact expression for the mutual information at noise intensity, $\sigma = 1$, given by Eqn. (5.8), and the various large N approximations to the mutual information at $\sigma = 1$, given by Eqns. (5.18), (5.19), (5.20) and (5.21). Eqns. (5.18) and (5.19) give the smallest error for small N , but once N gets close to 1000, all approximations except Eqn. (5.20) converge towards an error of zero. By contrast, Eqn. (5.20) converges to a nonzero error of $-0.5 \log_2 (0.5e^2 / \pi) \simeq -0.117$.




Figure 5.3. Error comparison for large N approximations to $I(x, y)$ for $\sigma = 1$. The two large N approximations that give the smallest error are Eqns. (5.18) and (5.19). This figure shows how for very small array size ($1 \leq N \leq 6$), Eqn. (5.19) gives the smallest absolute error, and for $N > 6$, Eqn. (5.18) gives the smallest absolute error. This discrepancy is due to the discarded terms that make these equations valid for large N . In Eqn. (5.19) two discarded terms cancel each other out for small N . For $6 \leq N \leq 20$, the difference in the errors gets larger, but for $N > 20$, gets smaller again, until as N gets larger than about 250, the difference between the two equations starts to converge asymptotically towards zero. From Fig. 5.2, the actual error in both these approximations approaches zero for large N .

5.3 $I(x, y)$ for Large N , Uniform Signal & Noise & $\sigma \leq 1$

Section 4.3.2 in Chapter 4 included the derivation given in Stocks (2001c) of an exact expression for the mutual information for uniform signal and noise and $\sigma \leq 1$. The relevant expressions, repeated here, are

$$P_y(n) = \begin{cases} \frac{\sigma}{N+1} + \frac{1}{2} - \frac{\sigma}{2} & \text{for } n = 0, N \\ \frac{\sigma}{N+1} & \text{for } 1 \leq n \leq N-1, \end{cases} \quad (5.22)$$

$$H(y) = - \left(\frac{2\sigma}{N+1} + 1 - \sigma \right) \log_2 \left(\frac{\sigma}{N+1} + \frac{1}{2} - \frac{\sigma}{2} \right) - \frac{(N-1)\sigma}{N+1} \log_2 \left(\frac{\sigma}{N+1} \right), \quad (5.23)$$

$$H(y|x) = \frac{\sigma}{N+1} \sum_{n=2}^N (N+1-2n) \log_2 n + \frac{N\sigma}{2 \ln 2}, \quad (5.24)$$

$$I(x, y) = - \left(\frac{2\sigma}{N+1} + 1 - \sigma \right) \log_2 \left(\frac{\sigma}{N+1} + \frac{1}{2} - \frac{\sigma}{2} \right) - \frac{(N-1)\sigma}{N+1} \log_2 \left(\frac{\sigma}{N+1} \right) - \frac{\sigma}{N+1} \sum_{n=2}^N (N+1-2n) \log_2 n - \frac{N\sigma}{2 \ln 2}. \quad (5.25)$$

In addition, Stocks (2001c) shows that a large N approximation can be made to these equations. His derivation is repeated here, and as was the case in Section 5.2, we find a similar approximation, with improved accuracy.

Large N mutual information

For large N , Eqn. (5.23) reduces to

$$H(y) \simeq \sigma \log_2 (N+1) + (1-\sigma)(1 - \log_2 (1-\sigma)) - \sigma \log_2 (\sigma). \quad (5.26)$$

Furthermore, the average conditional entropy is simply σ multiplied by Eqn. (5.7). Thus, using the large N approximation to Eqn. (5.7) given by Eqn. (5.17), we have

$$H(y|x) \simeq \frac{\sigma}{2} \log_2 \left(\frac{2\pi N}{e} \right). \quad (5.27)$$

Hence, using the same arguments as for the $\sigma = 1$ case of Section 5.2 the mutual information approximation for large N is

$$I(x, y) \simeq \frac{\sigma}{2} \log_2 \left(\frac{(N+2)e}{2\pi} \right) + (1-\sigma)(1 - \log_2 (1-\sigma)) - \sigma \log_2 (\sigma). \quad (5.28)$$

Thus, the mutual information is equal to σ times the $\sigma = 1$ mutual information plus the term $(1-\sigma)(1 - \log_2 (1-\sigma)) - \sigma \log_2 (\sigma)$, which reduces to zero for $\sigma = 1$.

Due to the same differences in the derivation as stated in Section 5.2, Eqn. (5.28) differs from Eqn. (9) given in Stocks (2001c), which can be written as

$$I(x, y) = \frac{\sigma}{2} \log_2 \left(\frac{(N+1)}{e} \right) + (1-\sigma)(1 - \log_2(1-\sigma)) - \sigma \log_2(\sigma), \quad (5.29)$$

so that the error is exactly that described in the $\sigma = 1$ case, only here it is a function of σ , i.e. $\frac{\sigma}{2} \log_2 \left(\frac{e^2}{2\pi} \right)$.

This error can be seen in Fig. (3) of Stocks (2001c). To illustrate this, and the accuracy of Eqn. (5.28), Fig. (3) of Stocks (2001c) is reproduced here as Fig. 5.4, with the mutual information given by Eqn. (5.28) superimposed. It is clear that Eqn. (5.28) gives a better approximation. Fig. 5.5 shows that the error between Eqn. (5.28) and the exact mutual information of Eqn. (5.25) asymptotically approaches zero for all $\sigma \leq 1$ as N increases.

Large N channel capacity

Differentiating Eqn. (5.28) with respect to σ and setting to zero obtains the extrema of the large N approximation of the mutual information. The result is a single extremum at

$$\sigma_o = \frac{\sqrt{(N+2)}}{\sqrt{(N+2)} + \sqrt{\left(\frac{8\pi}{e}\right)}}. \quad (5.30)$$

Taking the second derivative of Eqn. (5.28) results in an expression that is always negative and hence the extremum is a maximum. The value of the mutual information at this peak, can be found as the very simple formula

$$I_o(x, y) = 1 - \log_2(1 - \sigma_o) = \log_2 \left(2 + \sqrt{\frac{(N+2)e}{2\pi}} \right). \quad (5.31)$$

Clearly, as N becomes very large, the optimal value of σ given by Eqn. (5.30) approaches unity, and the mutual information at this σ given by Eqn. (5.31) approaches $0.5 \log_2((N+2)e/(2\pi))$, which agrees with Eqn. (5.28) at $\sigma = 1$, and Eqn. (5.17).

Recall that an exact expression for σ_o was given by Eqn. (4.59) in Section 4.3.2 of Chapter 4. Eqn. (5.30) and Eqn. (4.59) can be seen to be closely related when Eqn. (5.30) is rewritten as

$$\sigma_o = \frac{N+2}{N+2 + \sqrt{\left(\frac{8\pi(N+2)}{e}\right)}}. \quad (5.32)$$

The error between Eqn. (5.32) and Eqn. (4.59) also converges to zero for large N , again validating our approximations.

Stocks (2001c) also gives expressions for the optimal σ and the corresponding mutual information. Again, these are slightly different to Eqns (5.30) and (5.31), due to the slightly inaccurate terms in the large N approximation to the average conditional entropy. However the important qualitative result remains the same, which is that the maximum mutual information scales with half the logarithm of N , and the value of σ which achieves this asymptotically approaches unity. Fig. 5.4 illustrates this behaviour of σ_o and the corresponding $I_o(x, y)$, as N increases from 1 to 8000.




Figure 5.4. Mutual information for uniform signal and noise and large N . Plot showing mutual information against noise intensity, $\sigma \in [0, 1]$, for uniform signal and noise. The line plot shows the exact mutual information calculated using Eqn. (5.25). Plus signs are the large N approximation of Eqn. (5.28) and dots show the slightly less accurate large N approximation of Eqn. (5.29), given in Stocks (2001c). The dots plot the curve of the optimal value of σ obtained from Eqn. (5.30) against the corresponding mutual information of Eqn. (5.31), for N between 1 and 8000. As N becomes larger, it is clear that Eqn. (5.28) gives a more accurate approximation to Eqn. (5.25) than Eqn. (5.29).




Figure 5.5. Error in mutual information for uniform signal and noise and large N . Plot showing the difference between the exact mutual information for uniform signal and noise given by Eqn. (5.25) and the approximation given by Eqn. (5.28), as the array size, N , increases, for various values of noise intensity, σ . The error can be seen to approach zero as N increases, for each value of σ .

5.4 Mutual Information for Large N and Arbitrary σ

The previous two subsections give asymptotic results for (i) $\sigma = 1$ for arbitrary distributions, and (ii) for $\sigma \leq 1$ for uniform signal and noise. This section attempts to find an expression for the mutual information for arbitrary distributions and all σ . The $\sigma = 1$ result can be used to verify that any new expressions are correct at $\sigma = 1$.

5.4.1 The Gaussian Approximation to the Binomial Distribution

In the SSR model, for each value of the input, x , the probability that the output is $y = n$ is given by the binomial formula, as given by Eqn. (4.9), repeated here as

$$P(n|x) = \binom{N}{n} P_{1|x}^n (1 - P_{1|x})^{N-n}. \quad (5.33)$$

Consider only one value of x . For this value, $P(n|x)$ is a function of n . As N becomes large, the binomial distribution is known to approach a Gaussian distribution, with the same mean and variance as the binomial distribution (Kreyszig 1988), provided the mean of the binomial is sufficiently large. The mean is given by $NP_{1|x}$ and the variance is $NP_{1|x}(1 - P_{1|x})$. Thus, provided $0 \ll NP_{1|x} \ll N$,

$$P(n|x) \simeq \frac{1}{\sqrt{2\pi NP_{1|x}(1 - P_{1|x})}} \exp\left(-\frac{(n - NP_{1|x})^2}{2NP_{1|x}(1 - P_{1|x})}\right). \quad (5.34)$$

This approximation breaks down when $P_{1|x}$ is close to zero or unity, in which case, $P(n|x)$ can be approximated by the Poisson distribution, or the Edgeworth series approximation (Harrington 1955). However here we will use Eqn. (5.34), and find that it is valid for our needs, since for infinite support PDFs such as the Gaussian, logistic and Laplacian distributions, when $P_{1|x}$ is close to zero or unity, $P(x)$ is also very close to zero. Note that even though Eqn. (5.34) is a Gaussian function, since n is still discretely valued, Eqn. (5.34) still represents a discrete probability mass function rather than a PDF.

The large N approximation given by Eqn. (5.34) allows us to obtain a general large N approximation for the average conditional entropy. This is the focus of the following section.

5.4.2 Conditional Entropy for Large N

Using the binomial approximation

Recall in Chapter 4 that we used the notation $\hat{H}(y|x)$ to denote the conditional output entropy for a given value of x . The *average* conditional output entropy is then $H(y|x) = \int_x P(x)\hat{H}(y|x)dx$, where

$$\hat{H}(y|x) = - \sum_{n=0}^N P(n|x) \log_2 (P(n|x)). \quad (5.35)$$

Taking the log of the Gaussian approximation to the binomial given by Eqn. (5.34), and substituting into Eqn. (5.35) gives

$$\begin{aligned} \hat{H}(y|x) &\simeq - \sum_{n=0}^N P(n|x) \left(-0.5 \log_2 (2\pi NP_{1|x}(1 - P_{1|x})) - \frac{(n - NP_{1|x})^2}{2 \ln(2) NP_{1|x}(1 - P_{1|x})} \right) \\ &= 0.5 \log_2 (2\pi NP_{1|x}(1 - P_{1|x})) \\ &\quad + \frac{1}{2 \ln(2) NP_{1|x}(1 - P_{1|x})} \sum_{n=0}^N P(n|x)(n - NP_{1|x})^2 \\ &= 0.5 \log_2 (2\pi NP_{1|x}(1 - P_{1|x})) + \frac{1}{2 \ln(2) NP_{1|x}(1 - P_{1|x})} \text{var}[y|x] \\ &= 0.5 \log_2 (2\pi NP_{1|x}(1 - P_{1|x})) + \frac{1}{2 \ln(2)} \\ &= 0.5 \log_2 (2\pi e NP_{1|x}(1 - P_{1|x})), \end{aligned} \quad (5.36)$$

where we have used the fact that the variance of the binomial distribution, $P(n|x)$, is $NP_{1|x}(1 - P_{1|x})$, and the mean is $NP_{1|x}$.

Multiplying both sides of Eqn. (5.36) by $P(x)$ and integrating over all x gives

$$\begin{aligned} H(y|x) &\simeq 0.5 \log_2 (2\pi eN) + 0.5 \int_{x=-\infty}^{\infty} P(x) \log_2 (P_{1|x}(1 - P_{1|x})) dx \\ &= 0.5 \log_2 (2\pi eN) + \int_{x=-\infty}^{\infty} P(x) \log_2 P_{1|x} dx, \end{aligned} \quad (5.37)$$

since $P(x)$ and $R(\eta)$ are even functions about means of zero.

Eqn. (5.37) can also be written in terms of $Q(\tau)$. Recall from Chapter 4 that if we let $\tau = P_{1|x}$, then $P(x)dx = Q(\tau)d\tau$. Making this change of variable in Eqn. (5.37),

$$H(y|x) = 0.5 \log_2 (2\pi eN) + \int_{\tau=0}^{\tau=1} Q(\tau) \log_2 \tau d\tau. \quad (5.38)$$

Recalling that in general $Q(\tau)$ is a PDF, the integral in Eqn. (5.38) is the expected value of $\log_2(\tau)$ with respect to the PDF $Q(\tau)$.

Verification at $\sigma = 1$

Using the large N approximation to $H(y|x)$ derived in Section 5.2 for the specific case of $\sigma = 1$, we can verify Eqn. (5.38) for the case of $\sigma = 1$. At $\sigma = 1$, $Q(\tau) = 1$ and therefore $\int_{\tau=0}^{\tau=1} Q(\tau) \log_2(\tau) d\tau = -\log_2(e)$ and Eqn. (5.38) reduces to

$$H(y|x) = 0.5 \log_2 \left(\frac{2\pi N}{e} \right), \quad (5.39)$$

which agrees precisely with Eqn. (5.17).

Hoch's approach

An alternative approach, taken in Hoch *et al.* (2003a) and Hoch *et al.* (2003b), which does not explicitly use the Gaussian approximation to the binomial, is to use the Euler-Maclaurin summation formula (Spiegel and Liu 1999), to approximate $\hat{H}(y|x)$ as an integral for large N as

$$\hat{H}(y|x) = - \int_{n=0}^N P(n|x) \log_2(P(n|x)) dn. \quad (5.40)$$

This is the entropy of a continuous random variable defined on the support $n \in [0, N]$. If this random variable has a variance $\sigma_{y|x}^2$ then the entropy is less than or equal to the variance of a Gaussian random variable with the same variance. This is due to the well known result that the maximum entropy distribution for continuous random variables under a power constraint—that is, a specified variance in this case—is the Gaussian distribution (Cover and Thomas 1991). Hence, we have

$$\hat{H}(y|x) \leq 0.5 \log_2(2\pi e \sigma_{y|x}^2). \quad (5.41)$$

The variance in this case is again the binomial variance of $\text{var}[y|x] = NP_{1|x}(1 - P_{1|x})$. Therefore

$$\hat{H}(y|x) \leq 0.5 \log_2(2\pi e NP_{1|x}(1 - P_{1|x})). \quad (5.42)$$

Multiplying both sides of Inequality (5.42) by $P(x)$ and integrating over all x leaves

$$\begin{aligned} H(y|x) &\leq 0.5 \log_2(2\pi e N) + 0.5 \int_x P(x) \log_2(P_{1|x}(1 - P_{1|x})) dx \\ &= 0.5 \log_2(2\pi e N) + \int_x P(x) \log_2 P_{1|x} dx. \end{aligned} \quad (5.43)$$

This is identical to Eqn. (5.37) except that (5.43) is an inequality. The equality given by the Gaussian approximation to the binomial of Eqn. (5.37) will only hold for values of x for which $P(n|x)$ is exactly Gaussian.

A small N refinement

A refinement to the upper bound on $H(y|x)$ given by Eqn. (5.43) may be obtained for small N . Since $P(n|x)$ is actually a discrete distribution for a given x , we can make use of a formula for the differential entropy bound on discrete entropy (Cover and Thomas 1991), to get

$$\hat{H}(y|x) \leq 0.5 \log_2 \left(2\pi e \left(\text{var}[y|x] + \frac{1}{12} \right) \right). \quad (5.44)$$

This leads to

$$\hat{H}(y|x) \leq 0.5 \log_2 (2\pi e N P_{1|x} (1 - P_{1|x})) + 0.5 \log_2 \left(1 + \frac{1}{12 N P_{1|x} (1 - P_{1|x})} \right). \quad (5.45)$$

Thus, using the fact that $\ln(1+x) = x - x^2/2 + x^3/3 - x^4/4 \dots$ for $|x| \leq 1$, for values of x such that $P_{1|x}(1 - P_{1|x}) \geq 1/12N$,

$$\hat{H}(y|x) \leq 0.5 \log_2 (2\pi e N P_{1|x} (1 - P_{1|x})) + \frac{1}{24 \ln(2) N P_{1|x} (1 - P_{1|x})} + O\left(\frac{1}{N^2}\right). \quad (5.46)$$

The result of multiplying Inequality (5.46) by $P(x)$, and integrating over all x , reduces for large N to the bound given by Inequality (5.42).

Numerical verification

Consider the exact conditional entropy, $\hat{H}(y|x)$, given by Eqn. (5.35) and the approximation given by Eqn. (5.36). These are both functions of N and $P_{1|x}$. However, since $P_{1|x}$ is always between zero and unity, we do not need to specify a noise PDF to compare the approximation to the exact formula—we need only plot both as a function of $P_{1|x} \in [0, 1]$. Fig. 5.6 shows the absolute error between Eqns. (5.35) and (5.36), calculated numerically for $P_{1|x} \in [0.01, 0.99]$, for various values of N . It is clear from Fig. 5.6 that the absolute error decreases with increasing N , but is larger for $P_{1|x}$ near zero and unity. This illustrates the caveat that the Gaussian approximation to the binomial requires $0 \ll N P_{1|x} \ll N$ to be sufficiently large for high accuracy. Fig. 5.6 verifies this, by clearly showing that the absolute error in the approximation decreases as $P_{1|x}$ gets closer to 0.5 and as N gets larger.

We will see however, that the inaccuracy for $N P_{1|x} \rightarrow 0$ and $N P_{1|x} \rightarrow N$ becomes more important for $\sigma \ll 1$, that is, when the variance of the noise is much smaller than the variance of the signal. Assuming that $P(x)$ and $R(x)$ have long tails, such as in the

Gaussian distribution, if $\sigma \ll 1$ then $P_{1|x}$ gets very close to zero or unity for quite a large range of highly probable x values. Thus, we can expect that the accuracy of the approximation to $H(y|x)$ decreases for smaller σ .

This is indeed the case, as can be seen from Fig. 5.7, which shows plots of the approximation to $H(y|x)$ given by Eqn. (5.37), against increasing σ , compared with the exact $H(y|x)$ of Eqn. (5.3). These results were calculated numerically for the cases of Gaussian signal and noise, Laplacian signal and noise, and logistic signal and noise, and $\sigma = 0.2, 0.3, \dots, 1.6$. It is clear from Fig. 5.7 that the approximation to $H(y|x)$ is highly accurate for $\sigma > 0.7$, and although less accurate for smaller σ , the accuracy increases with increasing N .




Figure 5.6. Error in the large N approximation to $\hat{H}(y|x)$. This figure shows the absolute error between the exact conditional entropy—that is, the conditional entropy of y , given a particular value of x — $\hat{H}(y|x)$, of Eqn. (5.35), and the large N approximation of Eqn. (5.36), and demonstrates that it decreases with increasing N . The data were calculated numerically for various values of N for $P_{1|x} \in [0.01, 0.99]$. The error clearly gets larger for $P_{1|x} \rightarrow 0$ and $P_{1|x} \rightarrow 1$. As the absolute error gets very large at 0 and 1, the data were not calculated at these values.

Having obtained a large N approximation to the average conditional entropy that, unlike the result of Section 5.2 is valid for σ other than $\sigma = 1$, we now wish to find a similar approximation for the output entropy.

5.4 Mutual Information for Large N and Arbitrary σ




Figure 5.7. Approximate $H(y|x)$ for large N . This figure shows the exact average conditional output entropy, $H(y|x)$, of Eqn. (5.3) and the approximate average conditional entropy of Eqn. (5.37), for Gaussian signal and noise, Laplacian signal and noise, and logistic signal and noise. The exact $H(y|x)$ is shown with a thin solid line, while the approximate $H(y|x)$ is shown with circles and a thicker solid line to interpolate between values of noise intensity, σ . Clearly, the approximation is very accurate for $\sigma > 0.7$ in all cases. The accuracy for $\sigma < 0.7$ decreases with decreasing σ , but increases for increasing N .

5.4.3 Output Distribution and Entropy for Large N

This section aims to find a large N approximation to the output entropy. As discussed in the introduction to this chapter, since the output is discretely valued, the output entropy is that of a discrete random variable, regardless of how large N is. Thus, we require a large N approximation to the output probability mass function, $P_y(n)$.

Although the output of the SSR model is a discrete random variable, in the presence of noise the expected value of the output of the SSR array, y , given x , is a continuously valued variable, giving an *average* transfer function, $\bar{y}(x)$. Since $E[y|x] = NP_{1|x}$, the average transfer function for noise distributions with an even PDF about a mean of zero is $\bar{y}(x) = NF_R(x)$.

If we can find the PDF of the average transfer function, $P(\bar{y})$, for the continuous random variable, \bar{y} , then for large N we can use this PDF to approximate the actual output discrete probability mass function, $P_y(n)$.

PDF of the average transfer function

The PDF of the average transfer function can be easily derived from the input signal's PDF by a transformation from coordinate x , with PDF $P(x)$, to coordinate \bar{y} , with PDF $P(\bar{y})$.

In general, a transformation from coordinate a with PDF $P_a(a)$ to coordinate b with PDF $P_b(b)$, where $b = f(a)$ is a differentiable and invertible function, is given by (Fry 1928)

$$P_b(b) = P_a(f^{-1}(b)) \left| \frac{df^{-1}(b)}{db} \right|. \quad (5.47)$$

Thus, in this case we use the average transfer function, $\bar{y}(x)$, so that the inverse transfer function is $x = F_R^{-1} \left(\frac{\bar{y}}{N} \right)$, and the result for $P(\bar{y})$ is

$$\begin{aligned} P(\bar{y}) &= P \left(F_R^{-1} \left(\frac{\bar{y}}{N} \right) \right) \left| \frac{dF_R^{-1} \left(\frac{\bar{y}}{N} \right)}{d\bar{y}} \right| = P \left(F_R^{-1} \left(\frac{\bar{y}}{N} \right) \right) \left(\frac{1}{NR(x)} \right) \\ &= \frac{P \left(F_R^{-1} \left(\frac{\bar{y}}{N} \right) \right)}{NR \left(F_R^{-1} \left(\frac{\bar{y}}{N} \right) \right)} = \frac{P \left(F_R^{-1} \left(P_{1|x} \right) \right)}{NR \left(F_R^{-1} \left(P_{1|x} \right) \right)} \\ &= \frac{P(x)}{NR(x)}. \end{aligned} \quad (5.48)$$

5.4 Mutual Information for Large N and Arbitrary σ

Recall from Section 4.3.2 in Chapter 4 the definition of the PDF $Q(\tau)$, where $\tau = P_{1|x} \in [0, 1]$. It is clear that Eqn. (5.48) can be rewritten as

$$P(\bar{y}) = \frac{Q\left(\frac{\bar{y}}{N}\right)}{N}, \quad (5.49)$$

which is a PDF for the continuously valued variable, \bar{y} . Unlike $Q(\tau)$, which has support $[0, 1]$, the support of $P(\bar{y})$ is $[0, N]$.

If we make a change of probability measure so that $\tau = \frac{\bar{y}}{N}$, Eqn. (5.49) becomes

$$P(\tau) = Q(\tau). \quad (5.50)$$

This agrees with the discussion about $Q(\tau)$ in Chapter 4 where it was hypothesised that $Q(\tau)$ is the PDF of the continuous random variable given by $E[y|x]$.

Approximating the output distribution

For large N , we can assume that $\bar{y}(x) \rightarrow n$ and hence the PDF of Eqn. (5.49) satisfies

$$P(\bar{y})d\bar{y} \simeq P_y(n). \quad (5.51)$$

Thus, the PDF $P(\bar{y})$ approximates the discretely valued probability mass function,

$$P_y(n) \simeq \frac{Q\left(\frac{n}{N}\right)}{N}. \quad (5.52)$$

An alternative approach is to examine the large N behavior of the Gaussian approximation to $P(n|x)$ under a change of probability measure. Firstly, upon substituting for $\tau = P_{1|x}$, Eqn. (5.34) can be rewritten as

$$\begin{aligned} P(n|x) &\simeq \frac{1}{\sqrt{2\pi s^2}} \exp\left(-\frac{(\tau - \frac{n}{N})^2}{2s^2}\right) \frac{1}{N} \\ &= \frac{P(n|\tau)}{N}, \end{aligned} \quad (5.53)$$

where $s^2 = \frac{\tau(1-\tau)}{N}$. The term $P(n|\tau)$ is a function defined on $\tau \in [0, 1]$, and hence, it cannot be Gaussian. However, consider $0 \ll n \ll N$. In this region, $P(n|\tau)$ approaches zero for τ near zero or unity and hence $P(n|\tau)$ is approximately a Gaussian PDF with a variance that decreases with $\frac{1}{N}$. Such a Gaussian with a variance that approaches zero can be approximated by a delta function located at the mean of the Gaussian and thus,

$$P(n|\tau) \simeq \delta\left(\tau - \left(\frac{n}{N}\right)\right). \quad (5.54)$$

Thus, $P(n|\tau) = 1$ iff $\tau = n/N$. This means $P(n|x)$ approaches $1/N$ at $x = F_R^{-1}(\tau) = F_R^{-1}(\frac{n}{N})$. This value of x is the maximum likelihood value of the output, given x , and can be easily shown—see Eqn. (4.12)—without recourse to large N , to be the mode of $P(n|x)$ for each n .

Using the fact that $P_y(n) = \int_x P(x)P(n|x)dx$, changing variable from x to τ gives $P_y(n) \simeq \int_{\tau=0}^{\tau=1} Q(\tau)P(n|\tau)\frac{1}{N}d\tau$, which can be simplified by substitution of Eqn. (5.54) to

$$\begin{aligned} P_y(n) &\simeq \frac{1}{N} \int_{\tau=0}^{\tau=1} Q(\tau)\delta\left(\tau - \frac{n}{N}\right) d\tau \\ &= \frac{Q\left(\frac{n}{N}\right)}{N}, \end{aligned} \quad (5.55)$$

which verifies Eqn. (5.52). As with the Gaussian approximation to the binomial, we can expect that this approximation gets less accurate for values of n near 0 and N .

Numerical verification of large N approximation to $P_y(n)$

We have for the Gaussian, Laplacian and logistic cases, that when the noise intensity, $\sigma > 1$, $Q(0) = Q(1) = 0$, whereas for $\sigma < 1$, we have $Q(0) = Q(1) = \infty$. From Eqn. (5.52), this would mean that $P_y(0)$ and $P_y(N)$ are either zero or infinite. However, for finite N , there is some finite nonzero probability that all output states are on or off. Indeed, at $\sigma = 1$, we know that $P_y(n) = \frac{1}{N+1} \forall n$, and at $\sigma = 0$, $P_y(0) = P_y(N) = 0.5$. Furthermore, for finite N , the approximation of Eqn. (5.55) does not guarantee that $\sum_{n=0}^N P_y(n) = 1$. A little trial and error finds that two simple adjustments to the approximation deals with these problems. The two different cases of $\sigma < 1$ and $\sigma \geq 1$ require different adjustments to achieve a high accuracy that increases with increasing N .

Firstly, for $\sigma \geq 1$, we adjust Eqn. (5.52) to become

$$P'_y(n) = \frac{Q\left(\frac{n+1}{N+2}\right)}{\sum_{m=0}^N Q\left(\frac{m+1}{N+2}\right)} \quad n = 0, \dots, N. \quad (5.56)$$

Eqn. (5.56) ensures $\sum_{n=0}^N P'_y(n) = 1$ and also ensures that $P_y(0)$ and $P_y(N)$ are nonzero for finite N as required, while remaining identical to Eqn. (5.52) in the large N limit.

Secondly, for $\sigma < 1$ we adjust Eqn. (5.52) to become

$$P'_y(n) = \frac{Q\left(\frac{n}{N}\right)}{N} \quad n = 1, \dots, N-1. \quad (5.57)$$

5.4 Mutual Information for Large N and Arbitrary σ

Recall that the approximation is expected to be less valid for n near 0 and N for smaller σ . Hence, to increase the accuracy of our approximation, and to ensure $P_y(n)$ forms a valid probability mass function, we define our approximation as

$$P'_y(n) = \begin{cases} \frac{Q\left(\frac{n}{N}\right)}{N} & \text{for } n = 1, \dots, N-1 \\ 0.5 \left(1 - \sum_{m=1}^{N-1} \frac{Q\left(\frac{m}{N}\right)}{N}\right) & \text{for } n = 0, n = N. \end{cases} \quad (5.58)$$

For $\sigma < 1$, $P_y(0)$ and $P_y(N)$ are the most likely values of y . Eqn. (5.56) ensures that $P_y(0)$ and $P_y(N)$ are reasonably accurate.

Fig. 5.8 shows that the approximation to $P_y(n)$ is highly accurate for N as small as 63, for Gaussian signal and noise, and three values of σ . Similar results can be obtained for the Laplacian or logistic cases. Note from Fig. 5.8(a), where $\sigma = 0.4$, how $P_y(0)$ and $P_y(N)$ are the most inaccurate approximations. This is not surprising, given that we expected the approximation to decrease in accuracy for small σ , and n close to 0 or N .

The approximations to the output probability mass function given by Eqns. (5.56) and (5.58) can also be used to approximate the output entropy as

$$H(y) \simeq - \sum_{n=0}^N P'_y(n) \log_2 \left(P'_y(n) \right). \quad (5.59)$$

The result of this is shown in Fig. 5.9, where the absolute error between Eqn. (5.59) and the exact output entropy, calculated numerically, is shown for the cases of Gaussian signal and noise, Laplacian signal and noise, and logistic signal and noise. The accuracy can be seen to increase as N increases for all σ . The approximation is more accurate near $\sigma = 1$.

Consider the approximation for $P'_y(n)$ and $\sigma < 1$ given by Eqn. (5.58). Let $A = \frac{1}{N} \sum_{n=1}^{N-1} Q\left(\frac{n}{N}\right)$. Then the output entropy of Eqn. (5.59) is

$$\begin{aligned} H(y) &= -\frac{1}{N} \sum_{n=1}^{N-1} Q\left(\frac{n}{N}\right) \log_2 \left(\frac{Q\left(\frac{n}{N}\right)}{N} \right) - (1-A) \log_2 (0.5(1-A)) \\ &= -\frac{1}{N} \sum_{n=1}^{N-1} Q\left(\frac{n}{N}\right) \log_2 \left(Q\left(\frac{n}{N}\right) \right) + A \log_2 N + (1-A) (1 - \log_2 (1-A)). \end{aligned} \quad (5.60)$$

Eqn. (5.60) can also be expressed as

$$H(y) = -\frac{1}{N} \sum_{n=1}^{N-1} Q\left(\frac{n}{N}\right) \log_2 \left(Q\left(\frac{n}{N}\right) \right) + \log_2 N - 2P_y(0) \log_2 (NP_y(0)). \quad (5.61)$$

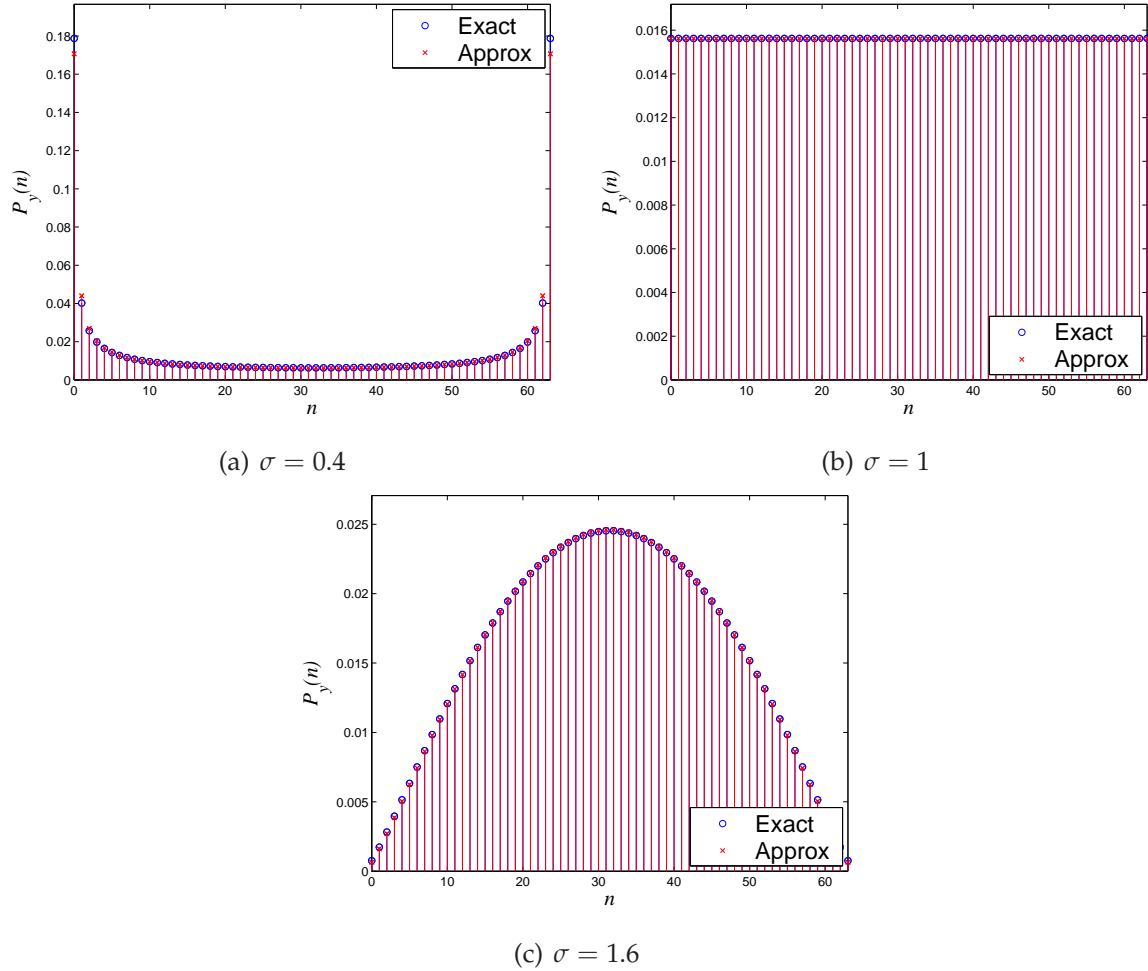


Figure 5.8. Large N approximation to $P_y(n)$. This figure shows for various values of noise intensity, σ , that for $N = 63$ and Gaussian signal and noise, the approximation to the output probability mass function, $P_y(n)$, given by Eqns. (5.56) and (5.56), gives a highly accurate approximation to the output probability mass function. The circles indicate the exact $P_y(n)$ obtained by numerical integration and the crosses show the approximations. Similar results can be obtained for the Laplacian or logistic cases. Fig. 5.8(a) shows that $P_y(0)$ and $P_y(N)$ are the most inaccurate approximations for $\sigma = 0.4$. This is not surprising, given the approximation is expected to decrease in accuracy for small σ , and n close to 0 or N .

5.4 Mutual Information for Large N and Arbitrary σ

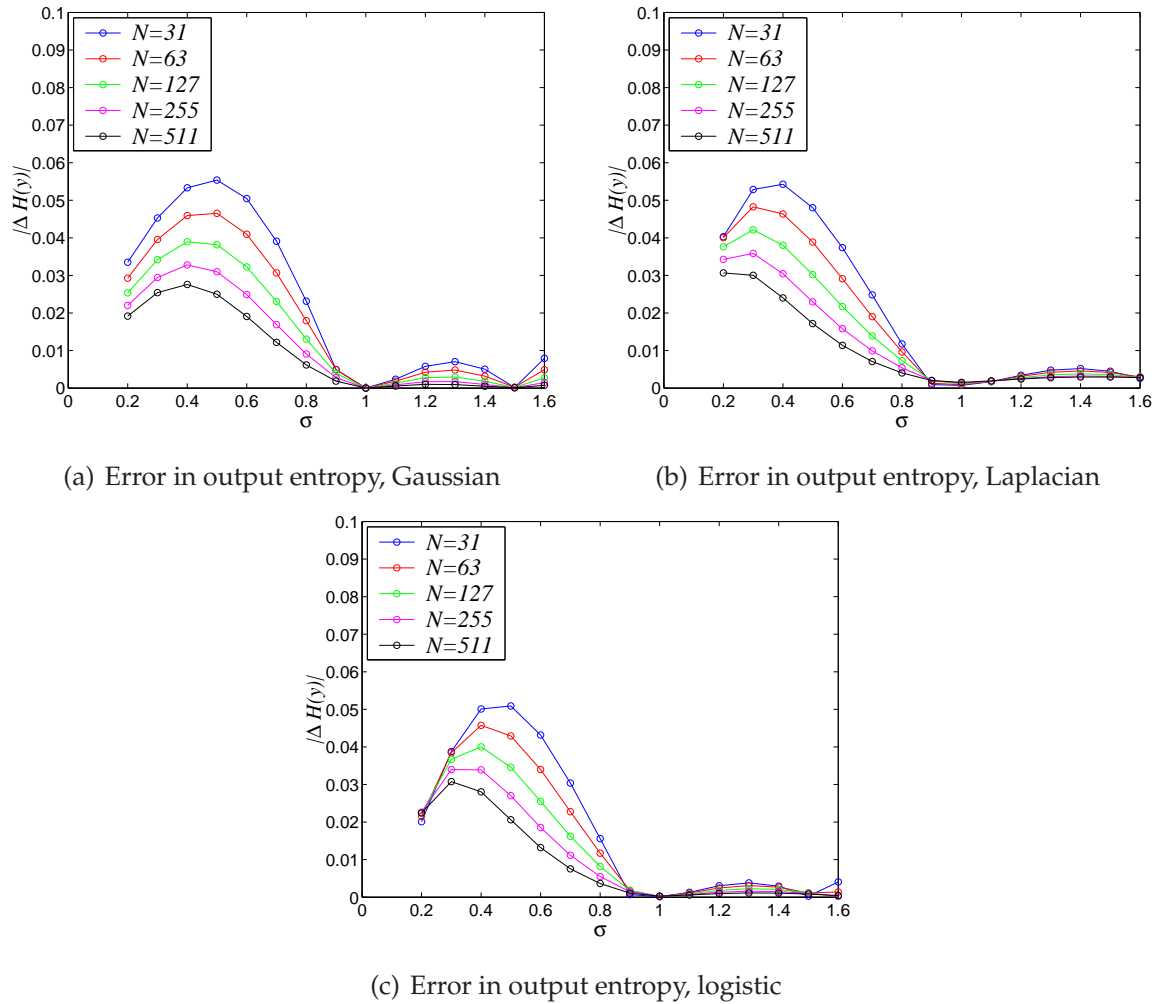


Figure 5.9. Absolute error in output entropy using large N approximation to $P_y(n)$. This figure shows that the approximation to the output probability mass function, $P_y(n)$, given by Eqns. (5.56) and (5.58), gives a highly accurate approximation to the output entropy, as expressed by Eqn. (5.59). The accuracy increases as the array size, N , increases. Notice that the approximation is more accurate near a noise intensity of $\sigma = 1$.

Approximating the output entropy

Consider the entropy of the discrete random variable y . Making use of Eqn. (5.55), we have

$$\begin{aligned}
 H(y) &= - \sum_{n=0}^N P_y(n) \log_2(P_y(n)) \\
 &= - \sum_{n=0}^N \frac{Q\left(\frac{n}{N}\right)}{N} \log_2\left(\frac{Q\left(\frac{n}{N}\right)}{N}\right) \\
 &= - \frac{1}{N} \sum_{n=0}^N Q\left(\frac{n}{N}\right) \log_2\left(Q\left(\frac{n}{N}\right)\right) + \frac{\log_2(N)}{N} \sum_{n=0}^N Q\left(\frac{n}{N}\right). \quad (5.62)
 \end{aligned}$$

Suppose that the summations above can be approximated by integrals, without any remainder terms. Carrying this out and then making the change of variable $\tau = n/N$ gives

$$\begin{aligned}
 H(y) &\simeq - \frac{1}{N} \int_{n=0}^N Q\left(\frac{n}{N}\right) \log_2\left(Q\left(\frac{n}{N}\right)\right) dn + \frac{\log_2(N)}{N} \int_{n=0}^N Q\left(\frac{n}{N}\right) dn \\
 &= - \int_{\tau=0}^{\tau=1} Q(\tau) \log_2(Q(\tau)) d\tau + \log_2 N \int_{\tau=0}^{\tau=1} Q(\tau) d\tau \\
 &= \log_2 N - \int_{\tau=0}^{\tau=1} Q(\tau) \log_2(Q(\tau)) d\tau \\
 &= \log_2 N + H_d(\bar{y}), \quad (5.63)
 \end{aligned}$$

where $H_d(\bar{y})$ is the differential entropy of the random variable \bar{y} . This analysis agrees with Theorem 9.3.1 of Cover and Thomas (1991), which states that the entropy of an M bit quantisation of a continuous random variable Z is approximately the sum of M and the entropy of Z . Here, we have \bar{y} as the continuous random variable that approximates the $N + 1$ state discrete output distribution. Hence the discrete output distribution is a $\log_2(N + 1)$ bit quantisation of \bar{y} , and has entropy approximately equal to the differential entropy of \bar{y} plus $\log_2(N + 1)$, which for large N agrees with Eqn. (5.63).

Performing a change of variable in Eqn. (5.63) of $\tau = F_R(x)$ gives

$$H(y) \simeq \log_2 N - \int_{x=-\infty}^{x=\infty} P(x) \log_2\left(\frac{P(x)}{R(x)}\right) dx = \log_2(N) - D(P(x)||R(x)), \quad (5.64)$$

where $D(P(x)||R(x))$ is the relative entropy between $P(x)$ and $R(x)$.

5.4 Mutual Information for Large N and Arbitrary σ

Thus, $H(y)$ for large N is approximately the sum of the number of output bits and the negative of the relative entropy between $P(x)$ and $R(x)$. Therefore, since relative entropy is always non-negative, the approximation to $H(y)$ given by Eqn. (5.64) is always less than or equal to $\log_2(N)$. This agrees with the known expression for $H(y)$ in the specific case of $\sigma = 1$ of $\log_2(N + 1)$, which holds for any N .

Fig. 5.10 shows the approximation of Eqn. (5.64), as well as the exact output entropy, for the Gaussian, Laplacian and logistic cases, for a range of σ and increasing N . The approximation is quite good for $\sigma > 0.7$, but worsens for smaller σ . However, as N increases the approximation improves. This indicates that approximating the summation in Eqn. (5.62) with the relative entropy between $P(x)$ and $R(x)$ gets more accurate for small σ with increasing N .

We are now in a position to combine the results of this section and the previous, to obtain a result for the mutual information.

5.4.4 Mutual Information for Large N

Section 5.4.2 obtained a large N approximation to the average conditional output entropy, given by Eqn. (5.38), and Section 5.4.3 obtained a large N approximation to the output entropy, given by Eqn. (5.63). These two equations can be combined to give a large N approximation to the mutual information as

$$\begin{aligned} I(x, y) &\simeq H(y) - H(y|x) \\ &= \log_2 N - \int_{\tau=0}^{\tau=1} Q(\tau) \log_2(Q(\tau)) d\tau - 0.5 \log_2(2\pi e N) - \int_{\tau=0}^{\tau=1} Q(\tau) \log_2 \tau d\tau \\ &= 0.5 \log_2 \left(\frac{N}{2\pi e} \right) - \int_{\tau=0}^{\tau=1} Q(\tau) \log_2(\tau Q(\tau)) d\tau \end{aligned} \quad (5.65)$$

$$= 0.5 \log_2 \left(\frac{N}{2\pi e} \right) - \int_{\tau=0}^{\tau=1} Q(\tau) \log_2(\tau) d\tau + H_d(\tau). \quad (5.66)$$

The integral on the RHS of Eqn. (5.65) is independent of N and therefore for large N , the mutual information scales with $0.5 \log_2(N)$. As noted previously in this chapter, and in Stocks (2001a), this is half the maximum channel capacity for an N comparator system. The integral on the RHS of Eqn. (5.65) is insignificant when compared to $\log_2(N)$, but its importance is that it determines how the mutual information varies from $0.5 \log_2 \left(\frac{N}{2\pi e} \right)$ as σ varies.

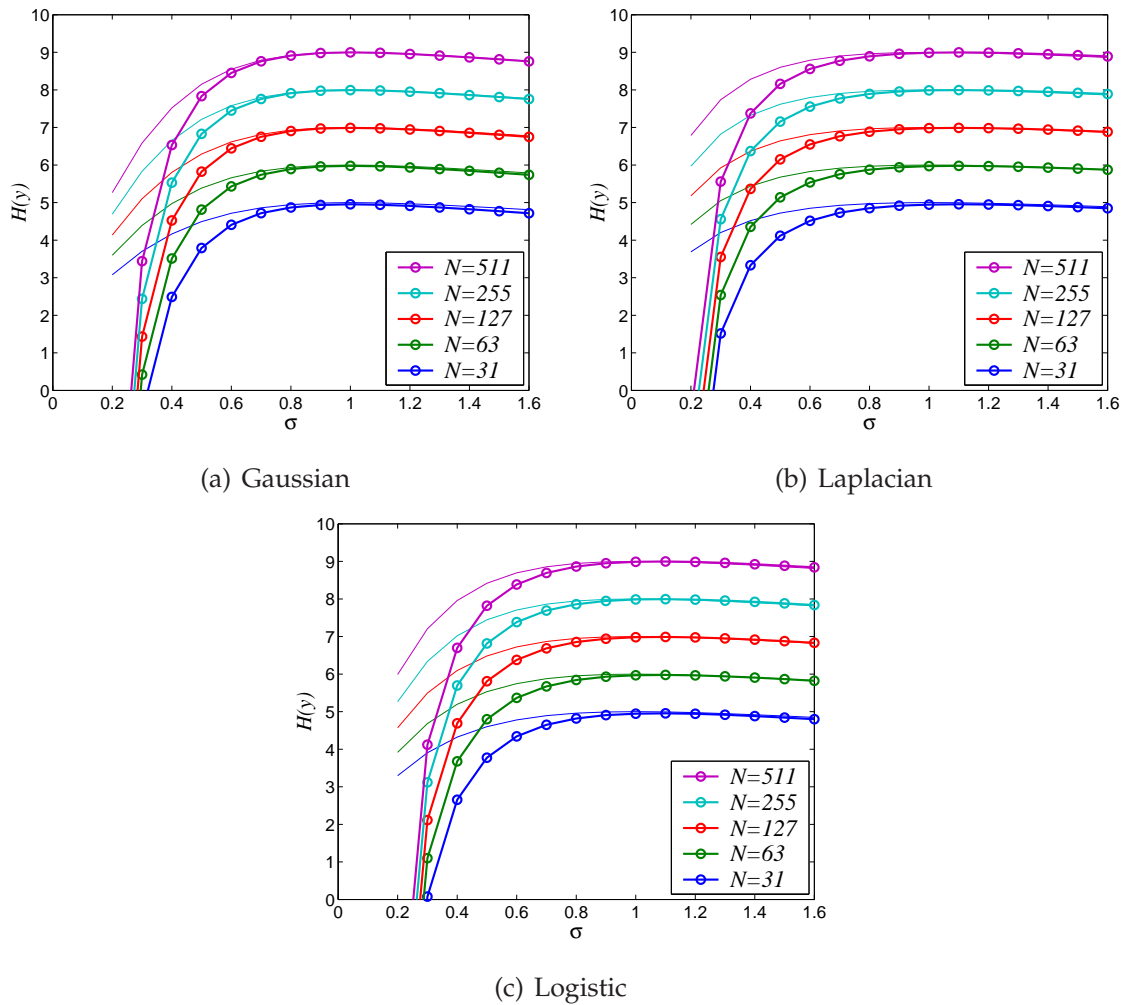


Figure 5.10. Large N approximation to output entropy. This figure shows the approximate output entropy given by Eqn. (5.64), as well as the exact output entropy calculated numerically. Clearly the approximation is quite good for noise intensities, $\sigma > 0.7$. The exact expression is shown by thin solid lines, and the approximation by circles, with a thicker solid line interpolating between values of σ as an aid to the eye. The approximation can be seen to always be a lower bound on the exact entropy.

5.4 Mutual Information for Large N and Arbitrary σ

For the specific case of $\sigma = 1$, $Q(\tau) = 1$ so that $-\int_{\tau=0}^{\tau=1} \log_2(\tau) d\tau = \log_2(e)$. Therefore

$$I(x, y) = 0.5 \log_2 \left(\frac{Ne}{2\pi} \right), \quad (5.67)$$

which for large N agrees precisely with Eqn. (5.19), validating the new formula in this specific case.

Note that Eqn. (5.65) can be rewritten in terms of x as

$$\begin{aligned} I(x, y) &= 0.5 \log_2 \left(\frac{N}{2\pi e} \right) - \int_{x=-\infty}^{x=\infty} P(x) \log_2 \left(\frac{P_{1|x}P(x)}{R(x)} \right) dx \\ &= 0.5 \log_2 \left(\frac{N}{2\pi e} \right) - \int_{x=-\infty}^{x=\infty} P(x) \log_2(P_{1|x}) dx - D(P(x) || R(x)). \end{aligned} \quad (5.68)$$

Fig. 5.11 shows the approximation of Eqn. (5.65), as well as the exact mutual information, for the Gaussian, Laplacian and logistic cases, for a range of σ and increasing N . As with the output entropy, and conditional output entropy approximations, the mutual information approximation is quite good for $\sigma > 0.7$, but worsens for smaller σ . However, as N increases the approximation improves.

5.4.5 Relationship to Fisher Information

Chapters 6 and 7 in this thesis make extensive use of Fisher information (Cover and Thomas 1991) in relation to the decoding of the SSR model. Hoch *et al.* (2003a) shows that the Fisher information for SSR is given by

$$J(x) = \left(\frac{dP_{1|x}}{dx} \right)^2 \frac{N}{P_{1|x}(1 - P_{1|x})} = \frac{NR(x)^2}{P_{1|x}(1 - P_{1|x})}. \quad (5.69)$$

Eqn. (5.69) is also derived in Section D.6 of Appendix D. Note that the Fisher information is always non-negative and has the same support as $R(x)$.

Rewriting the large N expression for mutual information of Eqn. (5.68) gives

$$\begin{aligned} I(x, y) &= -0.5 \log_2(2\pi e) - \int_{x=-\infty}^{x=\infty} P(x) \log_2 \left(\frac{P_{1|x}P(x)}{\sqrt{NR(x)}} \right) dx \\ &= H(x) - 0.5 \log_2(2\pi e) + 0.5 \int_{x=-\infty}^{x=\infty} P(x) \log_2 \left(\frac{NR(x)^2}{P_{1|x}(1 - P_{1|x})} \right) dx \\ &= H(x) - 0.5 \log_2(2\pi e) + 0.5 \int_{x=-\infty}^{x=\infty} P(x) \log_2(J(x)) dx \\ &= H(x) - 0.5 \int_{x=-\infty}^{x=\infty} P(x) \log_2 \left(\frac{2\pi e}{J(x)} \right) dx. \end{aligned} \quad (5.70)$$

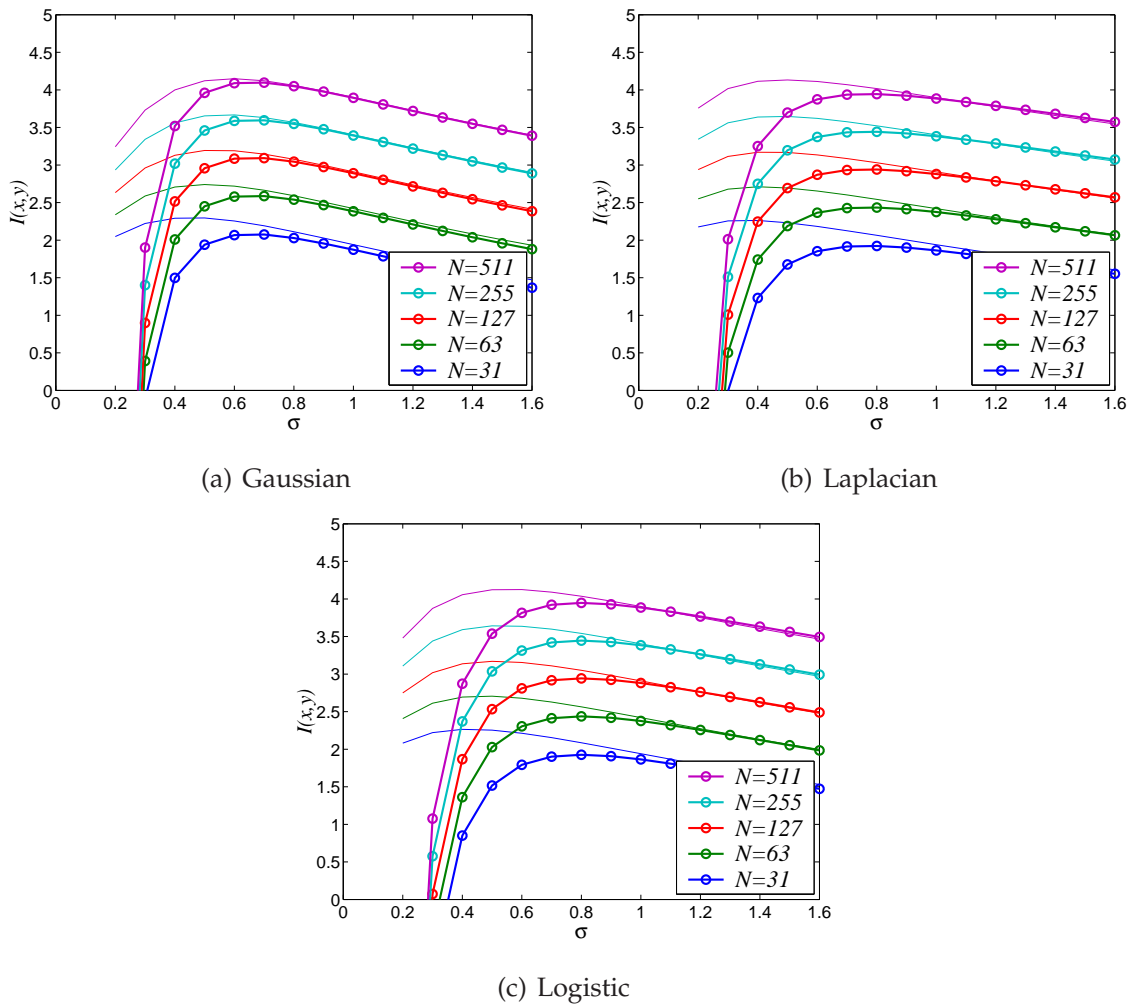


Figure 5.11. Large N approximation to mutual information. This figure shows the approximate output entropy given by Eqn. (5.65), as well as the exact mutual information calculated numerically. Clearly the approximation is quite good for noise intensities, $\sigma > 0.7$. The exact expression is shown by thin solid lines, and the approximation by circles, with a thicker solid line interpolating between values of σ as an aid to the eye. The approximation can be seen to always be a lower bound on the exact mutual information.

5.5 Channel Capacity for Large N

Eqn. (5.70) is precisely the same as that derived in Hoch *et al.* (2003a) as an asymptotic large N expression for the mutual information. Hoch *et al.* (2003a) derived this expression by considering a result from Brunel and Nadal (1998), which shows that in the limit of large N , the mutual information in a system becomes equal to the mutual information between the input signal and an *efficient* Gaussian estimator for that signal. *Efficient* has a precise technical meaning, which we will explore in Chapter 6. Furthermore, discussion of this result will also be examined later, since understanding it depends on the concept of optimally decoding the SSR system. Decoding is the focus of Chapter 6 for small N , and Chapter 7 for large N .

5.5 Channel Capacity for Large N

Inspection of Eqn. (5.65) shows that the large N approximation to the mutual information consists of a term that depends on N and a term that depends only on σ . Hence, for large N , this shows that the channel capacity occurs for the same value of σ —which, as before, we denote as σ_o —for all N . This fact is recognised in Hoch *et al.* (2003a) and Hoch *et al.* (2003b), which derives an analytical approximation for σ_o for the case of Gaussian signal and noise. The derivation is given in Hoch *et al.* (2003b), and depends on a Taylor expansion of the Fisher information inside the integral in Eqn. (5.70). However, before repeating and discussing this derivation, the following subsection shows that a channel capacity achieving input PDF can be found for any given noise PDF, whenever the approximation of Eqn. (5.65) holds.

5.5.1 General Channel Capacity

To begin, Eqn. (5.70) can be written as

$$I(x, y) = -0.5 \log_2(2\pi e) - \int_x P(x) \log_2 \left(\frac{P(x)}{\sqrt{J(x)}} \right) dx. \quad (5.71)$$

The integral in Eqn. (5.71) can be recognised as being very similar to the relative entropy between $P(x)$ and $\sqrt{J(x)}$. However, in general, the Fisher information as a function of x is not a PDF and thus the integral cannot be relative entropy. In the case of SSR though, remarkably—as shown in Section C.3 of Appendix C—the function,

$$S(x) = \frac{\sqrt{J(x)}}{\pi\sqrt{N}} = \frac{R(x)}{\pi\sqrt{P_{1|x}(1 - P_{1|x})}}, \quad (5.72)$$

is indeed a PDF, one which has the same support as $J(x)$, and for the noise distributions considered in this chapter, is a function of the single parameter, σ_η .

Note that the PDF, $S(x)$, can be written in terms of the Fisher information as

$$S(x) = \frac{\sqrt{J(x)}}{\int_x \sqrt{J(\phi)} d\phi}. \quad (5.73)$$

Such a PDF is therefore simply a normalisation of the square root of the Fisher information. A PDF with this property is known as Jeffrey's Prior (Rissanen 1996, Jaynes 2003).

Using Eqn. (5.72), Eqn. (5.71) can be written as

$$I(x, y) = 0.5 \log_2 \left(\frac{N\pi}{2e} \right) - D(P(x) || S(x)). \quad (5.74)$$

Thus, since relative entropy is always non-negative, the large N channel capacity is achieved for any given $R(x)$ when $P(x) = S(x)$ and is given by

$$C(x, y) = 0.5 \log_2 \left(\frac{N\pi}{2e} \right) \simeq 0.5 \log_2 N - 0.3956. \quad (5.75)$$

Eqn. (5.75) holds provided the conditions for the approximation given by Eqn. (5.70) hold, and gives an upper bound for the large N mutual information. In other words, when in the following subsection we consider matched signal and noise distributions, we can expect to find channel capacity that is less than or equal to that of Eqn. (5.75).

5.5.2 Capacity for Matched Signal and Noise

Unlike the previous subsection, we now consider the case of matched signal and noise distributions. From Eqn. (5.74), channel capacity occurs for the value of σ that minimises the relative entropy between $P(x)$ and $S(x)$. From Eqn. (5.70), it is also clear that this is equivalent to solving the following problem,

$$\lim_{N \rightarrow \infty} \sigma_o = \min_{\sigma} f(\sigma) = \int_{x=-\infty}^{x=\infty} P(x) \ln \left(\frac{1}{J(x)} \right) dx. \quad (5.76)$$

This is exactly the formulation stated in Hoch *et al.* (2003b). Recalling the change of variable defined by $\tau = P_{1|x} = F_R(x)$, Problem (5.76) can be equivalently expressed as

$$\lim_{N \rightarrow \infty} \sigma_o = \min_{\sigma} f(\sigma) = \int_{\tau=0}^{\tau=1} Q(\tau) \ln (\tau Q(\tau)) d\tau, \quad (5.77)$$

or as

$$\lim_{N \rightarrow \infty} \sigma_o = \min_{\sigma} f(\sigma) = D(P(x)||R(x)) + \int_{x=-\infty}^{x=\infty} P(x) \log_2(P_{1|x}) dx. \quad (5.78)$$

Thus, the channel capacity depends on a term that comes from the large N approximation to the output entropy, and a term—the relative entropy between the signal and noise PDFs—that comes from the large N approximation to the average conditional output entropy.

5.5.3 Gaussian Signal and Noise

This section gives a slightly different derivation of the optimal value of σ for large N and Gaussian signal and noise to that given in Hoch *et al.* (2003b). We begin with Problem (5.78). Solving this problem requires differentiating $f(\sigma)$ with respect to σ and solving for zero. This means

$$\frac{d}{d\sigma} D(P(x)||R(x)) + \frac{d}{d\sigma} \int_{x=-\infty}^{x=\infty} P(x) \log_2(P_{1|x}) dx = 0. \quad (5.79)$$

Recall from Table 4.2 in Chapter 4 that the entropy between $P(x)$ and $R(x)$ is

$$D(P(x)||R(x)) = \log_2(\sigma) + \frac{1}{2 \ln 2} \left(\frac{1}{\sigma^2} - 1 \right). \quad (5.80)$$

Therefore

$$\frac{d}{d\sigma} D(P(x)||R(x)) = \frac{1}{\ln 2} \left(\sigma^{-1} - \sigma^{-3} \right). \quad (5.81)$$

For the other term in Eqn. (5.79), we take the lead from Hoch *et al.* (2003b) and approximate $\ln(P_{1|x})$ by its second order Taylor series expansion (Spiegel and Liu 1999) about the mean of the noise. Hoch *et al.* (2003b) uses an arbitrary mean, however here, in line with the rest of this chapter, we set the signal and noise means to be zero. The result of this, after some algebra, and using the fact that for Gaussian noise, $P_{1|x} = 0.5$ at $x = 0$, and $R(0) = 1/\sqrt{2\pi\sigma_{\eta}^2}$, where σ_{η} is the standard deviation of the noise, is

$$-\ln(P_{1|x}) = \ln 2 - \sqrt{\frac{2}{\pi}} \frac{x}{\sigma_{\eta}} + \frac{x^2}{\pi\sigma_{\eta}^2} - \mathcal{O}\left(\frac{x^4}{\sigma_{\eta}^4}\right). \quad (5.82)$$

Multiplying Eqn. (5.82) by $P(x)$ and integrating over all x gives

$$-\int_{x=-\infty}^{x=\infty} P(x) \ln(P_{1|x}) dx \simeq \ln 2 - \sqrt{\frac{2}{\pi}} \frac{1}{\sigma_{\eta}} E[x] + \frac{1}{\pi\sigma_{\eta}^2} E[x^2]. \quad (5.83)$$

Since for a Gaussian signal, $E[x] = 0$ and $E[x^2] = \sigma_x^2$, Eqn. (5.83) becomes a function of σ ,

$$g(\sigma) = - \int_{x=-\infty}^{x=\infty} P(x) \log_2(P_{1|x}) dx \simeq 1 + \frac{1}{\pi\sigma^2 \ln 2}, \quad (5.84)$$

where as before, $\sigma = \sigma_\eta / \sigma_x$.

The approximation of Eqn. (5.84) appears to be quite accurate for all σ , as can be seen in Fig. 5.12. Numerical experiments show that the relative error is no more than about 10 percent for $\sigma > 0.2$. However, as we will see, this is inaccurate enough to cause the end result for the approximate channel capacity to significantly overstate the true channel capacity.

The derivative of Eqn. (5.84) with respect to σ is

$$-\frac{d}{d\sigma} \int_{x=-\infty}^{x=\infty} P(x) \log_2(P_{1|x}) dx \simeq -\frac{2}{\pi\sigma^3 \ln 2}. \quad (5.85)$$

Combining Eqn. (5.85) with Eqn. (5.81) and substituting into Eqn. (5.79) gives

$$\frac{2}{\pi\sigma^3 \ln 2} + \frac{1}{\ln 2} (\sigma^{-1} - \sigma^{-3}) = 0. \quad (5.86)$$

Solving Eqn. (5.86) gives the optimal value of σ as

$$\sigma_o \simeq \sqrt{1 - \frac{2}{\pi}} \simeq 0.6028, \quad (5.87)$$

as found in Hoch *et al.* (2003b).

An expression for the mutual information at σ_o can be found by substituting Eqn. (5.87) into Eqns. (5.84) and (5.80), and substituting the results into Eqn. (5.68). Carrying this out gives the large N channel capacity for Gaussian signal and noise as

$$C(x, y) \simeq 0.5 \log_2 \left(\frac{2N}{e(\pi - 2)} \right), \quad (5.88)$$

which can be written as $C(x, y) \simeq 0.5 \log_2 N - 0.3169$.

Although this expression for the channel capacity is close to correct, recall that the previous subsection showed that the channel capacity must be less than $0.5 \log_2 N - 0.3956$. Hence, Eqn. (5.88) gives a value for channel capacity that is too large. The reason for this is that the Taylor expansion approximation we used requires consideration of more terms. An expression including such higher order terms is given in Hoch *et al.* (2003b), however this also appears to not give quite the right answer. Hence, the next section solves Problem (5.78) numerically.




Figure 5.12. Approximation to $g(\sigma) = \int_{x=-\infty}^{x=\infty} P(x) \log_2(P_{1|x}) dx$. This figure shows the approximation given for Gaussian signal and noise by Eqn. (5.84), compared with the corresponding exact result, calculated numerically. Clearly, the approximation is quite accurate for most values of the noise intensity, σ , however, as shown in the text, it is inaccurate enough to cause the approximation to significantly overstate the true channel capacity.

5.5.4 Numerical Verification and Other Distributions

Due to the slight inaccuracy found in the previous subsection, and the fact that no analytical expression could be found for any other signal and noise distributions, we now solve Problem (5.78) numerically. The term $g(\sigma) = \int_x P(x) \log_2(P_{1|x}) dx$ can be found for any specified signal and noise distribution by numerical integration, just as carried out for the relative entropy plotted in Fig. 4.5 in Chapter 4 for cases other than Gaussian. Thus, the function of Problem (5.78) can be easily found as a function of x , and minimised. This function is plotted in Fig. 5.13.

Fig. 5.13 also shows the Gaussian case obtained using the approximation of Eqn. (5.82). It is clear that although this approximation is close to being correct, it does understate the true value of σ_o , and gives a more negative value of the function, which means overstating the channel capacity.

Table 5.1 gives the result of numerically calculating the value of σ_o , and the corresponding large N channel capacity. Note that in each case, $C(x, y) - 0.5 \log_2(N) < -0.3956$, as required by Eqn. (5.75). Note that the difference between capacity and $0.5 \log_2(N)$ is about 0.4 bits per sample in each case. Thus, in the limit of very large N , this shows that capacity is almost identical, regardless of the distribution. The value of σ_o at which

Table 5.1. Large N channel capacity and optimal σ for matched signal and noise.

Distribution	$C(x, y) - 0.5 \log_2(N)$	σ_o	$C(x, y) - I(x, y)_{\sigma=1}$
Gaussian	-0.3964	0.6563	0.208
Logistic	-0.3996	0.5943	0.205
Laplacian	-0.3990	0.5384	0.205

this capacity occurs though is different in each case. The results of Table 5.1 compare favourably with the results presented in Section 4.4 in Chapter 4 where capacity was found for N up to 1000. At $N = 1000$, for Gaussian signal and noise, $\sigma_o \simeq 0.607$, for Laplacian signal and noise, $\sigma_o \simeq 0.5$ and for logistic signal and noise $\sigma_o \simeq 0.564$. Since we also saw that σ_o increases with increasing N , these results are consistent with Table 5.1.

Difference Between Channel Capacity and $I(x, y)$ at $\sigma = 1$

Section 4.4 in Chapter 4 shows in Fig. 4.24(c) the difference between the channel capacity, and the mutual information at $\sigma = 1$. This figure shows that for $N = 1000$, this difference appears to be asymptotically converging towards a constant value for each matched signal and noise case. The value of this difference appears to be of the order of 0.22–0.24 bits per sample. Recall from Section 5.2 that at $\sigma = 1$, the mutual information is identical for each signal and noise pair considered here, and is approximately $I(x, y) = 0.5 \log_2(N) - 0.6444$. Thus, given that the channel capacity is slightly larger than this, as indicated by Table 5.1, for each case there is a constant difference between the channel capacity and the mutual information at $\sigma = 1$, for large N . This value is also listed in Table 5.1, and compares favourably with the results of Fig. 4.24(c) at $N = 1000$, which is simply not large enough to show the final asymptotic value of this difference.

The main conclusion to be drawn from this analysis though, is that the channel capacity is only of the order of 0.2 bits per sample larger at σ_o than it is at $\sigma = 1$. As N gets larger, this difference becomes more and more insignificant.

5.6 Chapter Summary

The initial section of this chapter gives a brief discussion of previous work on large N limit results for the SSR model.




Figure 5.13. Finding channel capacity for large N . This figure shows the function that requires minimisation to find the channel capacity, that is, $f(\sigma) = D(P(x)||R(x)) + \int_{x=-\infty}^{x=\infty} P(x) \log_2(P_{1|x}) dx$, as a function of noise intensity, σ , for the Gaussian, logistic and Laplacian cases. The Gaussian case obtained using the approximation of Eqn. (5.82) is also shown. Clearly, although the approximation gives σ_o close to the right answer, it underestimates the true value slightly, and also provides a minimum smaller than the true minimum, and hence significantly overstates the channel capacity.

Next, Section 5.2 re-derives and improves a large N expression for the mutual information first found in Stocks (2001a). This expression holds for the case of matched signal and noise, that is, for the case where $P(x) = R(\theta - x)$, so that the mutual information is independent of the actual signal and noise distribution.

Section 5.3 re-derives, and improves, a large N expression first found in Stocks (2001c) for the mutual information for the specific case of uniform signal and noise and $\sigma \leq 1$. It also gives an expression for the value of σ at which channel capacity occurs, in this case. The mutual information is found to scale with $0.5 \log_2 N$.

Section 5.4 then looks at the more general case of any σ , and any matched signal and noise distributions. Work on this topic was first published in Hoch *et al.* (2003a), who found an expression for the large N mutual information that is more accurate for large N and σ greater than about 0.7. We obtain an alternative derivation of this expression, by finding separate expressions for the output entropy, conditional output entropy, and output probability mass function. Furthermore, it is shown that all expressions can be written either in terms of the support of the signal distribution, x , and its associated PDF, $P(x)$, or in terms of the transformed random variable, τ , and its associated PDF, $Q(\tau)$, with support $\tau \in [0, 1]$.

Finally Section 5.5 uses the large N expression for mutual information to find an expression for a channel capacity achieving input PDF, for any given noise PDF. This expression is related to the Fisher information, and holds for distributions for which the large N mutual information approximation holds. This section also numerically finds the large N channel capacity for the cases of matched Gaussian signal and noise, Laplacian signal and noise, and logistic signal and noise. It is found in each case that the large N channel capacity is only about 0.2 bits per sample greater than the channel capacity at $\sigma = 1$, although the optimal noise intensity, σ , ranges between about 0.5 and 0.65. Furthermore, for the Gaussian signal and noise case, the optimal value of σ reported in Hoch *et al.* (2003a) is found to underestimate the true value by about 0.05.

5.6.1 Original Contributions for Chapter 5

This chapter included the following original contributions:

- An improved approximation to the large N average conditional output entropy, and mutual information at $\sigma = 1$, is obtained by making use of Stirling's formula. The improved mutual information approximation is shown to have an error that approaches zero as N increases, whereas the formula given in Stocks (2001a) is shown to have an error approaching about 0.117 bits per sample.
- An improved approximation to the large N average conditional output entropy, and mutual information for uniform signal and noise, and $\sigma \leq 1$ is derived. The improvement is due to the same reasoning as the improvement at $\sigma = 1$. The improved formula for the mutual information also leads to an improved formula for the optimal value of σ , and the channel capacity. As in Stocks (2001c), these new formulas show that the channel capacity occurs for $\sigma = 1$ at $N = \infty$.
- A new approximation to the conditional entropy of y given x , $\hat{H}(y|x)$, is given, based on the fact that the binomial distribution approaches a Gaussian for large N . This approximation is shown to give the same result as the approximation used in Hoch *et al.* (2003b), and is used to approximate the *average* conditional entropy, $H(y|x)$, for large N . The approximation to $\hat{H}(y|x)$ is numerically verified to become more accurate for increasing N , and the approximation to $H(y|x)$ is verified at $\sigma = 1$ by the results of Section 5.2.

- A derivation of the PDF of the average transfer function, \bar{y} , of the SSR model is given. This PDF is shown to be $Q(\frac{\bar{y}}{N})/N$.
- A large N approximation of the output probability mass function, $P_y(n)$, that is highly accurate for all σ is derived. This approximation is obtained in two different ways, and numerically verified to be accurate.
- A large N approximation to the output entropy, $H(y)$, is found, and is shown to be the difference between $\log_2(N)$ and the relative entropy between $P(x)$ and $R(x)$, or equivalently, the sum of $\log_2(N)$ and the entropy of τ .
- The new methods of approximating $H(y|x)$ and $H(y)$ are used to obtain the same large N approximation to $I(x, y)$ as expressed in Hoch *et al.* (2003b).
- An expression for a channel capacity achieving input PDF for any given noise PDF is found. The channel capacity for these conditions is derived as $C(x, y) = 0.5 \log_2(\frac{N\pi}{2e})$, thus giving an upper bound for the achievable channel capacity for SSR, under the conditions for which the large N mutual information approximation formula holds.
- Asymptotic large N expressions for the channel capacity, and optimal noise intensity, σ , for the Gaussian, logistic and Laplacian cases are found. It is shown that the difference between capacity and the mutual information at $\sigma = 1$ is of the order of 0.2 bits per sample for large N . It is also shown that the optimal σ for Gaussian signal and noise given in Hoch *et al.* (2003a) understates the true value by about 0.05.

5.6.2 Further Work

Possible future work and open questions arising from this Chapter might include:

- Analytical proofs that the large N approximations converge to zero error as N approaches infinity.
- Analytical proof that the large N approximation to the output probability mass function converges to the exact output probability mass function.
- Studies of large N approximations for one-sided signal and noise PDFs such as the Rayleigh or exponential PDFs, or finite support PDFs, such as the uniform case, when $\sigma > 1$.

- Studies of large N approximations for mixed signal and noise distributions.
- More rigorous justification of the use of the Fisher information approximation to the mutual information, and its relationship to the concept of stochastic complexity (Rissanen 1996) and Minimum Description Length (MDL) (Grunwald *et al.* 2005). Such work could begin by building on the work of Davisson and Leon-Garcia (1980), Clarke and Barron (1990), Barron and Cover (1991), Nadal and Parga (1994), Rissanen (1996) and Brunel and Nadal (1998), and may find an alternative derivation of the SSR channel capacity and large N mutual information in the MDL framework.

This concludes Chapter 5. Chapters 4 and 5 studied information transmission and encoding in the SSR model, by analytical and numerical studies of the mutual information and channel capacity. The following Chapter studies the SSR model as a quantisation or lossy source coding model, by examining decoding of the SSR model output, to achieve low distortion between the input and output signals.

Suprathreshold Stochastic Resonance: Decoding

THE initial research into suprathreshold stochastic resonance considers the viewpoint of information transmission. As discussed briefly in Chapter 4, the suprathreshold stochastic resonance effect can also be modelled as stochastic quantisation, and therefore provides a nondeterministic lossy compression of a signal. The reason for this is that the effect of independently adding noise to a common signal before thresholding the result a number of times, with the same static threshold value, is equivalent to quantising a signal with random thresholds. This observation leads naturally to measuring and describing the performance of suprathreshold stochastic resonance with standard quantisation theory. In a context where a signal is to be reconstructed from its quantised version, this requires a *reproduction value* or *reproduction point* to be assigned to each possible state of a quantised signal. The quantising operation is often known as the *encoding* of a signal, and the assignment of reproduction values as the *decoding* of a signal. This chapter examines various methods for decoding the output of the suprathreshold stochastic resonance model, and evaluates the performance of each technique, as the input noise intensity and array size changes. As it is the performance criterion most often used in conventional quantisation theory, the measure used is the *mean square error distortion* between the original input signal, and the decoded output signal.

6.1 Introduction

We begin this Chapter by very briefly reviewing the SSR model introduced in Chapter 4. We then introduce the concept of *decoding* the output of a quantiser's encoding to reconstruct the input signal, and consider measuring the performance of such a reconstruction. Three measures are considered, the *mean square error* distortion, the *signal-to-quantisation-noise ratio* (SQNR) and the *correlation coefficient*. The latter two measures are however shown to be very closely related to the Mean Square Error (MSE) distortion. The remainder of this Chapter is outlined in Section 6.1.3.

This Chapter contains original work on the description of the SSR model as a quantiser that has been published, in part, in the open literature (McDonnell *et al.* 2002a, McDonnell *et al.* 2002b, McDonnell *et al.* 2003a, McDonnell *et al.* 2005c, McDonnell *et al.* 2005d).

6.1.1 Summary of Relevant Results from Chapters 4 and 5

Chapter 4 introduced the array of threshold devices in which Suprathreshold Stochastic Resonance (SSR) occurs. The model is shown in Fig. 4.1, and consists of N binary threshold devices. As in Chapters 4 and 5, we consider each threshold device to have an identical value, θ , and to receive the same input signal, x , which is a sequence of independent samples from the random signal with Probability Density Function (PDF) given by $P(x)$ and a standard deviation being a function of σ_x . However, *iid* additive noise, with PDF $R(\eta)$, and a standard deviation being a function of σ_η , is present at the input to each threshold, so that the output of each threshold device is unity if $x + \eta > \theta$ or zero otherwise. The overall output of the SSR model is the sum of all the threshold outputs, $y \in [0, N]$, which is an integer encoding of the input signal.

We saw in Chapter 4 that the performance of the SSR model can be measured by the mutual information between x and y , and that the maximum mutual information occurs for a nonzero noise intensity. Calculating the mutual information depends on knowing the *transition probabilities*, $P(y = n|x)$, which we abbreviate to $P(n|x)$. For the SSR model, this conditional probability mass function is given by the binomial distribution, in terms of the probability that any given threshold device is on, given x , denoted as $P_{1|x}$. Furthermore, the output probability mass function can be found by integrating the joint input-output probability density function, $P(x, y) = P(x)P(n|x)$ to get $P_y(n) = \int_x P(x)P(n|x)dx$.

A convenient parameterisation of the noise intensity is given by the ratio of the noise standard deviation to the signal standard deviation, which for ‘matched’ signal and noise distributions is $\sigma = \sigma_\eta / \sigma_x$. For such a case, the mutual information can be expressed as a function of σ , and is therefore independent of the signal’s variance for a given σ . In this chapter however, we will see that sometimes the performance of the decoded SSR array depends on both σ_η and σ_x .

Furthermore, under certain conditions on the support of $P(x)$ and $R(x)$, the mutual information can be expressed, by a transformation using the Inverse Cumulative Distribution Function (ICDF) of the noise, $F_R^{-1}(\cdot)$, in terms of a PDF with support between zero and unity,

$$Q(\tau) = \frac{P(x)}{R(\theta - x)} \Big|_{x=\theta - F_R^{-1}(1-\tau)}. \quad (6.1)$$

Chapter 5 shows that $Q(\tau)$ is the PDF of the *average* transfer function of the SSR model divided by N —that is, the PDF of the random variable, $E[y|x]/N$ —and how, for large N , the mutual information is highly dependent on the entropy of $Q(\tau)$.

In this chapter, we will consider only cases where $P(x)$ and $R(x)$ are even functions about a mean of zero and the threshold value is $\theta = 0$.

The next subsection introduces the concept of decoding a quantised signal to reconstruct an approximation to the input signal, and how the performance of such a reconstruction can be measured by the MSE distortion.

6.1.2 Decoding a Quantised Signal

The output of the SSR model, y , is a non-deterministic integer encoding of the input signal, x . The set of possible values of y are the integers from 0 to N , and is therefore a quantisation of x . We consider now any scalar quantisation scheme. To enable a reconstruction of the input signal from such a quantisation, each value of y must be assigned a *reconstruction point*. Label the n -th reconstruction point, that is, the point corresponding to $y = n$, as \hat{y}_n . Therefore, the decoded output is the discretely valued signal \hat{y} , which has possible values $\hat{y}_0, \dots, \hat{y}_n, \dots, \hat{y}_N$.

Assuming suitable values of $\hat{y} = \{\hat{y}_n\}$ have been set, then it is possible to define an error signal between the input x , and the reconstructed input, \hat{y} . Let this error be

$$\epsilon = x - \hat{y}. \quad (6.2)$$

Clearly, this error is a function of x . Since x is a continuously valued variable, some measure of the *average* error—or *distortion*—is required.

Mean square error distortion

There are many possible ways to define such an average error, the most common of which are the absolute error and the Mean Square Error (MSE) distortion. The MSE distortion is also sometimes known as the quantisation noise. We will only focus on the MSE distortion, as this is very commonly used in estimation and quantisation theory (Gray and Neuhoff 1998).

Firstly, it is sometimes convenient to consider the conditional MSE distortion for a given value of the input, x , which we label as

$$\begin{aligned}
 D(x) &= \text{E}[\epsilon^2|x] \\
 &= \text{E}[(x - \hat{y})^2|x] \\
 &= \sum_{n=0}^N P(n|x)(x - \hat{y}_n)^2.
 \end{aligned} \tag{6.3}$$

Note also that

$$\begin{aligned}
 D(x) &= \text{var}[\epsilon|x] + \text{E}[\epsilon|x]^2 \\
 &= \text{var}[\epsilon|x] + \text{E}[\hat{y} - x|x]^2 \\
 &= \text{var}[\epsilon|x] + (\text{E}[\hat{y}|x] - x)^2 \\
 &= \text{var}[\epsilon|x] + b_{\hat{y}}(x)^2,
 \end{aligned} \tag{6.4}$$

where $b_{\hat{y}}(x)$ is the *bias* of the decoded output, \hat{y} , as a function of x . The concept of bias is defined and discussed in Section 6.7, but for now we simply comment that the above equation shows that the conditional MSE distortion is the sum of the conditional variance of the error and the square of the bias. This illustrates that the distortion is strongly dependent on both the variance of the error and the bias.

The MSE distortion is then defined as the average value of $D(x)$ over all x , that is,

$$\begin{aligned}
 \text{MSE} &= \text{E}[\epsilon^2] \\
 &= \text{E}[D(x)] \\
 &= \sum_{n=0}^N \int_{x=-\infty}^{\infty} P(x, y)(x - \hat{y}_n)^2 dx \\
 &= \int_{x=-\infty}^{x=\infty} P(x) \left(\sum_{n=0}^N P(n|x)(x - \hat{y}_n)^2 \right) dx.
 \end{aligned} \tag{6.5}$$

Thus, using this definition, provided the decoding and the transition probabilities are known, the MSE distortion can be calculated numerically.

Also, the MSE can be expressed as

$$\begin{aligned} \text{MSE} &= \int_x P(x)D(x)dx \\ &= \int_x P(x) \left(\text{var}[\epsilon|x] + E[\epsilon|x]^2 \right) dx \\ &= E[\text{var}[\epsilon|x]] + E[b_{\hat{y}}(x)^2]. \end{aligned} \quad (6.6)$$

We will call the first term in Eqn. (6.6) the *average error variance* and the second term the *mean square bias*.

The expression given by Eqn. (6.5) can also be simplified to give

$$\begin{aligned} \text{MSE} &= E[\epsilon^2] \\ &= E[\hat{y}^2] - 2E[x\hat{y}] + E[x^2], \end{aligned} \quad (6.7)$$

where $E[\hat{y}^2]$ is the mean square value of the decoded output, \hat{y} , $E[x^2]$ is the mean square value of the input, x , and $E[x\hat{y}]$ is the correlation between x and \hat{y} . In the SSR model, since the input signal has a mean of zero, the mean square value of x is also its variance. Without exception, the decoded output mean will also always be zero in this Chapter and therefore the mean square value of the decoded output is also always the variance of the decoded output. In contrast, the *encoded* output signal for SSR has a non-zero mean—which is shown in Section D.2.1 of Appendix 6 to be $N/2$ —and therefore has a variance

$$\text{var}[y] = E[y^2] - N^2/4. \quad (6.8)$$

Output signal-to-quantisation-noise ratio

The lossy source coding and quantisation literature often defines an output signal-to-noise ratio measure as the ratio of the input signal power to the MSE distortion—or quantisation noise—power (Proakis and Manolakis 1996, Gray and Neuhoff 1998, Lathi 1998). Since we are only analysing signals that are stationary random variables with zero means, the input signal's power is simply its mean square value. Thus, the output SQNR is

$$\text{SQNR} = \frac{E[x^2]}{\text{MSE}} = \frac{E[x^2]}{E[(x - \hat{y})^2]}. \quad (6.9)$$

Note that it is conventional to express SNR measures in terms of decibels (dB), which for the SQNR is $10 \log_{10}(\text{SQNR})$.

Correlation coefficient

The correlation coefficient between two random variables, x and \hat{y} is defined as

$$\rho_{x\hat{y}} = \frac{\text{cov}[x\hat{y}]}{\sqrt{\text{var}[x]\text{var}[\hat{y}]}} \quad (6.10)$$

where $\text{cov}[x\hat{y}] = E[x\hat{y}] - E[x]E[\hat{y}]$ is the covariance between x and \hat{y} . Since, here, $E[x] = 0$ and $E[\hat{y}] = 0$,

$$\rho_{x\hat{y}} = \frac{E[x\hat{y}]}{\sqrt{E[x^2]E[\hat{y}^2]}} \quad (6.11)$$

Correlation coefficient is a measure of the linearity between two random variables. It is equal to unity if an exact linear relationship can be found (Yates and Goodman 2005). It is zero if the two random variables are independent. It will be shown that simple relationships exist between the correlation coefficient and SQNR for optimal linear and nonlinear decoding, and that therefore the correlation coefficient and SQNR are equivalent measures for such decoding schemes.

6.1.3 Chapter Structure

Having introduced the concept of decoding a quantised signal, and measures of the performance of such decodings, we are now ready to discuss in detail their application to the SSR model. However, firstly in Section 6.2, we will use this theory to discuss noise removal via averaging for analog signals. This is for later comparison with the performance of the SSR effect—which can be seen as an averaging of N single bit digital signals—for noise removal.

Next, Sections 6.3 and 6.4 consider *linear* decoding, while Section 6.5 considers *nonlinear* decoding. It is shown in Section 6.6 that although nonlinear decoding can provide a smaller MSE, linear decoding is simpler to specify and calculate.

Section 6.7 shows how the output of the decoded SSR model can be considered as an estimate of the input, x , and how the performance of its decoding can be understood using estimation theory. The main result is to derive an expression for the *Fisher information* of the SSR model, and to use the Fisher information to derive theoretical lower bounds on the MSE performance of the system. Section 6.7 also demonstrates how the performance of the decoded SSR model depends strongly both on *bias* and *variance*.

Finally, Section 6.8 discusses an approach proposed in Stocks (2001a) to measuring the SSR model's output in terms of an SNR measure that is different to the SQNR.

Some results relevant to this chapter are given in more detail in Appendix D.

6.2 Averaging Without Thresholding

It is well known that averaging N independently noisy versions of a signal reduces the noise by a factor of $1/\sqrt{N}$. We begin by proving this result.

Suppose x is a random signal consisting of a sequence of samples drawn from some continuously valued probability distribution with PDF $P(x)$, zero mean and variance σ_x^2 . Suppose $i = 1, \dots, N$ and η_i is a set of uncorrelated random signals drawn from a distribution with PDF $R(\eta)$, zero mean, and variance σ_η^2 . Let $y_i = x + \eta_i$ be the sum of the signal and the i -th noise signal. Then the error between y_i and x is $\epsilon_i = y_i - x = \eta_i$, and the MSE distortion between the raw signal, x , and the noisy signal is the noise variance, $D(x) = \sigma_\eta^2 \forall x$.

Suppose now that the N independently noisy signals are averaged so that

$$\begin{aligned} y(x) &= \frac{1}{N} \sum_{i=1}^N y_i \\ &= \frac{1}{N} \sum_{i=1}^N (x + \eta_i) \\ &= x + \frac{1}{N} \sum_{i=1}^N \eta_i. \end{aligned} \tag{6.12}$$

The error between $y(x)$ and x has a mean of zero and is given by

$$\epsilon(x) = y(x) - x = \frac{1}{N} \sum_{i=1}^N \eta_i \quad \forall x. \tag{6.13}$$

The mean square error for a given x is then

$$D(x) = E[\epsilon(x)^2] = \frac{1}{N^2} E \left[\left(\sum_{i=1}^N \eta_i \right)^2 \right]. \tag{6.14}$$

6.2 Averaging Without Thresholding

Since each noise signal is uncorrelated with all other noise signals, Eqn. (6.14) simplifies to

$$\begin{aligned} D(x) &= \frac{1}{N^2} \mathbb{E} \left[\sum_{i=1}^N \eta_i^2 \right] \\ &= \frac{1}{N^2} \sum_{i=1}^N \mathbb{E}[\eta_i^2] \\ &= \frac{1}{N^2} N \sigma_\eta^2 \\ &= \frac{\sigma_\eta^2}{N} \quad \forall x. \end{aligned} \tag{6.15}$$

Thus, the conditional MSE distortion is reduced by a factor of N , and is constant for all values of x . Hence, the MSE distortion is also the conditional mean square distortion, and is independent of the signal variance. This fact also means that the signal, x , need not be a random signal at all, but can be a deterministic signal of any form. Note that each y_i is an *iid* random variable, with a mean of x , and variance σ_η^2 .

The square root of the MSE distortion is often referred to as the root-mean-square (rms) error. Thus, the well known result of a $\frac{1}{\sqrt{N}}$ reduction in noise due to averaging refers to using the RMS error as the noise measure.

In terms of SNR, the SNR after averaging is

$$\text{SNR} = \frac{\sigma_x^2}{\text{MSE}} = \frac{N}{\sigma^2}, \tag{6.16}$$

where $\sigma^2 = \sigma_\eta^2 / \sigma_x^2$ is the reciprocal of the SNR of a single noisy input signal, y_i . Thus, the SNR after averaging is N times the SNR of a single noisy version of the signal, y_i .

In Section 6.8 we will compare this result with measures of the output SNR in the SSR model.

One can envisage situations where thresholding first before averaging may be advantageous, or indeed the only possible way of averaging, most particularly, in digital systems. In particular, the advantage of thresholding first is that it provides compression to a discrete number of possible values.

We now proceed to discussing specific decoding methods for quantisation schemes.

6.3 Linear Decoding Theory

Consider linear decoding of a quantiser's encoding stage output, $y \in [0, N]$. Such a decoding operation can be written as

$$\hat{y} = ay + b, \quad y \in 0, \dots, N, \quad (6.17)$$

where a and b are constants for all values of y . Such linear decoding will give a set of outputs that are evenly spaced. We label the n -th value of \hat{y} as \hat{y}_n .

Assume in general that y is a non-deterministic function of an input signal, x and that $E[y] = N/2$. If we impose the condition that the mean of the decoded output, \hat{y} , is also zero, then this requires that $E[\hat{y}] = aE[y] + b = 0$. Therefore setting $a = -2b/N$ is sufficient to provide a zero-meaned decoding. Changing notation by letting $c = -b$ gives

$$\hat{y} = \frac{2c}{N}y - c, \quad y \in 0, \dots, N. \quad (6.18)$$

Thus, the mean square value of \hat{y} is also its variance, since the mean is zero.

Conditional linear MSE distortion

As well as the overall MSE distortion, we will also be interested in the *conditional* moments of the decoding and the conditional MSE distortion for a given value of x ; each value of x has its own average distortion.

Unlike the overall output, the conditional mean of \hat{y} given x is, in general, non-zero, since

$$E[\hat{y}|x] = \frac{2c}{N}E[y|x] - c. \quad (6.19)$$

Consequently, the conditional variance of \hat{y} will not equal the conditional mean square value of \hat{y} for a given x .

The conditional mean square value of the decoding is

$$E[\hat{y}^2|x] = \frac{4c^2}{N^2}E[y^2|x] - \frac{4c^2}{N}E[y|x] + c^2, \quad (6.20)$$

and the conditional variance is

$$\begin{aligned} \text{var}[\hat{y}|x] &= \frac{4c^2}{N^2} \left(E[y^2|x] - E[y|x]^2 \right) \\ &= \frac{4c^2}{N^2} \text{var}[y|x]. \end{aligned} \quad (6.21)$$

6.3 Linear Decoding Theory

Eqns. (6.20) and (6.21) show that the conditional variance of the decoded output is proportional to the conditional variance of the 'un-decoded' output, while, since y and \hat{y} have different means, the mean square value is not.

It is straightforward to show that the variance of the error, given x , is equal to the variance of \hat{y} given x , so that

$$\text{var}[\epsilon|x] = \text{var}[\hat{y}|x] = \frac{4c^2}{N^2} \text{var}[y|x]. \quad (6.22)$$

The conditional MSE distortion is

$$\begin{aligned} D(x) &= \text{var}[\hat{y}|x] + E[\hat{y} - x|x]^2 \\ &= \text{var}[\hat{y}|x] + E[\hat{y}|x]^2 - 2xE[y|x] + x^2 \\ &= E[\hat{y}^2|x] - 2xE[\hat{y}|x] + x^2. \end{aligned} \quad (6.23)$$

Alternatively,

$$\begin{aligned} D(x) &= E[(\hat{y} - x)^2|x] \\ &= E[\hat{y}^2|x] - 2xE[\hat{y}|x] + x^2. \end{aligned} \quad (6.24)$$

Substituting Eqns. (6.19) and (6.20) into Eqn. (6.23) and simplifying gives

$$D(x) = \frac{4c^2}{N^2} E[y^2|x] - \frac{4c}{N} E[y|x](x + c) + (x + c)^2. \quad (6.25)$$

Linear mean square error distortion

The average MSE distortion is the expected value of the conditional distortion,

$$\begin{aligned} \text{MSE} &= E[D(x)] \\ &= \int_x P(x)D(x)dx \\ &= E[(\hat{y} - x)^2] \\ &= E[\hat{y}^2] - 2E[x\hat{y}] + E[x^2], \end{aligned} \quad (6.26)$$

as described by Eqn. (6.7) in Section 6.1.2. Substituting Eqn. (6.18) into Eqn. (6.26) and simplifying gives

$$\text{MSE} = \frac{4c^2}{N^2} E[y^2] - \frac{4c}{N} E[xy] + E[x^2] - c^2. \quad (6.27)$$

Correlation coefficient

Consider the linear decoding of Eqn. (6.17). It is straightforward to show that the correlation coefficient between x and \hat{y} is identical for any a and b in this equation. Hence, to find an expression for the correlation coefficient, it is sufficient to set $a = 1$ and $b = 0$, and obtain the correlation coefficient between x and the ‘un-decoded’ output y , which is

$$\rho_{x\hat{y}} = \rho_{xy} = \frac{E[xy] - E[x]E[y]}{\sqrt{\text{var}[x]\text{var}[y]}}. \quad (6.28)$$

Since $E[x] = 0$ and $E[y] = N/2$,

$$\rho_{x\hat{y}} = \rho_{xy} = \frac{E[xy]}{\sqrt{E[x^2] \left(E[y^2] - \frac{N^2}{4} \right)}}. \quad (6.29)$$

Due to its invariance to linear decoding, the correlation coefficient is not a relevant measure when comparing different specific linear decoding schemes. However because of this fact, it does provide a means of measuring the effectiveness of the best possible linear decoding. It will also prove useful as a means of comparing the effectiveness of a nonlinear decoding when compared to a linear decoding.

Having now defined a linear decoding scheme, and derived general expressions for the MSE, SQNR and correlation coefficient, the next three subsections give three different possible ways for specifying the value of c in Eqn. (6.18).

6.3.1 Constant Decoding

If either the signal variance or the noise variance are unknown, decoding that is independent of either is necessary. Since the aim is to minimise the MSE distortion, a naive decoding choice is one that has the same range of possible values of x as the input signal so that,

$$\min \hat{y} = \min x \quad \text{and} \quad \max \hat{y} = \max x. \quad (6.30)$$

However, for signal distributions that have infinite support, such as the Gaussian distribution, it is obvious that there should be no reconstruction point at $\pm\infty$. Hence, one solution for such a source is to set the minimum and maximum reconstruction points to a set number of standard deviations.

6.3.2 Matched Moments Linear Decoding

The previous subsection considers a simple method of decoding for the event that the signal variance is unknown. In this and the following subsection, we consider decoding that depends on knowledge of the variance of the input signal, and the linear correlation coefficient between the input and output. The aim is to find the linear decoding scheme that provides the minimum MSE distortion.

Assume the decoding is linear, as given by Eqn. (6.17). Under the hypothesis that the output moments should be equal to the input moments, a naive approach is to find a and b in Eqn. (6.17) such that $E[\hat{y}] = E[x]$ and $E[\hat{y}^2] = E[x^2]$. For the distributions we are considering, $E[x] = 0$, and $E[y] = N/2$. This immediately imposes $b = -aN/2$. We then have

$$E[\hat{y}^2] = a^2 \left(E[y^2] - \frac{N^2}{4} \right) = E[x^2]. \quad (6.31)$$

This gives

$$\begin{aligned} a &= \sqrt{\frac{E[x^2]}{E[y^2] - \frac{N^2}{4}}} \\ &= \sqrt{\frac{E[x^2]}{\text{var}[y]}}. \end{aligned} \quad (6.32)$$

Consequently, the decoded output signal is

$$\hat{y} = \sqrt{\frac{E[x^2]}{\text{var}[y]}} \left(y - \frac{N}{2} \right). \quad (6.33)$$

With reference to Eqn. (6.18),

$$c = \frac{N}{2} \sqrt{\frac{E[x^2]}{\text{var}[y]}}, \quad (6.34)$$

and substituting Eqn. (6.34) into Eqn. (6.27) gives

$$\begin{aligned} \text{MSE} &= 2 \left(E[x^2] - \frac{\sqrt{E[x^2]}E[xy]}{\sqrt{\text{var}[y]}} \right) \\ &= 2E[x^2] \left(1 - \frac{E[xy]}{\sqrt{E[x^2]\text{var}[y]}} \right) \\ &= 2E[x^2](1 - \rho_{xy}). \end{aligned} \quad (6.35)$$

Hence, the MSE of this decoding scheme can be simply expressed in terms of the mean square value of the input signal, and the correlation coefficient of a linear decoding

scheme. Furthermore, rearranging Eqn. (6.35) gives an expression for the SQNR that is independent of $E[x^2]$ as

$$\text{SQNR} = \frac{1}{2(1 - \rho_{xy})}. \quad (6.36)$$

6.3.3 Wiener Optimal Linear Decoding

What is the optimal linear decoding scheme? The only unknown in Eqn. (6.18) is c . The optimal value of c can be found by differentiating Eqn. (6.27) with respect to c , setting the result to zero, and solving for c , the result being

$$c = \frac{NE[xy]}{2\text{var}[y]}. \quad (6.37)$$

Such decoding is known as Wiener decoding, and is the optimal linear decoding scheme for the MSE distortion (Yates and Goodman 2005). Note that with this decoding, it is straightforward to show that

$$E[\hat{y}^2] = E[x\hat{y}] = \frac{E[xy]^2}{\text{var}[y]}. \quad (6.38)$$

Substituting Eqn. (6.37) into Eqn. (6.27) and simplifying gives

$$\begin{aligned} \text{MSE} &= E[x^2] - \frac{E[xy]^2}{\text{var}[y]} \\ &= E[x^2] - E[\hat{y}^2]. \end{aligned} \quad (6.39)$$

Thus, with this decoding, the MSE distortion is independent of the correlation between x and y . We have also that

$$\begin{aligned} \text{MSE} &= E[x^2] \left(1 - \frac{E[xy]^2}{E[x^2]\text{var}[y]} \right) \\ &= E[x^2](1 - \rho_{xy}^2). \end{aligned} \quad (6.40)$$

Eqn. (6.40) can be found in Yates and Goodman (2005). Just as with matched-moments linear decoding, the MSE distortion given by Eqn. (6.40) is entirely dependent on the correlation coefficient and the mean square value of the input. It is also straightforward to show that with this decoding the correlation between the encoded output, y , and the error, ϵ , is zero, that is, $E[\epsilon y] = 0$ (Yates and Goodman 2005).

Subtracting Eqn. (6.40) from Eqn. (6.35) gives the difference between the MSE distortions of matched-moments decoding and Wiener decoding as

$$2E[x^2](1 - \rho_{xy}) - E[x^2](1 - \rho_{xy}^2) = E[x^2](\rho_{xy} - 1)^2 > 0. \quad (6.41)$$

6.4 Linear decoding for SSR

This shows that the MSE distortion obtained by Wiener decoding is always smaller than that obtained with matched-moments decoding. Since Eqn. (6.40) was derived to give the smallest possible linear MSE distortion, Inequality (6.41) simply acts as verification of this fact.

The SQNR for Wiener decoding can be obtained from rearrangement of Eqn. (6.40) as

$$\text{SQNR} = \frac{1}{1 - \rho_{xy}^2}. \quad (6.42)$$

In decibels, this can be written as

$$10 \log_{10}(\text{SQNR}) = -10 \log_{10}(1 - \rho_{xy}^2). \quad (6.43)$$

6.4 Linear decoding for SSR

Section 6.3 does not say anything specific about the SSR model; it only discusses the MSE, SQNR and correlation coefficient between the signals x and \hat{y} . We assumed the decoded signal, \hat{y} , is given by the linear transform of Eqn. (6.18), where y is a non-deterministic function of x , with a mean of $N/2$ and that x has a mean of zero and a PDF that is an even function of x .

Recall we also assume in this chapter that the input of the SSR model has a zero-meant and even PDF, $P(x)$. Section D.2.1 in Appendix D gives a proof that for the SSR model, and such an input signal, that $E[y] = N/2$. Thus, the equations derived in Section 6.3 apply to the SSR model. The remainder of this section simplifies the equations of Section 6.3 for the SSR model, and then considers some specific cases of signal and noise distributions.

6.4.1 Decoding Measures for SSR

MSE distortion

Expressions for the mean of y given x and the mean square value of y given x , for the SSR model, are given in Section D.1 of Appendix D in terms of $P_{1|x}$, by Eqns. (D.3) and (D.6). Substituting these into Eqn. (6.25) gives the conditional MSE distortion for SSR in terms of $P_{1|x}$ as

$$D(x) = 4c^2 P_{1|x} (1 - P_{1|x}) \left(\frac{1 - N}{N} \right) - 4c P_{1|x} x + (x + c)^2. \quad (6.44)$$

An expression for the correlation between x and y , for the SSR model, is given in Section D.3.3 of Appendix D in terms of $P_{1|x}$, by Eqn. (D.25). Substituting this, and Eqn. (D.17) into Eqn. (6.27) gives the MSE distortion for SSR as

$$\text{MSE} = \frac{4c^2(N-1)}{N} \text{E}[P_{1|x}^2] - 4c \text{E}[xP_{1|x}] + \text{E}[x^2] - \frac{c^2(N-2)}{N}. \quad (6.45)$$

Thus, from Eqn. (6.45), the MSE distortion of a linear decoding of the SSR model can be obtained for any c , provided $\text{E}[P_{1|x}^2]$ and $\text{E}[xP_{1|x}]$ are known.

Recall the change of probability measure transformation used in Chapters 4 and 5, $\tau = F_R(x) = P_{1|x}$, where $F_R(\cdot)$ is the Cumulative Distribution Function (CDF) of the noise. Using this same change of measure, the quantity $\text{E}[P_{1|x}^2]$ can be recognised as the mean square value of the PDF, $Q(\tau)$, that is, $\text{E}[\tau^2]$.

SQNR

Using the definition given by Eqn. (6.9), the SQNR for a linear decoding of the SSR model is

$$\text{SQNR} = \frac{\text{E}[x^2]}{\frac{4c^2(N-1)}{N} \text{E}[P_{1|x}^2] - 4c \text{E}[xP_{1|x}] + \text{E}[x^2] - \frac{c^2(N-2)}{N}}. \quad (6.46)$$

Correlation coefficient

Substituting Eqns. (D.17) and (D.25) into Eqn. (6.29), and recalling $\text{E}[P_{1|x}] = 0.5$, gives the correlation coefficient for SSR and a linear decoding as

$$\rho_{xy} = \rho_{x\hat{y}} = \frac{\sqrt{N} \text{E}[xP_{1|x}]}{\sqrt{\text{E}[x^2] \left((N-1) \text{E}[P_{1|x}^2] - \frac{(N-2)}{4} \right)}}. \quad (6.47)$$

So as with the MSE distortion, the correlation coefficient can be calculated given knowledge of $\text{E}[P_{1|x}^2]$ and $\text{E}[xP_{1|x}]$.

Having now obtained expressions for the MSE distortion, SQNR and correlation coefficient for the SSR model with a linear decoding, it is time to examine these measures for specific signal and noise distributions. We find that exact closed-form analytical expressions are available in certain cases. In other cases, numerical evaluation of the distortion, SQNR or correlation coefficient is required.

6.4.2 Specific Signal and Noise Distributions

In this subsection we derive closed form analytical expressions for $E[P_{1|x}^2]$ and $E[xP_{1|x}]$ for the cases of Gaussian signal and noise, Uniform signal and noise and Laplacian signal and noise. Using these expressions, formulas for the linear decoding correlation coefficient are derived. No such exact expressions can be found for logistic signal and noise, but results obtained using numerical integration are obtained for this distribution. This section does not explicitly state expressions for the MSE distortion or SQNR, but such analytical expressions can easily be found by use of Eqns. (6.45) or (6.46) and the equations stated in this section, once a value of c is specified. Specifying c is left for Sections 6.4.3, 6.4.4 and 6.4.5.

Gaussian signal and noise

For Gaussian signal and noise, the mean square value of $P_{1|x}$ can be derived as follows.

$$\begin{aligned}
 E[P_{1|x}^2] &= \int_{-\infty}^{\infty} P_{1|x}^2 P(x) dx \\
 &= \int_{-\infty}^{\infty} \left(\frac{1}{2} + \frac{1}{2} \operatorname{erf} \left(\frac{x}{\sqrt{2}\sigma_{\eta}} \right) \right)^2 P(x) dx \\
 &= \int_{-\infty}^{\infty} \left(\frac{1}{4} + \frac{1}{2} \operatorname{erf} \left(\frac{x}{\sqrt{2}\sigma_{\eta}} \right) + \frac{1}{4} \operatorname{erf}^2 \left(\frac{x}{\sqrt{2}\sigma_{\eta}} \right) \right) P(x) dx \\
 &= \frac{1}{4} \int_{-\infty}^{\infty} P(x) dx + \frac{1}{2} \int_{-\infty}^{\infty} \operatorname{erf} \left(\frac{x}{\sqrt{2}\sigma_{\eta}} \right) P(x) dx + \frac{1}{4} \int_{-\infty}^{\infty} \operatorname{erf}^2 \left(\frac{x}{\sqrt{2}\sigma_{\eta}} \right) P(x) dx \\
 &= \frac{1}{4} + 0 + \frac{1}{4} \int_{-\infty}^{\infty} \operatorname{erf}^2 \left(\frac{x}{\sqrt{2}\sigma_{\eta}} \right) P(x) dx. \tag{6.48}
 \end{aligned}$$

The second term above is zero, since $P(x)$ is even and $\operatorname{erf}(x)$ is odd. We make use of the following result proven in Section D.4 of Appendix D,

$$\int_{x=-\infty}^{x=\infty} \exp(-a^2 x^2) \operatorname{erf}^2(x) dx = \frac{2}{a\sqrt{\pi}} \arctan \left(\frac{1}{a\sqrt{a^2 + 2}} \right), \tag{6.49}$$

which gives

$$\begin{aligned}
E[P_{1|x}^2] &= \frac{1}{4} + \frac{1}{4\sqrt{2\pi\sigma_x^2}} \int_{-\infty}^{\infty} \operatorname{erf}^2\left(\frac{x}{\sqrt{2}\sigma_\eta}\right) \exp\left(-\frac{x^2}{2\sigma_x^2}\right) dx \\
&= \frac{1}{4} + \frac{\sigma}{4\sqrt{\pi}} \int_{-\infty}^{\infty} \operatorname{erf}^2(\tau) \exp(-\sigma^2\tau^2) d\tau \\
&= \frac{1}{4} + \frac{1}{2\pi} \arctan\left(\frac{1}{\sigma\sqrt{\sigma^2+2}}\right) \\
&= \frac{1}{4} + \frac{1}{2\pi} \arcsin\left(\frac{1}{\sigma^2+1}\right). \tag{6.50}
\end{aligned}$$

The conversion from arctan to arcsin holds provided $\sigma > 0$, which it always does. Hence, $E[P_{1|x}^2]$ is a function of the parameter σ . It is shown in Section D.3.1 of Appendix D that the correlation between the inputs to any two devices—that is, the correlation between $x + \eta_i$ and $x + \eta_j$ —in the SSR model is $\rho_i = \frac{1}{\sigma^2+1}$. Therefore

$$E[P_{1|x}^2] = \frac{1}{4} + \frac{1}{2\pi} \arcsin(\rho_i). \tag{6.51}$$

Note that substituting Eqn. (6.51) into Eqn. (D.13) gives the variance of the output of the SSR model encoding as

$$\operatorname{var}[y] = \frac{N(N-1)}{2\pi} \arcsin(\rho_i) + \frac{N}{4}. \tag{6.52}$$

For a decoding given by $c = N$, so that \hat{y} can take values between $-N$ and N , with a mean of zero, we have

$$\operatorname{var}[\hat{y}] = E[\hat{y}^2] = N + N(N-1) \frac{2}{\pi} \arcsin(\rho_i). \tag{6.53}$$

Eqn. (6.53) is precisely the same as an equation derived in Remley (1966) for the output mean square signal power for Gaussian signal and Gaussian noise in a Digital Multi-beam Steering (DIMUS) sonar array—see also Section 4.2.5 in Chapter 4. This result shows that the DIMUS sonar array model is equivalent to the decoded SSR model. However, the results presented in this chapter are far more general, as nowhere in the DIMUS literature is the output of the DIMUS array considered to be a quantised approximation to a random input signal, nor is the output response considered as a function of varying input noise.

Note that Remley (1966) derives Eqn. (6.53) in a different manner to that given here, by making use of a result known as Van Vleck's relation¹² (Price 1958, Van Vleck and

¹²In Remley (1966), Van Vleck's relation is referred to as "Van Fleck's" relation, even though the paper referenced actually uses the correct spelling of "Van Vleck"—see Van Vleck and Middleton (1966).

6.4 Linear decoding for SSR

Middleton 1966). Van Vleck's result is a formula for the normalised autocorrelation, ρ_2 , of "clipped" Gaussian noise—that is, a one bit quantisation of the noise—in terms of the normalised autocorrelation of the unclipped noise, ρ_1 ,

$$\rho_2(\tau) = \frac{2}{\pi} \arcsin(\rho_1(\tau)). \quad (6.54)$$

When applied to the multi-threshold DIMUS sonar array, with Gaussian signal and noise, ρ_1 is the correlation coefficient between the inputs to any two of the thresholds, that is, the correlation coefficient between $x + \eta_i$ and $x + \eta_j$, where η_i and η_j are the independent Gaussian noise samples on two thresholds. At the output, ρ_2 is the correlation coefficient between the binary variables, $2y_i - 1$ and $2y_j - 1$. The possible values of these variables are -1 and 1 .

Section D.3.2 in Appendix D derives the correlation coefficient between any two thresholds' outputs in the SSR model as Eqn. (D.20). Substituting Eqn. (6.51) into Eqn. (D.20) gives

$$\begin{aligned} \rho_o &= 4\mathbb{E}[P_{1|x}^2] - 1 \\ &= \frac{2}{\pi} \arcsin(\rho_i), \end{aligned} \quad (6.55)$$

which is exactly Van Vleck's result, thus indicating how Remley (1966) derived Eqn. (6.53).

We also have

$$\begin{aligned} \mathbb{E}[xP_{1|x}] &= \int_{-\infty}^{\infty} xP(x)P_{1|x}dx \\ &= \int_{-\infty}^{\infty} xP(x) \left(\int_{-\infty}^x R(\eta)d\eta \right) dx \\ &= \int_{-\infty}^{\infty} R(\eta) \left(\int_{\eta}^{\infty} xP(x)dx \right) d\eta \\ &= \int_{-\infty}^{\infty} R(\eta) \left(\int_{\eta}^{\infty} \frac{x}{\sqrt{2\pi}\sigma_x} \exp\left(-\frac{x^2}{2\sigma_x^2}\right) dx \right) d\eta \\ &= \frac{\sigma_x}{\sqrt{2\pi}} \int_{-\infty}^{\infty} R(\eta) \exp\left(-\frac{\eta^2}{2\sigma_x^2}\right) d\eta \\ &= \frac{\sigma_x}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}\sigma_\eta} \exp\left(-\frac{\eta^2}{2\sigma_\eta^2}\right) \exp\left(-\frac{\eta^2}{2\sigma_x^2}\right) d\eta \\ &= \frac{1}{2\pi\sigma} \int_{-\infty}^{\infty} \exp\left(-\frac{\eta^2}{2} \left(\frac{1+\sigma^2}{\sigma_\eta^2}\right)\right) d\eta, \end{aligned} \quad (6.56)$$

since $\sigma = \sigma_\eta/\sigma_x$. The final integrand is a Gaussian PDF, with variance $\sigma_\eta^2/(1 + \sigma^2)$. Hence the integral from $-\infty$ to $+\infty$ is unity times the normalising factor,

$$E[xP_{1|x}] = \frac{1}{2\pi\sigma} \sqrt{2\pi \left(\frac{\sigma_\eta^2}{1 + \sigma^2} \right)} = \frac{\sigma_x}{\sqrt{2\pi(1 + \sigma^2)}}, \quad (6.57)$$

therefore

$$E[xP_{1|x}] = \frac{\sigma_x}{\sqrt{2\pi}} \sqrt{\rho_i}. \quad (6.58)$$

Interestingly, this expression is a function not just of σ but also σ_x . However, referring to the correlation coefficient, as given by Eqn. (6.47), since $E[x^2] = \sigma_x^2$ for a Gaussian signal, σ_x will be cancelled out of Eqn. (6.47), and therefore the correlation coefficient is a function only of σ . Upon substitution of Eqns. (6.50) and (6.57) into Eqn. (6.47) and simplification,

$$\begin{aligned} \rho_{xy} &= \frac{\sqrt{N}}{\sqrt{\sigma^2 + 1} \sqrt{(N-1) \arcsin\left(\frac{1}{\sigma^2+1}\right) + \frac{\pi}{2}}} \\ &= \frac{\sqrt{N\rho_i}}{\sqrt{(N-1) \arcsin(\rho_i) + \frac{\pi}{2}}} \\ &= \frac{\sqrt{N \sin(\rho_o)}}{\sqrt{(N-1)\rho_o + 1}}. \end{aligned} \quad (6.59)$$

However, referring to the MSE distortion given by Eqn. (6.45), σ_x will not necessarily be eradicated from the MSE distortion. In fact, we will see later in this section that in general, the MSE will be a function of the mean square value of the signal.

Uniform signal and noise

For uniform signal and noise, it is necessary to derive expressions for $E[P_{1|x}^2]$ and $E[xP_{1|x}]$ separately for the cases of $\sigma \leq 1$ and $\sigma \geq 1$. The two expressions will be equal at $\sigma = 1$.

For $\sigma \leq 1$,

$$\begin{aligned} E[P_{1|x}^2] &= \int_{-\infty}^{\infty} P_{1|x}^2 P(x) dx \\ &= \frac{1}{\sigma_x} \int_{-\sigma_\eta/2}^{\sigma_\eta/2} \left(\frac{1}{2} + \frac{x}{\sigma_\eta} \right)^2 dx + \frac{1}{\sigma_x} \int_{\sigma_\eta/2}^{\sigma_x/2} dx \\ &= \frac{1}{2} - \frac{\sigma}{6}, \end{aligned} \quad (6.60)$$

and

$$\begin{aligned}
 E[xP_{1|x}] &= \int_{-\infty}^{\infty} xP(x)P_{1|x}dx \\
 &= \int_{-\sigma_{\eta}/2}^{\sigma_{\eta}/2} x \left(\frac{1}{2} + \frac{x}{\sigma_{\eta}} \right) \frac{1}{\sigma_x} dx + \int_{\sigma_{\eta}/2}^{\sigma_x/2} \frac{x}{\sigma_x} dx \\
 &= \frac{\sigma_x}{24} (3 - \sigma^2). \tag{6.61}
 \end{aligned}$$

For $\sigma \geq 1$

$$\begin{aligned}
 E[P_{1|x}^2] &= \int_{-\infty}^{\infty} P_{1|x}^2 P(x) dx \\
 &= \frac{1}{\sigma_x} \int_{-\sigma_x/2}^{\sigma_x/2} \left(\frac{1}{2} + \frac{x}{\sigma_{\eta}} \right)^2 dx \\
 &= \frac{1}{4} + \frac{1}{12\sigma^2}, \tag{6.62}
 \end{aligned}$$

and

$$\begin{aligned}
 E[xP_{1|x}] &= \int_{-\infty}^{\infty} xP(x)P_{1|x}dx \\
 &= \frac{1}{\sigma_x} \int_{-\sigma_x/2}^{\sigma_x/2} x \left(\frac{1}{2} + \frac{x}{\sigma_{\eta}} \right) dx \\
 &= \frac{\sigma_x}{12\sigma}. \tag{6.63}
 \end{aligned}$$

Thus, the correlation coefficient for uniform signal and noise is

$$\rho_{x,y} = \begin{cases} \frac{\sqrt{N}(3-\sigma^2)}{2\sqrt{\sigma(2-2N)+3N}} & (\sigma \leq 1), \\ \frac{\sqrt{N}}{\sqrt{3\sigma^2+N-1}} & (\sigma \geq 1), \end{cases} \tag{6.64}$$

which again is only dependent on the ratio σ .

Laplacian signal and noise

For Laplacian signal and noise, the mean square value of $P_{1|x}$ can be derived as

$$E[P_{1|x}^2] = \frac{\sigma^2 + 3\sigma + 4}{4(\sigma + 1)(\sigma + 2)}. \tag{6.65}$$

Substituting Eqn. (6.65) into Eqn. (D.20) gives the correlation coefficient between the output of any two thresholds as

$$\rho_o = \frac{2}{(\sigma + 1)(\sigma + 2)}. \tag{6.66}$$

The expected value of $xP_{1|x}$ can be derived as

$$E[xP_{1|x}] = \frac{\sigma_x(2\sigma + 1)}{2\sqrt{2}(\sigma + 1)^2}. \quad (6.67)$$

Substitution of Eqns. (6.65) and (6.67) into Eqn. (6.47) gives the linear decoding correlation coefficient for Laplacian signal and noise as

$$\begin{aligned} \rho_{xy} &= \frac{\sqrt{N}(2\sigma + 1)\sqrt{(\sigma + 1)(\sigma + 2)}}{\sqrt{2}(\sigma + 1)^2\sqrt{2N + \sigma(\sigma + 3)}} \\ &= \frac{\sqrt{N}(2\sigma + 1)}{\sqrt{\rho_o(\sigma + 1)^2\sqrt{2N + \sigma(\sigma + 3)}}}. \end{aligned} \quad (6.68)$$

Illustrations

Fig. 6.1 shows the correlation coefficient between the output of any two threshold devices, ρ_o , the quantity, $E[P_{1|x}^2]$ and the quantity, $E[xP_{1|x}]$ for each of four cases of matched signal and noise distributions. The Gaussian, uniform and Laplacian cases were plotted from the exact formulas derived in this section, while the logistic case was calculated by numerical integration. Fig. 6.1(a) also shows the correlation coefficient between the inputs of any two comparators, ρ_i , which is the same for all distributions. All quantities are decreasing functions of σ . Note from Fig. 6.1(a) how ρ_i is always greater than ρ_o , which shows that the one bit quantisation that results from thresholding decreases the correlation between the signals at any two devices. Fig. 6.1(b) shows that at $\sigma = 1$, $E[P_{1|x}^2]$ —and consequently, ρ_o —is identical in all cases.

Fig. 6.2 shows the linear correlation coefficient for the Gaussian, uniform, Laplacian and logistic cases. Except for the logistic case, the correlation coefficient was calculated from the exact formulas specified in this section. The logistic case was calculated by numerical integration of $E[P_{1|x}^2]$ and $E[xP_{1|x}]$. It is clear in all cases that the correlation coefficient has a maximum for a nonzero value of σ . The optimal value of σ increases with increasing N . For the Gaussian, Laplacian and logistic cases, the maximum value of the correlation coefficient occurs for $\sigma > 1$ for sufficiently large N , while for the uniform case, the optimal value of σ appears to get closer and closer to unity as N increases. The results shown in Fig. 6.2 illustrate that stochastic resonance occurs in the linear decoding correlation coefficient for the SSR model.

To compare the linear correlation coefficient for different distributions, Fig. 6.3 shows ρ_{xy} for the four different matched signal and noise distributions for $N = 127$. The uniform case gives the largest correlation coefficient, however the same qualitative behaviour can be seen in each case.

6.4 Linear decoding for SSR

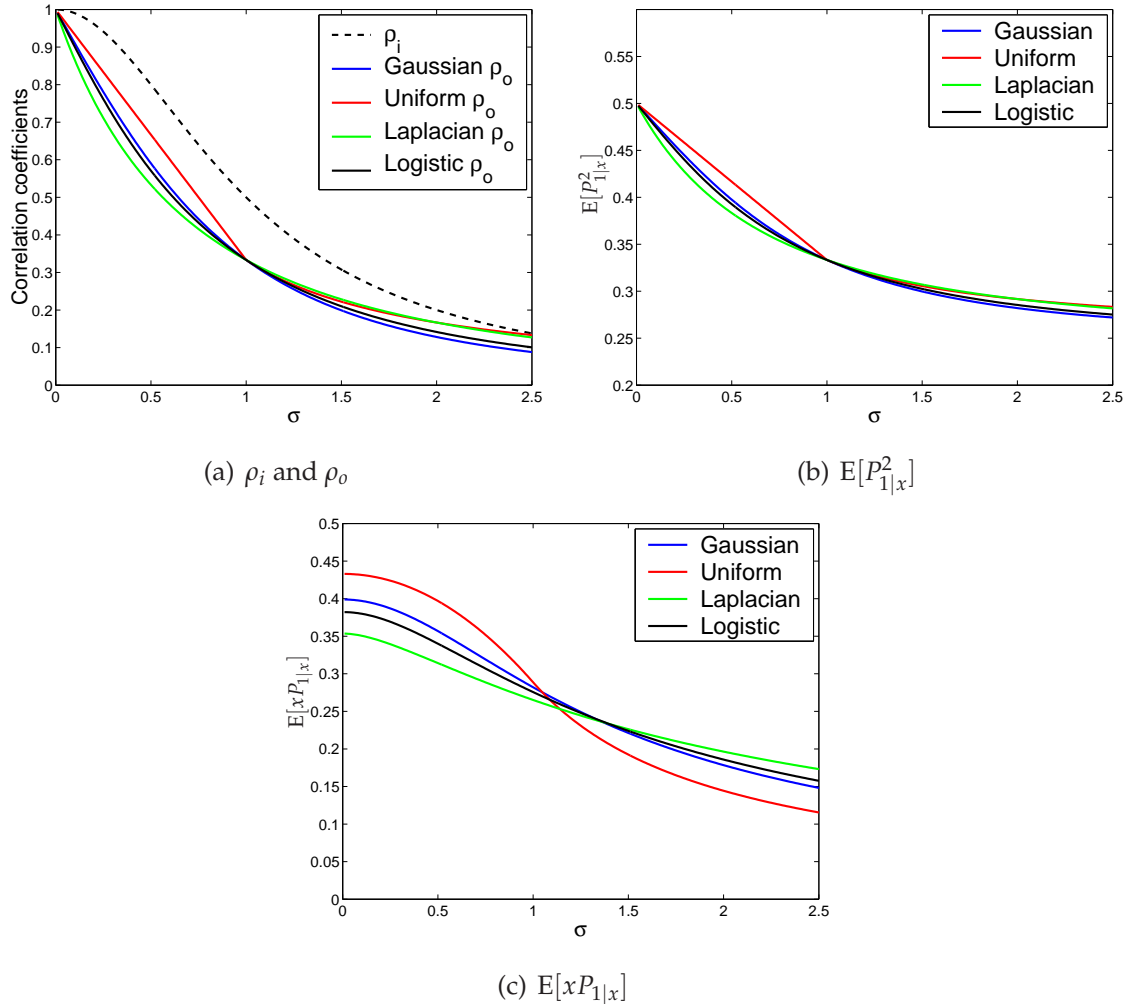


Figure 6.1. Quantities for linear decoding. This figure shows the correlation coefficient between the output of any two threshold devices, ρ_o , the quantity, $E[P_{1|x}^2]$ and the quantity, $E[xP_{1|x}]$ for four cases of matched signal and noise distributions. The Gaussian, uniform and Laplacian cases were calculated from the exact formulas, and the logistic case was calculated by numerical integration. Note that Fig. 6.1(a) also shows the correlation coefficient between the inputs of any two comparators, ρ_i . All quantities are decreasing functions of σ . Note how ρ_i is always greater than ρ_o , which shows that thresholding decreases the correlation between the signals at any two devices. At $\sigma = 1$, $E[P_{1|x}^2]$ —and consequently, ρ_o —is identical in all cases.

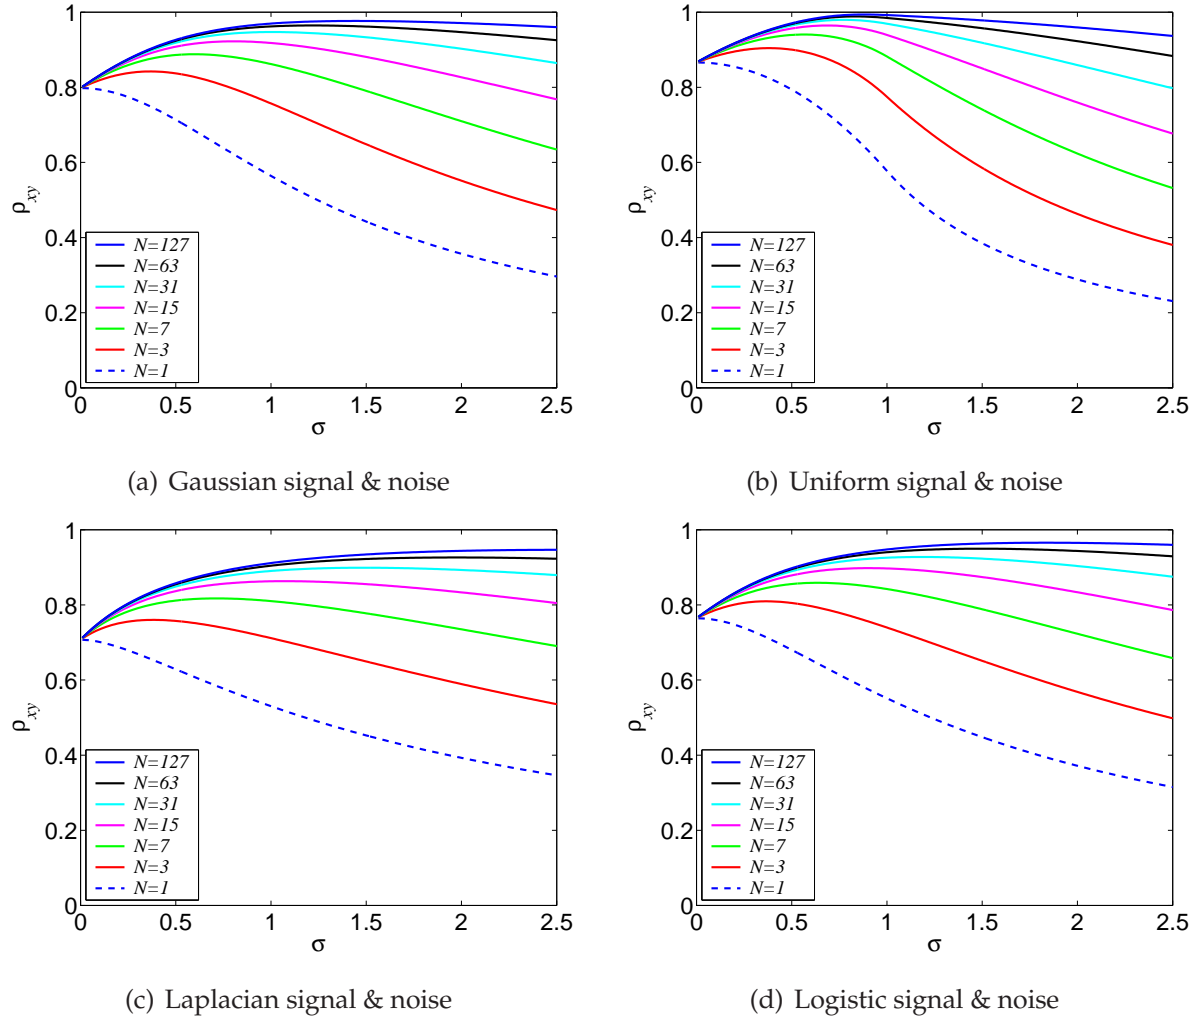


Figure 6.2. Linear decoding correlation coefficient. Plot of the linear decoding correlation coefficient, ρ_{xy} , against increasing σ for four different matched signal and noise distributions and various values of N . Apart from the logistic case, for which ρ_{xy} was calculated only numerically, each plot was calculated from the exact formulas derived in this section, and also validated numerically. Notice how the correlation coefficient increases with increasing N , and has a maximum for a nonzero value of σ . The optimal value of σ also increases with increasing N , and except for the uniform case, becomes larger than unity. In the uniform case, the optimal value of σ gets closer to unity with increasing N .

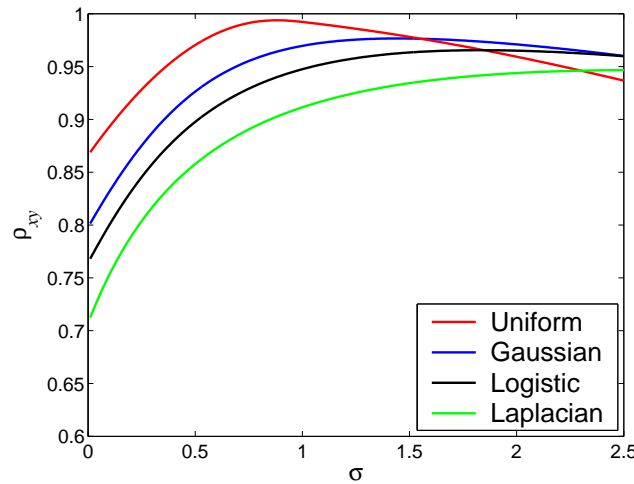


Figure 6.3. Linear decoding correlation coefficient for $N = 127$. This plot shows the linear decoding correlation coefficient, ρ_{xy} , for $N = 127$ and four matched signal and noise distributions. The uniform case gives the largest correlation coefficient, however each case shows the same qualitative behaviour.

6.4.3 Constant Decoding

This section examines the case of sub-optimal constant decoding for the cases of uniform signal and noise, and Gaussian signal and noise. Such a decoding remains constant for all σ and a given signal variance.

Uniform signal and noise

Without loss of generality, we set the support of the uniformly distributed signal to be between $\pm\sigma_x/2$. Thus, we set $c = \sigma_x/2$, so that the decoded output varies between $\pm\sigma_x/2$, as does the input signal. Substitution of Eqns. (6.60), (6.61), (6.62) and (6.63) into Eqn. (6.45) gives the MSE distortion as

$$\text{MSE} = \begin{cases} \sigma_x^2 \left(\frac{\sigma^2+1}{12} - \frac{\sigma(N-1)}{6N} \right) & (\sigma \leq 1), \\ \sigma_x^2 \left(\frac{\sigma^2(3+N)-2N\sigma+N-1}{12N\sigma^2} \right) & (\sigma \geq 1). \end{cases} \quad (6.69)$$

From Eqn. (6.69), for a given non-zero value of σ_x , the MSE distortion for $\sigma \leq 1$ is a quadratic function of σ and has a minimum of

$$\text{MSE}_o = \frac{\sigma_x^2(2N-1)}{12N^2}, \quad (6.70)$$

at a minimising σ of $\sigma_o = (N-1)/N$. It is clear from Eqn. (6.70) that the MSE will decrease with increasing N .

For $\sigma \geq 1$, it is straightforward to show that the MSE of Eqn. (6.69) is strictly increasing. Hence, the MSE is minimised for a non-zero value of σ , just as Chapter 4 showed that the mutual information is maximised for a non-zero value of σ . Here, σ_o is independent of the size of the signal variance, however the minimum distortion is a function of σ_x^2 , and therefore the signal variance.

An alternative constant decoding is one that sets the maximum and minimum reconstruction points slightly smaller than the maximum and minimum possible signal values. Such a decoding provides the same reconstruction points that are optimal for an optimally quantised uniform source, when no input noise is present. This situation will be discussed in more detail in Chapter 8, but for now we simply let $c = \frac{\sigma_x N}{2(N+1)}$, and derive the MSE distortion for this decoding as

$$\text{MSE} = \begin{cases} c^2 \left(\frac{3N-2\sigma(N-1)}{3N} \right) + c\sigma_x \left(\frac{\sigma^2-3}{6} \right) + \frac{\sigma_x^2}{12} & \sigma \leq 1 \\ \frac{4c^2(N-1+3\sigma^2)-4cN\sigma_x\sigma+\sigma^2\sigma_x^2N}{12N\sigma^2} & \sigma \geq 1. \end{cases} \quad (6.71)$$

This has a minimum at $\sigma_o = 2(N-1)c/(N\sigma_x) = (N-1)/(N+1)$, which gives a minimum MSE distortion at this value of

$$\text{MSE}_o = \frac{\sigma_x^2}{12(N+1)^2} \left(\frac{2N^2 - N + 1}{N + 1} \right). \quad (6.72)$$

Again, the MSE is minimised for a nonzero value of σ , while the MSE will get smaller with increasing N .

Fig. 6.4 shows the MSE distortion for the two constant decoding schemes given by Eqns. (6.69) and (6.71), with $\sigma_x = 1$. It is clear that the second decoding gives a smaller MSE, although as N increases, the difference between the MSE distortion for the two decoding schemes becomes smaller. It is also clear from this figure that the optimal value of σ gets close to unity as N increases, while the MSE distortion becomes closer to zero at this value of σ , as the theory predicted.

Gaussian signal and noise

Choosing an appropriate value of c is more problematical for signals with infinite support PDFs, such as a Gaussian signal. One possible method is to select c such that the maximum and minimum output values are equal to a certain number of standard deviations of the input signal. Hence, let $c = k\sigma_x$, where k is the number of standard deviations desired. Substitution of Eqns. (6.58) and (6.51) into Eqn. (6.45) gives

$$\text{MSE} = \sigma_x^2 \left(\frac{k^2}{N} + \frac{2(N-1)}{\pi N} \arcsin(\rho_i) - \frac{2\sqrt{2}k}{\sqrt{\pi}} \sqrt{\rho_i + 1} \right). \quad (6.73)$$

6.4 Linear decoding for SSR




Figure 6.4. Constant decoding MSE distortion for uniform signal and noise. This figure shows the MSE distortion for uniform signal and noise, and the two constant decoding schemes given by Eqns. (6.69) and (6.71), with $\sigma_x = 1$. The optimal value of σ is nonzero, and thus SR occurs, for $N > 1$, as expected. As N increases, the MSE distortion at the optimal σ gets closer to zero, and the optimal value of σ gets closer to unity. Fig. 6.4(c) shows that the absolute difference between the MSE of each decoding becomes smaller as N increases. Note that if $\sigma_x \neq 1$, then the actual values of the reconstruction points scale in proportion. This is also the case for all similar plots in this chapter.

How to choose k ? One might desire the maximum and minimum values of the output to be about those of the input. In this case, the temptation is to set k to 3 for example. However, the variance of the output will not in general be the same as that of the input. For example, for small σ the most probable output states will be those close to 0 or N . Hence, if the output is decoded with $c = 3\sigma_x$, the output will very often be near $\pm 3\sigma_x$, whereas the input is not often near that value. If k is to remain fixed for all σ , then the best solution appears to be to find the value of k which minimises the distortion as a function of σ , and then using that value of k find the value of σ that minimises the distortion, and then use that value of k for all σ . The end result will be a function of N .

However, as we will see shortly, this procedure is very close to a method for finding the optimal linear decoding scheme as a function of σ . So instead, we only point out that, as with the uniform case, the behaviour of the MSE distortion as σ varies for a fixed k will strongly depend on the actual value of k . Due to this fact, we concentrate on decoding schemes that are not constant for all σ , as described in the next two subsections.

6.4.4 Matched Moments Linear Decoding

The previous subsection only considers decoding schemes that have reproduction points that are constant for all σ . Section 6.3.2 derives an expression for a linear decoding that ensures the first two moments of the decoded output match the first two moments of the input distribution. This decoding is a function of σ , and the value of c is given by Eqn. (6.34). The MSE distortion is a function of the linear correlation coefficient, as given by Eqn. (6.35). Given these expressions, and the expressions for $E[P_{1|x}^2]$ and $E[xP_{1|x}]$ derived earlier for matched Gaussian, uniform and Laplacian signal and noise, the MSE distortion and the decoding, $\hat{y}(n)$, can be calculated exactly. For the logistic case, the MSE distortion and $\hat{y}(n)$ can be calculated numerically.

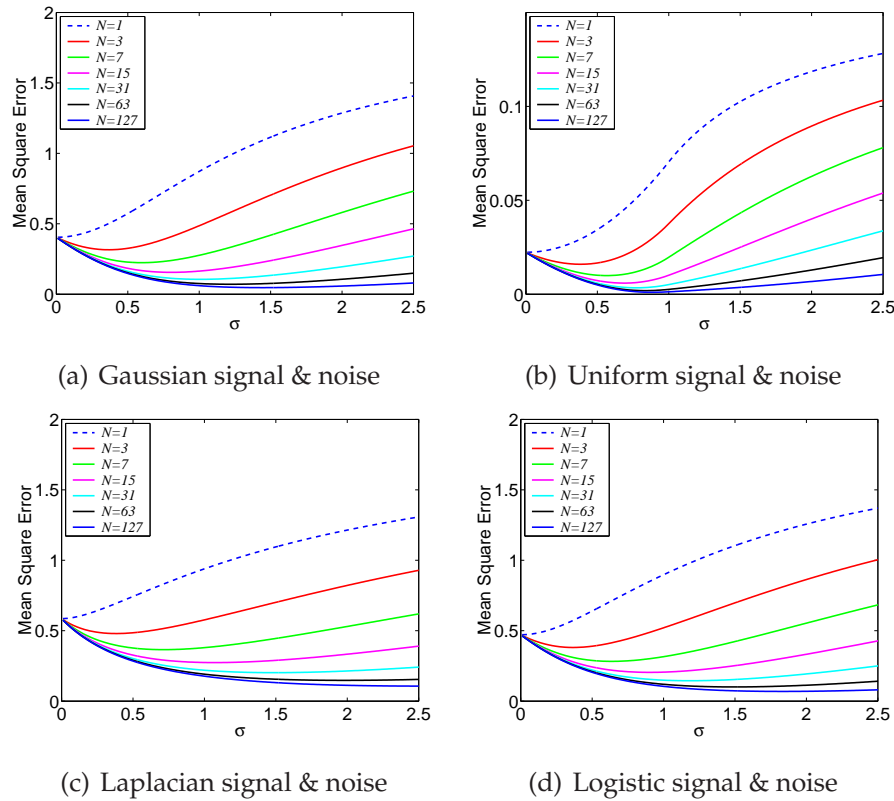


Figure 6.5. MSE distortion for matched moments decoding. This figure shows plots of the MSE distortion against increasing σ for various values of N using matched moments linear MSE decoding, with $\sigma_x = 1$. The optimal value of σ is nonzero, and thus SR occurs, as expected. As N increases, the MSE distortion at the optimal σ gets closer to zero, and the optimal value of σ increases, other than for the uniform case, to a value larger than unity.

6.4 Linear decoding for SSR

Fig. 6.5 shows the MSE distortion for matched-moments decoding for each of the four matched signal and noise distributions considered in this chapter, with $\sigma_x = 1$. Note that for the uniform case, the MSE distortion is far smaller than for the other cases. This is due to the fact that the variance of a uniform signal that has support between $\pm\sigma_x/2$ is $\text{var}[x] = \sigma_x^2/12$, whereas the variance for the other distributions is σ_x^2 .

Fig. 6.6 shows for each case and $N = 15$ the actual values of the decoded output, $\hat{y}(n)$, as a function of σ , with $\sigma_x = 1$. Again, the maximum and minimum values of $\hat{y}(n)$ are much smaller for the uniform case, for the same reasons as the MSE distortion. It can be seen that the reconstruction points are evenly spaced for any given value of σ .

6.4.5 Wiener Optimal Linear Decoding

Section 6.3.2 derives an expression that minimises the MSE distortion for a linear decoding. This decoding is a function of σ , and the value of c is derived in Eqn. (6.37). As shown by Eqn. (6.40), the MSE distortion is a function of the linear correlation coefficient. As with matched-moments decoding, given these expressions, and the expressions for $E[P_{1|x}^2]$ and $E[xP_{1|x}]$ derived earlier for matched Gaussian, uniform and Laplacian signal and noise, the MSE distortion and the decoding, $\hat{y}(n)$, can be calculated exactly. For the logistic case, the MSE distortion and $\hat{y}(n)$ can be calculated numerically.

Fig. 6.7 shows the optimal linear MSE distortion for each of the four matched signal and noise distributions considered in this chapter and $\sigma_x = 1$. Fig. 6.8 shows for each case and $N = 15$ the actual values of the decoded output, $\hat{y}(n)$, as a function of σ . It can be seen that the reconstruction points are evenly spaced for any given value of σ .

As shown by Inequality (6.41), Wiener decoding always gives a smaller MSE distortion than matched moments decoding. Furthermore, Wiener decoding was derived as the decoding that gives the minimum MSE distortion of all linear decoding schemes. However, as mentioned earlier, it is possible to obtain smaller MSE distortions by using nonlinear decoding schemes, that is, reconstruction points that are not a linear function of n , and are therefore not evenly spaced. This is the focus of the next section.




Figure 6.6. Reconstruction points for matched moments linear decoding, $N = 15$. Plot of the reconstruction points, $\hat{y}(n)$, against increasing σ , for $N = 15$ and $\sigma_x = 1$, using matched moments linear MSE decoding. The qualitative behaviour is the same for each distribution; the difference between the maximum and minimum reconstruction points increases with increasing σ . Note that if $\sigma_x \neq 1$, then the actual values of the reconstruction points scale in proportion. This is also the case for all similar plots in this chapter.

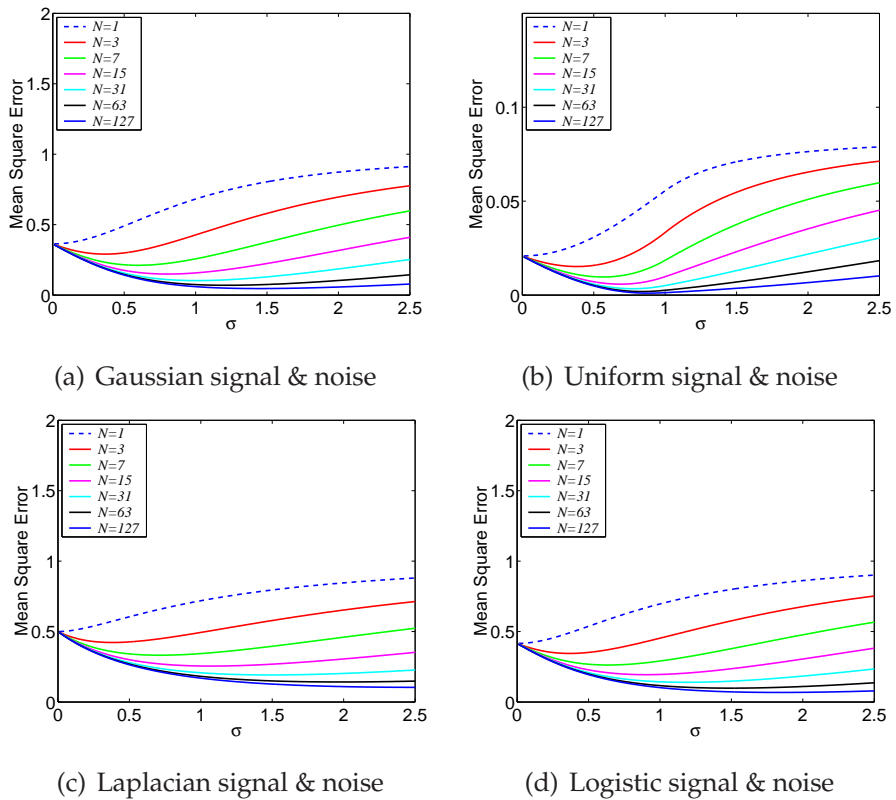


Figure 6.7. MSE distortion for optimal linear (Wiener) decoding. Plot of the MSE distortion against increasing σ for various values of N and $\sigma_x = 1$, using the optimal linear, or Wiener, MSE decoding. The optimal value of σ is nonzero, and thus SR occurs, as expected. As N increases, the MSE distortion at the optimal σ gets closer to zero, and the optimal value of σ increases, other than for the uniform case, to a value larger than unity.




Figure 6.8. Reconstruction points for optimal linear (Wiener) decoding, $N = 15$. Plot of the reconstruction points against increasing σ for $N = 15$ and $\sigma_x = 1$, using the optimal linear, or Wiener MSE decoding. The qualitative behaviour is the same for each distribution; the difference between the maximum and minimum reconstruction points increases with increasing σ .

6.5 Nonlinear Decoding Schemes

Unlike the linear decoding schemes of Sections 6.3 and 6.4, the use of a nonlinear decoding scheme produces a reconstruction of the input signal with reconstruction points that are not equally spaced. We will see that such a nonlinear scheme can provide a smaller MSE distortion than any linear scheme.

This section provides the theory behind three commonly used methods of using nonlinear decoding to estimate a signal from a nondeterministic measurement of it. The applicability of these methods depend on whether knowledge of various probability distributions is available. One of the probability distributions required is the *backward conditional probability distribution*, and we firstly define this distribution, before considering each method, and their application to the decoding of the SSR model.

Backward conditional probabilities

Recall that the key distribution for calculating the mutual information or MSE distortion is the transition probabilities, $P(n|x)$ where $n \in 0, \dots, N$. This section will also use the *Backward Conditional Probability Distribution* (BCPD), $P(x|n)$. The BCPD is also known as the *a posteriori* distribution. The following formula gives the relationship between this distribution and the input and output distributions,

$$P(x) = \sum_{n=0}^N P(x|n)P_y(n). \quad (6.74)$$

In general, the only means available for calculating the BCPD is to use Eqns. (4.3) and (4.4) from Section 4.3 in Chapter 4, after firstly calculating the transition probabilities. That is

$$P(x|n) = \frac{P(x, y)}{P_y(n)} = \frac{P(x)P(n|x)}{P_y(n)} \quad (6.75)$$

$$= \frac{P(x)P(n|x)}{\int_{-\infty}^{\infty} P(n|x)P(x)dx}. \quad (6.76)$$

6.5.1 MAP Decoding

The *Maximum A Posteriori* (MAP) criteria (Yates and Goodman 2005) is often used in estimation applications. Although it does not necessarily provide the optimal estimation, it can be convenient to use and often provides MSE distortions close to optimal.

The MAP criteria states that an estimate x , given observation $y = n$, should be chosen so that the estimate is

$$\hat{y}_n = \arg \max_x P(x|n). \quad (6.77)$$

Given that the *a posteriori* distribution, $P(x|n)$, is given by Eqn. (6.76), this criteria can be simplified to

$$\hat{y}_n = \arg \max_x P(x)P(n|x), \quad (6.78)$$

that is, the MAP estimator, given observation $y = n$ is the value of x that maximises the joint probability density function of x and y for $y = n$.

For SSR, $P(n|x)$ is given by the binomial formula as in Eqn. (4.9), repeated here as

$$P(n|x) = \binom{N}{n} P_{1|x}^n (1 - P_{1|x})^{N-n} \quad n = 0, \dots, N. \quad (6.79)$$

Finding the value of x that maximises the joint distribution can be achieved by differentiating $P(x, y) = P(n|x)P(x)$ with respect to x and setting the result to zero. The derivative of $P(n|x)$ with respect to x is given by Eqn. (B.1) in Appendix B. Applying this formula leads to the following criteria for a stationary point of the joint density function, $P(y, x)$, with respect to x ,

$$\begin{aligned} \frac{d}{dx} (P(x)P(n|x)) &= P'(x)P(n|x) + \\ &P(x)P(n|x) \left(\frac{n - NP_{1|x}}{P_{1|x}(1 - P_{1|x})} \right) R(x) = 0, \quad n = 0, \dots, N. \end{aligned} \quad (6.80)$$

For $P(x)$ and $R(x)$ with infinite support, $P(n|x)$ cannot be zero and hence, Eqn. (6.80) reduces to

$$P'(x) + P(x) \left(\frac{n - NP_{1|x}}{P_{1|x}(1 - P_{1|x})} \right) R(x) = 0, \quad n = 0, \dots, N. \quad (6.81)$$

In general there is no closed form solution to Eqn. (6.81), although it is possible to simplify it for a specified $P(x)$ and solve numerically.

However, in practice, the simplest way to find the MAP decoding for SSR with given signal and noise distributions, is to numerically calculate $P(x, y) = P(x)P(n|x)$, and find the value of x for which $P(x, y)$ is a maximum for each n . This value of x is then the MAP decoding for SSR for each n . The result of this procedure for varying σ and N is shown in Fig. 6.9, which shows the MSE distortion, and Fig. 6.10, which shows the reconstruction points for $N = 15$ for the cases of Gaussian, Laplacian and logistic

6.5 Nonlinear Decoding Schemes

signal and noise with $\sigma_x = 1$. The uniform case is not considered, as there is not always a unique maximum of the joint distribution.

Note that for the Laplacian case and $N = 1$, in Fig. 6.9(b), the MSE distortion saturates at unity for $\sigma \geq 1$. This is due to a general property of the MSE distortion, that a decoding can always be chosen such that $\text{MSE} = \text{E}[x^2]$, which in this case is unity. The decoding that achieves this is to set $\hat{y}_n = \text{E}[x] \forall n$; when $\sigma \geq 1$, the MAP criteria sets the decoding of both output states to be zero, so that the output is *always* zero. Hence, the MSE distortion is identical to the mean square value of the input signal. This is the only case we will see where this happens, but note that for σ sufficiently large, this situation will happen regardless of N and the signal and noise distribution.




Figure 6.9. MSE distortion for maximum *a posteriori* decoding. Plot of the MSE distortion against increasing noise intensity, σ , for various values of array size, N , and signal standard deviation set to $\sigma_x = 1$, using the maximum *a posteriori* decoding. Note that for the Laplacian case and $N = 1$, the MSE distortion saturates at unity for $\sigma \geq 1$, since the MAP criteria sets both output reproduction points to the signal mean of zero, and thus the decoded output is always zero, meaning the MSE distortion is the mean square value of the input, i.e. unity.

6.5.2 Maximum Likelihood Decoding

In contrast with the MAP criteria, the Maximum Likelihood (ML) estimator (Yates and Goodman 2005) does not directly use the *a priori* distribution of the signal to be estimated, that is, $P(x)$. The ML estimator for x given observation $y = n$ is given by

$$\hat{y}_n = \arg \max_x P(n|x), \quad (6.82)$$

that is, the ML estimator, given observation $y = n$ is the value of x that maximises the probability mass function of y given x .




Figure 6.10. Reconstruction points for maximum *a posteriori* decoding, $N = 15$. Plot of the reconstruction points against increasing noise intensity, σ , for $N = 15$, and the signal standard deviation set to $\sigma_x = 1$, using the maximum *a posteriori* decoding.

As with the MAP estimator, the ML estimator cannot be shown to minimise the MSE distortion, but in practice often provides a fairly good approximation to it.

The values of the ML reconstruction points for SSR can be obtained by differentiating the transition probabilities, $P(n|x)$, with respect to x and setting the result to zero. The result, as given in Appendix B by Eqn. (B.1), is

$$\frac{n - NP_{1|x}}{P_{1|x}(1 - P_{1|x})} = 0, \quad n = 0, \dots, N. \quad (6.83)$$

This reduces to $x = F_R^{-1}(n/N)$, and therefore the ML reconstruction points are

$$\hat{y}_n = F_R^{-1}\left(\frac{n}{N}\right), \quad n = 0, \dots, N. \quad (6.84)$$

For $n = N$ and $n = 0$, the transition probabilities are maximised at $\pm\infty$. Hence $\hat{y}_N = \infty$ and $\hat{y}_0 = -\infty$. Clearly this is not satisfactory, since for small σ , the largest output probabilities are $P_y(0)$ and $P_y(N)$. There is no entirely satisfactory way around this problem, although one possibility is to define a decoding which approximates the ML decoding for large N ,

$$\hat{y}_n = F_R^{-1}\left(\frac{n+1}{N+2}\right), \quad n = 0, \dots, N. \quad (6.85)$$

Although this method does not give the ML reconstruction points, as N gets larger, the reconstruction points get closer and closer to the ML points.

A numerical approach could also provide an approximation to the ML solution. For example, if the reproduction points of Eqn. (6.84) for $1 \leq n \leq N-1$ are used, finite points for $n = 0$ and $n = N$ can be found by assuming $\hat{y}_0 = -\hat{y}_N$ and then varying

\hat{y}_N from \hat{y}_{N-1} towards infinity and calculating the MSE distortion for each value of \hat{y}_N . The value of \hat{y}_N that provides the minimum distortion is then used. However, this will turn out to be the optimal reproduction point for $n = 0$ and $n = N$, as discussed in the following subsection. Thus, due to these difficulties, we do not present results for a ML decoding in this chapter. However, Chapter 7, which discusses decoding for large N will show that in certain circumstances, the optimal nonlinear decoding approaches the ML decoding for large N .

6.5.3 Minimum Mean Square Error Distortion Decoding

We label the signal that results from an optimal MSE decoding as \hat{x} . As shown in Section D.5 of Appendix D, the decoding that gives the minimum possible MSE distortion is the decoding that consist of values

$$\hat{x}_n = E_x[x|n] = \int_x xP(x|n)dx = \frac{1}{P_y(n)} \int_x xP(n|x)P(x)dx, \quad n = 0, \dots, N. \quad (6.86)$$

Like the MAP and ML decoding schemes, and unlike the linear decoding schemes considered earlier, this decoding varies for each value of the encoding, y . Since this is the theoretical best MSE distortion decoding, we call the resultant distortion the Minimum Mean Square Error (MMSE). Section D.5 of Appendix D also shows that with this decoding, $E[x\hat{x}] = E[\hat{x}^2]$ and that therefore by inspecting Eqn. (6.7), the MMSE is given by

$$\begin{aligned} \text{MMSE} &= E[x^2] - E[\hat{x}_n^2] \\ &= E[x^2] - \sum_{n=0}^N E[x|n]^2 P_y(n) \\ &= E[x^2] - \sum_{n=0}^N \frac{1}{P_y(n)} \left(\int_x xP(n|x)P(x)dx \right)^2. \end{aligned} \quad (6.87)$$

Good references that discuss such a decoding include Gershenfeld (1999) and Yates and Goodman (2005). Given the initial motivation of SSR as a model for a population of neurons, note also that MMSE decoding has also previously been used in computational neuroscience research, for example, Bethge *et al.* (2002). Just as we pointed out for the case of the optimal linear decoding, the encoded output, y , is uncorrelated with the error, $\epsilon = x - \hat{x}$ for MMSE decoding. This is proven in Section D.5 of Appendix D.

Section D.5 of Appendix D also shows that $E[\hat{x}] = 0$. Hence, the correlation coefficient between the decoded output signal, \hat{x} , and the input signal, x , is

$$\begin{aligned}\rho_{x,\hat{x}} &= \frac{E[x\hat{x}] - E[x]E[\hat{x}]}{\sqrt{\text{var}[x]\text{var}[\hat{x}]}} \\ &= \frac{E[x\hat{x}]}{\sqrt{E[x^2]E[\hat{x}^2]}} \\ &= \sqrt{\frac{E[\hat{x}^2]}{E[x^2]}}.\end{aligned}\quad (6.88)$$

Rearranging Eqn. (6.88) shows that the MMSE can be written in terms of the correlation coefficient as

$$\text{MMSE} = E[x^2] \left(1 - \rho_{x,\hat{x}}^2\right). \quad (6.89)$$

Therefore, as with the optimal linear decoding scheme, given by Wiener decoding, the MMSE distortion can be written in terms of the mean square value of the input signal, and the correlation coefficient between the input signal and the decoded output signal. The difference between Eqns. (6.89) and (6.40) is that Eqn. (6.40) gives the MSE in terms of the correlation coefficient for *linear* decoding, and Eqn. (6.89) gives the MSE in terms of the correlation coefficient for optimal *nonlinear* decoding.

Furthermore, as with optimal linear decoding, the SQNR for MMSE decoding can be written in terms of the correlation coefficient as

$$\text{SQNR} = \frac{1}{1 - \rho_{x,\hat{x}}^2}. \quad (6.90)$$

Fig. 6.11 shows the MMSE distortion for various N and increasing σ , and Fig. 6.12 shows the reconstruction points against increasing σ for $N = 15$, for the cases of Gaussian, Laplacian and logistic signal and noise, and $\sigma_x = 1$. The case of uniform signal and noise will be considered shortly. As N increases, the MMSE distortion at the optimal σ gets closer to zero, and the optimal value of σ increases, to a value larger than unity. However, for increasing $\sigma \geq 1$, the rate of decrease of the MMSE is quite slow, and gets slower for increasing N .

If the reconstruction points shown in Fig. 6.12 are compared with those of the MAP decoding shown in Fig. 6.10, it can be seen that the points are almost identical, apart from the largest and smallest ones for any given σ , which correspond to $n = 0$ and $n = N$.

6.5 Nonlinear Decoding Schemes




Figure 6.11. MSE distortion for MMSE decoding. Plot of the MSE distortion against increasing σ for various values of N using MMSE decoding and $\sigma_x = 1$. The optimal value of σ is nonzero, and thus SR occurs, as expected. As N increases, the MSE distortion at the optimal σ gets closer to zero, and the optimal value of σ increases, to a value larger than unity. For increasing $\sigma \geq 1$, the rate of decreases of the MMSE is quite slow, and gets slower for increasing N .




Figure 6.12. Reconstruction points for MMSE decoding, $N = 15$. Plot of the reconstruction points against increasing σ for various values of N using MMSE decoding and $\sigma_x = 1$. When compared with the reconstruction points obtained with the MAP decoding shown in Fig. 6.10, it can be seen that the points are almost identical, apart from the largest and smallest ones for any given σ , which correspond to $n = 0$ and $n = N$.

Fig. 6.13 shows the optimal correlation coefficient given by Eqn. (6.88) for various N , against increasing σ , for Gaussian, Laplacian and logistic signal and noise. As N increases, the correlation coefficient gets quite close to unity, especially for $\sigma \geq 1$. As with the MMSE distortion, the rate of decrease of the correlation coefficient as σ increases past unity is quite slow, and gets slower with increasing N . A comparison with the correlation coefficient for a linear decoding is left for the next subsection. Firstly, however we derive an exact result for MMSE decoding reconstruction points and MMSE distortion for uniform signal and noise, and $\sigma \leq 1$.




Figure 6.13. Correlation coefficient for MMSE decoding. Plot of MMSE decoding correlation coefficient against increasing σ for various values of N . As N increases, the correlation coefficient gets quite close to unity, especially for $\sigma \geq 1$. As with the MMSE, the rate of decrease of the correlation coefficient as σ increases past unity is quite slow, and gets slower with increasing N .

Exact MMSE decoding for uniform signal and noise

We now derive exact expressions for the MMSE distortion and reconstruction points for uniform signal and noise, and $\sigma \leq 1$. As discussed in Chapter 4, the output probability mass function, in this case, for $N > 1$, is

$$P_y(n) = \frac{\sigma}{N+1}, \quad n = 1, \dots, N-1 \quad (6.91)$$

and

$$P_y(0) = P_y(N) = \sigma \left(\frac{1}{N+1} - \frac{1}{2} \right) + \frac{1}{2} = \frac{1}{\sigma_x} F_R^{-1} \left(\frac{1}{N+1} \right) + \frac{1}{2}. \quad (6.92)$$

Note that if $N = 1$ then $P_y(0) = P_y(N) = 0.5$.

The optimal reproduction points, \hat{x}_n , are given by Eqn. (6.86). We define for convenience $A_n = \int_x x P(n|x) P(x) dx$, $n = 0, \dots, N$. Using Eqn. (4.9) from Chapter 4, and

6.5 Nonlinear Decoding Schemes

Eqn. (B.14) from Appendix B, integration of the RHS of Eqn. (6.86) gives

$$A_n = \frac{\sigma_x}{2} \sigma^2 \left(\frac{2n - N}{N^2 + 3N + 2} \right), \quad n = 1, \dots, N - 1, \quad N > 1 \quad (6.93)$$

$$A_N = \int_x x P(N|x) P(x) dx = \frac{\sigma_x}{8} \left(1 - \sigma^2 \left(\frac{N^2 - N + 2}{N^2 + 3N + 2} \right) \right), \quad (6.94)$$

and

$$A_0 = \int_x x P(0|x) P(x) dx = -\frac{\sigma_x}{8} \left(1 - \sigma^2 \left(\frac{N^2 - N + 2}{N^2 + 3N + 2} \right) \right), \quad (6.95)$$

This gives

$$\hat{x}_n = \frac{\sigma_x}{2} \sigma \left(\frac{2n - N}{N + 2} \right) = \frac{\sigma_\eta}{2} \left(\frac{2n - N}{N + 2} \right), \quad n = 1, \dots, N - 1, \quad N > 1 \quad (6.96)$$

and

$$\hat{x}_N = -\hat{x}_0 = \frac{\sigma_x}{4} \left(\frac{1 - \sigma^2 \left(\frac{N^2 - N + 2}{N^2 + 3N + 2} \right)}{1 - \sigma \left(\frac{N - 1}{N + 1} \right)} \right). \quad (6.97)$$

So for all n except $n = 0$ and $n = N$, the optimal decoding can be stated independently of either the signal variance or the noise variance, but not both. However, the ratio \hat{x}/σ_x is a function of only σ , so the optimal reconstruction points scale proportionally to σ_x .

Rearranging Eqn. (6.96) gives

$$\hat{x}_n = \sigma_\eta \left(\frac{n + 1}{N + 2} - \frac{1}{2} \right) = F_R^{-1} \left(\frac{n + 1}{N + 2} \right), \quad n = 1, \dots, N - 1, \quad N > 1. \quad (6.98)$$

Thus, except for the cases of $y = 0$ and $y = N$, the optimal decoding is a linear decoding, as given by Eqn. (6.18) with $c = \frac{N\sigma_\eta}{2(N+2)}$.

Note that when $\sigma = 1$, Eqn. (6.97) reduces to

$$\hat{x}_N = -\hat{x}_0 = \frac{\sigma_x}{2} \left(\frac{N}{N + 2} \right) = F_R^{-1} \left(\frac{N + 1}{N + 2} \right). \quad (6.99)$$

Therefore, when $\sigma = 1$, $\hat{x}_n = F_R^{-1} \left(\frac{n+1}{N+2} \right)$ for all n , and MMSE decoding is a linear decoding, as given by Eqn. (6.18) with $c = \frac{N\sigma_\eta}{2(N+2)}$. Thus, Wiener linear decoding must give the same MSE performance as MMSE decoding at $\sigma = 1$.

From Eqn. (6.87), for $N > 1$ the MMSE distortion is

$$\begin{aligned}
\text{MMSE} &= \frac{\sigma_x^2}{12} - E[\hat{x}^2] \\
&= \frac{\sigma_x^2}{12} - 2P_y(N)\hat{x}_N^2 - \sum_{n=1}^{N-1} P_y(n)\hat{x}_n^2 \\
&= \frac{\sigma_x^2}{12} - 2A_N\hat{x}_N - \sum_{n=1}^{N-1} A_n\hat{x}_n \\
&= \frac{\sigma_x^2}{12} - 2A_N\hat{x}_N - \sum_{n=1}^{N-1} \frac{\sigma_x}{2}\sigma^2 \left(\frac{2n-N}{N^2+3N+2} \right) \frac{\sigma_x}{2}\sigma \left(\frac{2n-N}{N+2} \right) \\
&= \frac{\sigma_x^2}{12} - 2A_N\hat{x}_N - \frac{\sigma_x^2\sigma^3}{4} \frac{1}{(N+1)(N+2)^2} \sum_{n=1}^{N-1} (2n-N)^2 \\
&= \frac{\sigma_x^2}{12} - 2A_N\hat{x}_N - \frac{\sigma_x^2\sigma^3}{12} \frac{N(N-1)(N-2)}{(N+1)(N+2)^2} \\
&= \frac{\sigma_x^2}{12} \left(1 - \frac{\frac{3}{4} \left(1 - \sigma^2 \left(\frac{N^2-N+2}{N^2+3N+2} \right) \right)^2}{1 - \sigma \left(\frac{N-1}{N+1} \right)} - \sigma^3 \frac{N(N-1)(N-2)}{(N+1)(N+2)^2} \right). \quad (6.100)
\end{aligned}$$

Fig. 6.14 shows for uniform signal and noise the MMSE distortion given by Eqn. (6.100), the corresponding MMSE correlation coefficient, and the values of the optimal MMSE reconstruction points given by Eqns. (6.96) and (6.97), with $E[x^2] = 1/12$. Fig. 6.14 also shows these quantities calculated by numerical integration, to verify the exact results for $\sigma \leq 1$, and to show the continuation to $\sigma > 1$.




Figure 6.14. MMSE decoding for uniform signal and noise. Plot of the MMSE distortion, correlation coefficient, and reconstruction points for $N = 15$, for uniform signal and noise with $\sigma_x = 1$. The circles show the values calculated from the exact expressions given for $\sigma \leq 1$ in this Section, and the lines show numerical values. The optimal value of σ is nonzero, and gets closer to unity as N increases.

6.6 Decoding Analysis

The purpose of this section is to compare the performance of all the different decoding schemes considered in this Chapter. Firstly, we will compare the MSE distortion for matched-moments and Wiener linear decoding schemes, and the MAP and MMSE nonlinear decoding schemes, for the case of $N = 127$ and Gaussian, Laplacian, and logistic signal and noise. For the uniform case, we will compare the MSE distortion of matched-moments and Wiener linear decoding, and MMSE nonlinear decoding for $N = 127$. In all cases we plot for $\sigma_x = 1$. We will also compare the linear correlation coefficient with the nonlinear correlation coefficient obtained with MMSE decoding.

Secondly, we briefly look at the *average transfer function* of the decoded SSR model, for $N = 7$ and Gaussian signal and noise.

6.6.1 Distortion Comparison

Fig. 6.15 shows the MSE distortion plotted against increasing σ , for the four decoding schemes for the cases of Gaussian, Laplacian and logistic signal and noise. Firstly, as expected, Wiener linear decoding gives a smaller MSE distortion for all σ than matched-moments decoding. The difference between the MSE distortion for these two decoding schemes is expressed by Eqn. (6.41). Secondly, the optimal MSE distortion, obtained by MMSE decoding, is clearly verified as giving a smaller MSE distortion than the other decoding schemes.

Thirdly, MAP decoding can be seen to have quite poor distortion performance for small σ . However, for σ greater than about 0.1 – 0.2, the MAP decoding MSE distortion is smaller than the MSE distortion for the linear decoding schemes and, although not visible in Fig. 6.15, actually gets closer and closer to the MMSE distortion as σ increases. This verifies the known empirical fact that, although it is not provable, MAP decoding often provides a MSE distortion quite close to optimal.

Fig. 6.16 shows the linear and MMSE nonlinear correlation coefficients plotted against increasing σ . Clearly, the nonlinear correlation coefficient is always greater than the linear correlation coefficient. Interestingly, the difference between the two coefficients gets quite small for $\sigma > 2$ for the Gaussian case. The gap between each is far larger for the Laplacian and logistic cases.



Figure 6.15. Comparison of MSE distortion for various decoding schemes, $N = 127$. Plot of the MSE distortion against increasing σ for all four considered decoding schemes, for $N = 127$ with $\sigma_x = 1$. This figure verifies that Wiener decoding is superior to matched-moments decoding, and that MMSE decoding is superior to all decoding schemes. It is also apparent that the MAP decoding is quite poor for small σ , but for larger σ has a performance very close to that of the optimal MMSE decoding. There are no remarkable differences between each of the signal and noise pairs.

Fig. 6.17 compares the MSE distortion and correlation coefficient for uniform signal and noise and $N = 127$. Again, it is clear that MMSE decoding is optimal, and Wiener decoding is better than matched-moments decoding. However, as predicted in Section 6.5.3, for $\sigma = 1$, the performance of Wiener decoding and MMSE decoding is identical, since at this point, MMSE decoding is a linear function of n .

Recall the definition of SQNR given by Eqn. (6.9). Figs. 6.18(a)-6.18(c) show the SQNR plotted in dB against increasing σ for $N = 127$. As we would expect, the SQNR is maximised for the same value of σ that minimises the MSE distortion.

It is also possible to define an input SNR measure for the SSR model. If the input signal variance is σ_x^2 and the noise signal variance is σ_η^2 , then the input SNR in decibels is

$$\text{SNR}_i = 10 \log_{10} \left(\frac{\sigma_x^2}{\sigma_\eta^2} \right) = -20 \log_{10} (\sigma) \quad \text{dB.} \quad (6.101)$$

Figs. 6.18(d)-6.18(f) show the output SQNR in dB plotted against the input SNR in dB. Since the MSE distortion is minimised for σ near unity, this corresponds to an input SNR around 0 dB. This is clearly illustrated in these subfigures.

6.6 Decoding Analysis




Figure 6.16. Comparison of correlation coefficients for $N = 127$. Comparison of the linear correlation coefficient with the MMSE decoding correlation coefficient, against increasing σ , for $N = 127$. This figure verifies that the correlation coefficient is larger for the nonlinear MMSE decoding than it is for any linear decoding, for all σ .

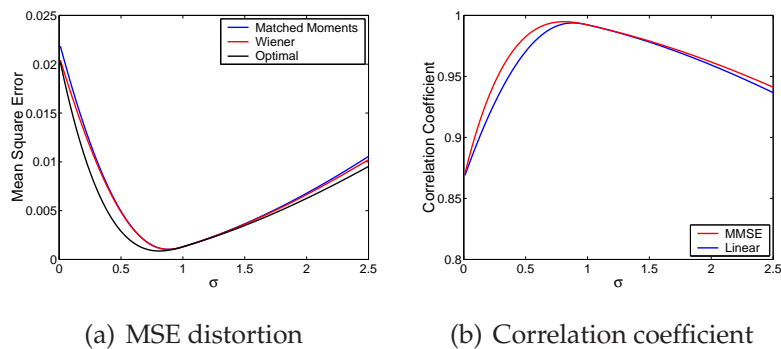


Figure 6.17. Decoding comparison for uniform signal and noise, $N = 127$. Plot of the MSE distortion and correlation coefficient against increasing σ for uniform signal and noise and $N = 127$. This figure verifies that Wiener decoding is superior to matched-moments decoding, and that MMSE decoding is superior to Wiener decoding for all σ . However, it can be seen that for $\sigma = 1$, the MMSE and Wiener decoding schemes have the same performance.

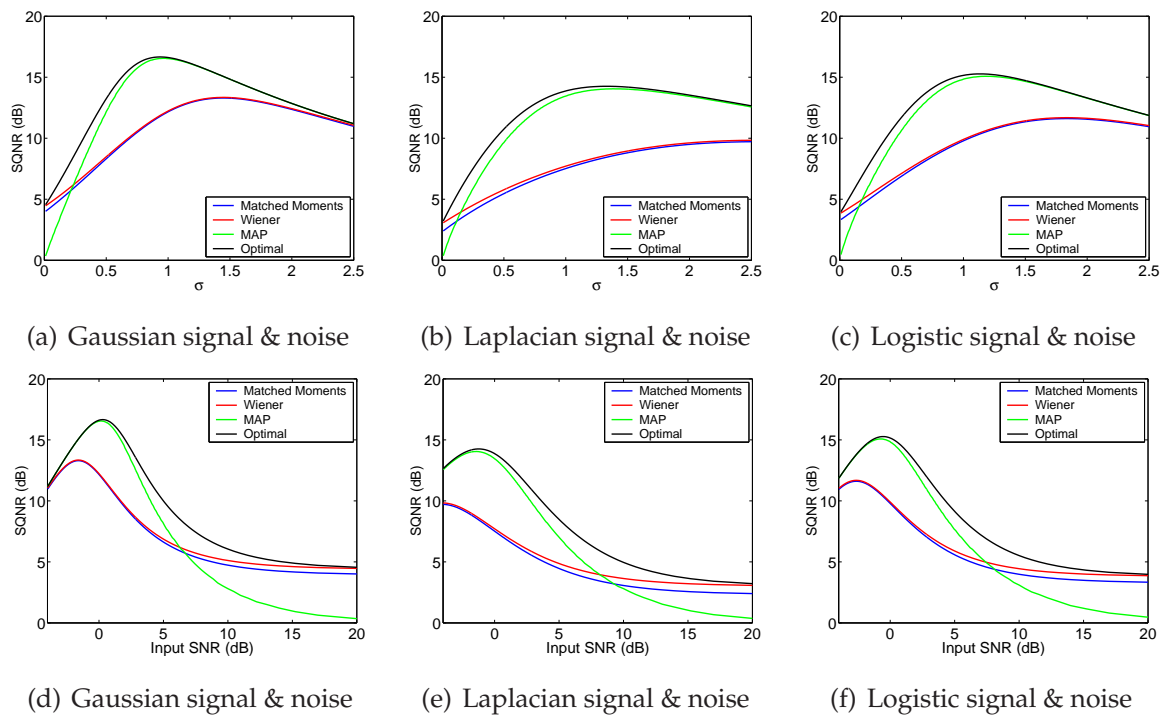


Figure 6.18. SQNR for each decoding for $N = 127$. Plots of the SQNR in dB against increasing σ and the SQNR against increasing input SNR in dB for all four considered decoding schemes, for $N = 127$. As is obviously the case from its definition, the SQNR is maximised for the same value of σ that minimises the MSE distortion, which when expressed as the input SNR is near 0 dB.

6.6.2 Transfer Functions

Recall that in Section 4.5.1 of Chapter 4 that we plot the *average* transfer function of the SSR model's encoding, y , which is the expected value of y given x . We now plot the average transfer function after decoding of the SSR model.

Fig. 6.19(a) shows the *average* transfer function for Wiener linear decoding, with $N = 7$, for various values of noise standard deviation, σ_η , and Gaussian noise. Fig. 6.19(b) shows the average transfer function for MMSE nonlinear decoding. Also shown for comparison, with thick black lines, is the optimal deterministic transfer function for a 3-bit quantisation of a Gaussian source. This plot was obtained using the Lloyd Method I algorithm—see Section 8.5 in Chapter 8 for more details. Notice that the average transfer function of each decoding scheme can be seen to be almost identical, however, the maximum and minimum values of the average output are smaller in magnitude than for the optimal noiseless quantiser.

Fig. 6.19(c) shows the variance corresponding to each value of σ_η as a function of x , for Wiener decoding, and Fig. 6.19(d) shows the variance for MMSE decoding. This variance can be seen to be smaller for more values of x for MMSE decoding, particularly for values of x close to the signal mean.

This section and the previous two sections examined and compared various decoding schemes and the corresponding MSE distortions and correlation coefficients for SSR. The optimal decoding scheme was discussed, and verified to provide a smaller MSE distortion than any other decoding scheme. The next section investigates relationships between this optimal decoding scheme, and various ideas from estimation theory, including Fisher information, and the concept of bias.

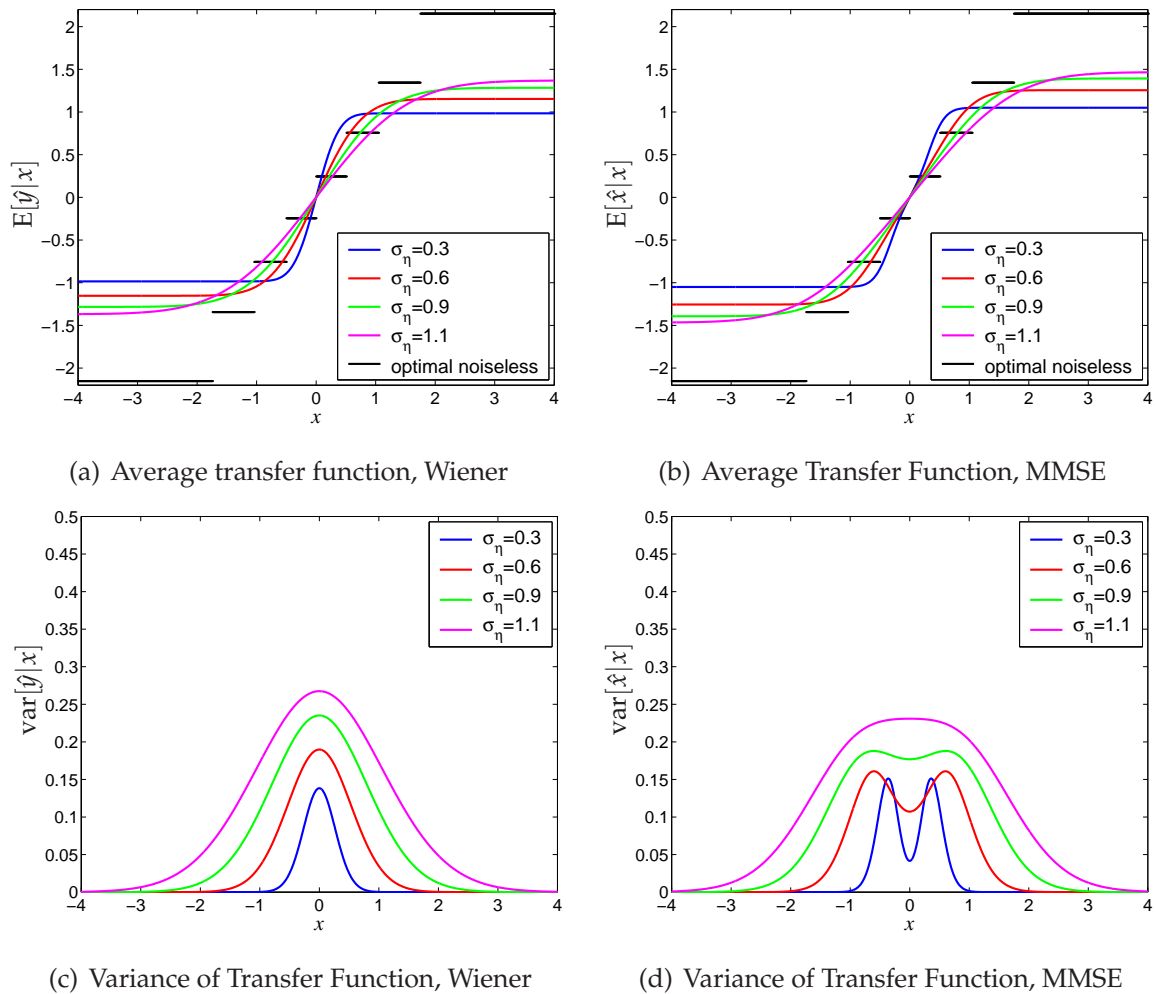


Figure 6.19. Decoding transfer functions and variance. Fig. 6.19(a) shows the *average* transfer function for Wiener linear decoding, with $N = 7$, for various values of noise standard deviation, σ_η , and Gaussian noise. Fig. 6.19(b) shows the average transfer function for MMSE nonlinear decoding. Also shown for comparison with thick black lines is the optimal deterministic transfer function for a 3-bit quantisation of a Gaussian source. Fig. 6.19(c) shows the variance corresponding to each value of σ_η as a function of x , for Wiener decoding, and Fig. 6.19(d) shows the variance for MMSE decoding. The average transfer function of each decoding scheme can be seen to be almost identical, however, the variance can be seen to be smaller for more values of x , particularly for values of x close to the signal mean.

6.7 An Estimation Perspective

This section examines the encoding and decoding of an input signal in the SSR model from the point of view of estimation theory. Before proceeding to some results, we firstly introduce the concepts of *estimation*, *bias*, *Fisher information*, the *Cramer-Rao bound* and the *information bound*, and then show how they apply to the SSR model. The main result obtained is a method for calculating a lower bound on the MSE distortion for any given decoding of the SSR model.

6.7.1 Definitions

Estimators

Suppose one can make a measurement, z , of some scalar parameter, θ . If there is some random or systematic error in the measurement, then z is a point *estimate* of the parameter, θ (Lehmann and Casella 1998). The error of any given measurement is $\epsilon = z - \theta$.

Bias of an estimator

The *bias*, $b_z(\theta)$, of an estimator, z , for a parameter, θ , is the expected value of the error, ϵ , of the estimator. Suppose z is discretely valued, and has a probability mass function for a given θ of $P_z(z|\theta)$. Then the bias is

$$\begin{aligned} b_z(\theta) &= E[\epsilon|\theta] \\ &= E[z|\theta] - \theta \\ &= \sum_z z P_z(z|\theta) - \theta. \end{aligned} \tag{6.102}$$

An *unbiased estimator* is an estimator for which the bias is zero, that is $E[z|\theta] = \theta$ (Cover and Thomas 1991).

If the parameter, θ , is also a continuous random variable, with a PDF given by $P_\theta(\theta)$, then the mean square value of the bias is

$$E[b_z^2(\theta)] = \int_\theta b_z^2(\theta) P_\theta(\theta) d\theta. \tag{6.103}$$

Fisher information

Consider a random variable Z that estimates a parameter θ . Let $P(Z = z|\theta)$ be the conditional probability density function of the random variable Z given measurement

θ . The *Fisher information* is defined as the variance of the gradient of the log-likelihood function with respect to z (Cover and Thomas 1991). The gradient of the log-likelihood function is often called the *score function*,

$$U(Z, \theta) = \frac{d \ln P(Z = z|\theta)}{d\theta}. \quad (6.104)$$

and hence the Fisher information is

$$J(\theta) = \text{var}[U(Z, \theta)] = \text{var} \left[\frac{d \ln P(Z = z|\theta)}{d\theta} \right]. \quad (6.105)$$

Thus the Fisher information is a function of the parameter, θ . From this definition, it is possible to derive an expression for the Fisher information in terms of a second derivative with respect to θ as

$$J(\theta) = \text{E} \left[-\frac{d^2 \ln P(Z = z|\theta)}{d^2\theta} \right]. \quad (6.106)$$

Cramer-Rao bound

Fisher information can be related to the variance of an estimator by the *Cramer-Rao bound* (Cover and Thomas 1991). This bound holds for an unbiased estimator but is easily generalised to biased estimators, as discussed below. The Cramer-Rao bound states that the reciprocal of the Fisher information is a lower bound on the variance of the error of an unbiased estimator and thus gives the smallest possible variance on an unbiased estimator for θ . Note however that this smallest possible bound may not be achievable. The Cramer-Rao bound for an unbiased estimator is given by

$$\text{var}[\epsilon|\theta] \geq \frac{1}{J(\theta)}. \quad (6.107)$$

Section D.7 of Appendix D gives a proof of the Cramer-Rao bound.

Efficient estimators

An unbiased estimator is said to be *efficient* if it meets the Cramer-Rao bound with equality.

Information bound for a biased estimator

A result analogous to the Cramer-Rao bound holds for a biased estimator (Cover and Thomas 1991, Lehmann and Casella 1998). A lower bound on the variance of the error

6.7 An Estimation Perspective

of a biased estimator, z , for the parameter, θ , can be expressed in terms of the expected value of the estimate given θ and the Fisher information of θ as

$$\text{var}[\epsilon|\theta] \geq \frac{\left(\frac{d}{d\theta} \mathbb{E}[z|\theta]\right)^2}{J(\theta)}. \quad (6.108)$$

This generalisation of the Cramer-Rao bound is sometimes known as the *information bound* or *information inequality* (Lehmann and Casella 1998). Section D.7 of Appendix D gives a proof of the information bound.

Noting that $\text{var}[z|\theta] = \text{var}[z|\theta] - \text{var}[\theta|\theta] = \text{var}[z - \theta|\theta] = \text{var}[\epsilon|\theta]$, Inequality (6.108) is equivalent to stating

$$\text{var}[z|\theta] \geq \frac{\left(\frac{d}{d\theta} \mathbb{E}[z|\theta]\right)^2}{J(\theta)}. \quad (6.109)$$

The information bound can be rewritten in terms of the MSE distortion for a given θ as

$$\begin{aligned} \mathbb{E}[\epsilon^2|\theta] &\geq \frac{\left(\frac{d}{d\theta} \mathbb{E}[z|\theta]\right)^2}{J(\theta)} + \mathbb{E}[\epsilon|\theta]^2 \\ &= \frac{\left(1 + \frac{db_z(\theta)}{d\theta}\right)^2}{J(\theta)} + b_z^2(\theta). \end{aligned} \quad (6.110)$$

This version of the information bound is a function of both the bias and the Fisher information. As with Eqn. (6.4), the information bound states how a lower bound for the MSE distortion is a tradeoff between the mean square bias, and—via the Fisher information term—a term related to the conditional error variance, regardless of the actual decoding. Note that setting $b_z(\theta) = 0$ reduces Inequality (6.110) to the Cramer-Rao bound, and thus the Cramer-Rao bound is a specific case of the more general information bound.

6.7.2 Application to the SSR model

We now apply the above theory to describing signal encoding and decoding in the SSR model.

Bias for the SSR model

Let the output of the SSR model, y , be an estimate for the input. Recalling that $\mathbb{E}[y|x] = NP_{1|x}$, the bias is then

$$b_y(x) = \mathbb{E}[\epsilon(x)|x] = \mathbb{E}[y|x] - x = \sum_{n=0}^N nP(n|x) - x = NP_{1|x} - x. \quad (6.111)$$

Thus, the SSR encoding, y , can be described as a biased estimator for the input signal, x .

The expected value of the bias, taken over all possible values of the parameter, x , is $E[b_y(x)] = E[y] - E[x] = N/2$. The mean square value of the bias is

$$E[b_y^2(x)] = \int_x P(x) b_y^2(x) dx = \int_x (E[y|x] - x)^2 P(x) dx. \quad (6.112)$$

Fisher information for the SSR model

Again, we consider the output of the SSR model, $y \in 0, \dots, N$, to be an estimate for the input, x . The score is then the random variable

$$U(x, y) = \frac{d}{dx} \ln(P(n|x)), \quad n = 0, \dots, N, \quad (6.113)$$

which is a function of both x and y , and is simply the gradient with respect to x of the log-likelihood function for each value of y . It is straightforward to show that the mean over all y of the score for a particular value of x is zero. Therefore the variance of the score for a particular value of x is the mean square value with respect to n . This defines the Fisher information for the SSR encoding,

$$J(x) = E[U(x, y)^2] = \sum_{n=0}^N P(n|x) U(x, n)^2 = \sum_{n=0}^N P(n|x) \left(\frac{d}{dx} \ln(P(n|x)) \right)^2. \quad (6.114)$$

It is also straightforward to show that the Fisher information can be written as

$$J(x) = E \left[- \frac{d^2 \ln(P(n|x))}{dx^2} \right]. \quad (6.115)$$

Simplification of Eqn. (6.114) gives

$$J(x) = \sum_{n=0}^N \frac{\left(\frac{dP(n|x)}{dx} \right)^2}{P(n|x)}. \quad (6.116)$$

Section D.6 of Appendix D gives two derivations for the Fisher information of the SSR model in terms of $P_{1|x}$, the result being

$$J(x) = \left(\frac{dP_{1|x}}{dx} \right)^2 \frac{N}{P_{1|x}(1 - P_{1|x})}. \quad (6.117)$$

This is in agreement with the formula derived for the Fisher information in Hoch *et al.* (2003a) and Hoch *et al.* (2003b), as discussed in Section 5.4.5 of Chapter 5.

6.7 An Estimation Perspective

Eqn. (6.117) can be simplified by recalling that for $\theta = 0$, and the noise PDF even about a mean of zero that $P_{1|x} = F_R(x)$. Hence

$$\left(\frac{dP_{1|x}}{dx}\right)^2 = R(x)^2, \quad (6.118)$$

and the Fisher information is

$$J(x) = \frac{NR(x)^2}{P_{1|x}(1 - P_{1|x})}. \quad (6.119)$$

Thus, the Fisher information increases linearly with N . The increase in information with N makes sense since, from the estimation point of view, N corresponds to the number of times a noisy measurement is taken of the same parameter x . With more measurements, more information about the parameter is obtained.

6.7.3 SSR Encoding

Since the output of the SSR model, y , is a biased estimator for the input, x , the relevant inequality giving a bound on the variance of the error of y as an estimator is the information bound rather than the Cramer-Rao bound. Recalling that $E[y|x] = NP_{1|x}$, we have from Inequality (6.108)

$$\text{var}[\epsilon|x] = \text{var}[y|x] \geq N^2 \frac{\left(\frac{dP_{1|x}}{dx}\right)^2}{J(x)} = N^2 \frac{R(x)^2}{J(x)}. \quad (6.120)$$

Substituting for the Fisher information given by Eqn. (6.119) we have

$$\begin{aligned} \text{var}[y|x] &\geq N^2 R(x)^2 \frac{P_{1|x}(1 - P_{1|x})}{NR(x)^2} \\ &= NP_{1|x}(1 - P_{1|x}) \\ &= \text{var}[y|x], \end{aligned} \quad (6.121)$$

where the last step follows by recalling the fact that $P(n|x)$ is given by Eqn. (6.79) and the variance of such a binomial distribution is $NP_{1|x}(1 - P_{1|x})$. Thus, the SSR model provides an estimate for x with a conditional MSE distortion equal to its minimum possible MSE distortion, *given its bias*. Thus, the SSR model meets the information bound with equality, and although the definition of *efficient* applies only to unbiased estimators, the SSR encoding, y , can be said to be an *efficient biased estimator*. Discussion of such terminology, and relevant theoretical background can be found in Lehmann and Casella (1998).

Despite the above result, we know that decoding the SSR model by assigning reproduction points to its output results in a smaller conditional MSE distortion for a given x than the raw conditional variance of y . This fact is accounted for once it is noted that for a decoded output, the expected value of y given x in Inequality (6.120) needs to be replaced by the expected value of the decoding, \hat{y} , given x . Assuming the decoded output is still a biased estimator, the bound on the conditional MSE distortion is

$$\begin{aligned} E[\epsilon^2|x] &\geq \frac{\left(\frac{d}{dx}E[\hat{y}|x]\right)^2}{J(x)} + E[\epsilon|x]^2 \\ &= \frac{\left(\frac{d}{dx}E[\hat{y}|x]\right)^2}{J(x)} + (E[\hat{y}|x] - x)^2. \end{aligned} \quad (6.122)$$

Note that the Fisher information is independent of the decoding. This can be seen by inspecting its definition, from which it is clear that it is only dependant on the transition probabilities, $P(n|x)$. Thus, from Inequality (6.122), the bound on the conditional MSE distortion is entirely dependent on $E[\hat{y}|x]$. The next two subsections consider the information bound for both linear and nonlinear decoding of the SSR model.

6.7.4 SSR Linear Decoding

Bias of linear decoding

Recall that a zero mean linear decoding of the SSR model output can be written as in Eqn. (6.18). The bias of this estimate is $b_{\hat{y}}(x) = E[\hat{y}|x] - x = \frac{2c}{N}E[y|x] - c - x$. Thus, recalling that $E[y|x] = NP_{1|x}$, the bias can be written as

$$b_{\hat{y}}(x) = 2cP_{1|x} - c - x. \quad (6.123)$$

This is clearly non-zero in general and hence \hat{y} is a biased estimator, unless $P_{1|x} = x/2c + 0.5$, that is $P_{1|x}$ is linear with x . This can occur for the case of uniform signal and noise, since for $\theta = 0$,

$$P_{1|x} = \begin{cases} 0 & \text{for } x < \sigma_\eta/2, \\ x/\sigma_\eta + 1/2 & \text{for } -\sigma_\eta/2 \leq x \leq \sigma_\eta/2, \\ 1 & \text{for } x > \sigma_\eta/2. \end{cases} \quad (6.124)$$

In this situation, \hat{y} is an unbiased estimator for all x when $c = \sigma_\eta/2$ and x is in the range $[-\sigma_\eta/2, \sigma_\eta/2]$, so that $\sigma_x \leq \sigma_\eta$. For other noise distributions, the bias may be

6.7 An Estimation Perspective

close to zero for many values of x . For example, for Gaussian noise $P_{1|x}$ is given by the complimentary error function, which is known to be approximately linear for values of x near its mean. In general, however, we assume that \hat{y} will be biased, since the information bound holds whether an estimator is biased or unbiased.

Information bound for linear decoding

We now show that if a linearly decoded output is taken as an estimate for x , then the estimate is an efficient biased estimator, that is, the estimate gives a MSE distortion that meets the information bound with equality.

Differentiating the bias given by Eqn. (6.123) with respect to x and adding one gives

$$1 + b'_{\hat{y}}(x) = 2c \frac{d}{dx} P_{1|x} = 2cR(x). \quad (6.125)$$

Substituting Eqn. (6.125) into Inequality (6.110) gives

$$E[\epsilon^2|x] \geq \frac{4c^2R(x)^2}{J(x)} + 4c^2P_{1|x}^2 - 4cP_{1|x}(c+x) + (c+x)^2. \quad (6.126)$$

Substituting Eqn. (6.119) into Inequality (6.126) gives

$$\begin{aligned} E[\epsilon^2|x] &\geq \frac{4c^2P_{1|x}(1-P_{1|x})}{N} + 4c^2P_{1|x}^2 - 4cP_{1|x}(c+x) + (c+x)^2 \\ &= 4c^2P_{1|x}(1-P_{1|x}) \left(\frac{1-N}{N} \right) - 4cP_{1|x}x + (c+x)^2. \end{aligned} \quad (6.127)$$

The final expression is exactly that previously derived for the conditional MSE distortion in Eqn. (6.44). Hence, a linearly decoded output of the SSR model, \hat{y} , gives an average conditional MSE distortion that is equal to the theoretical limit given by the information bound.

6.7.5 SSR Nonlinear Decoding

For nonlinear decoding no such result demonstrating biased efficiency, as in the linear case, can be found. However, the bound given by Inequality (6.122), can be used to verify numerical calculations of the MSE distortion as follows. Firstly, notice that for any given decoding, \hat{y} , Inequality (6.122) becomes

$$E[\epsilon^2|x] \geq \frac{\left(\frac{d}{dx} E[\hat{y}|x] \right)^2}{J(x)} + (E[\hat{y}|x] - x)^2. \quad (6.128)$$

Now

$$\begin{aligned} \frac{d}{dx} (\mathbb{E}[\hat{y}|x]) &= \sum_{n=0}^N \hat{y}_n \frac{d}{dx} (P(n|x)) \\ &= \frac{R(x)}{P_{1|x}(1 - P_{1|x})} \sum_{n=0}^N \hat{y}_n P(n|x) (n - NP_{1|x}). \end{aligned} \quad (6.129)$$

Therefore

$$\begin{aligned} \frac{\left(\frac{d}{dx} (\mathbb{E}[\hat{y}|x])\right)^2}{J(x)} &= \frac{\left(\sum_{n=0}^N \hat{y}_n P(n|x) (n - NP_{1|x})\right)^2}{NP_{1|x}(1 - P_{1|x})} \\ &= \frac{\left(\sum_{n=0}^N \hat{y}_n P(n|x) (n - \mathbb{E}[y|x])\right)^2}{NP_{1|x}(1 - P_{1|x})} \\ &= \frac{(\text{cov}[\hat{y}y|x])^2}{\text{var}[y|x]} \\ &= \rho_{\hat{y}y|x} \text{var}[\hat{y}|x]. \end{aligned} \quad (6.130)$$

For example, if $\hat{y} = y$ then the RHS of Eqn. (6.130) becomes simply the variance of y given x , and we have the case of the information bound being met with equality, as in Eqn. (6.121). Furthermore, for a linear decoding, the correlation coefficient between y and \hat{y} will be unity, and we have the same result as the previous subsection, that the information bound is met with equality. Note that Eqn. (6.130) implies that

$$\text{var}[\epsilon|x] \geq \rho_{\hat{y}y|x} \text{var}[\hat{y}|x]. \quad (6.131)$$

Since $\text{var}[\hat{y}|x] = \text{var}[\epsilon|x]$, this inequality now states $\rho_{\hat{y}y|x} \leq 1$, which is a demonstration of the validity of our arguments, since the correlation coefficient is always smaller than unity. Furthermore, the bound is only met with equality if $\rho_{\hat{y}y|x} = 1$.

Average Information Bound

The main point of this Section is to obtain an expression suitable for numerical calculation of a lower bound on the MSE distortion, and comparing this to the known MMSE distortion.

Taking the expected value of Inequality (6.128) with respect to the input PDF, $P(x)$, gives a bound on the MSE distortion, in terms of two terms, the first term being the expected value of Eqn. (6.130) and the second being the mean square value of the bias. Since we know that the best possible decoding is MMSE decoding, with reconstruction

6.7 An Estimation Perspective

points given by $\hat{x}_n = E[x|n]$, we will use \hat{x} as the decoding in the RHS of Inequality (6.128). Carrying this out gives

$$\begin{aligned} \text{MSE} &\geq E \left[\frac{\left(\sum_{n=0}^N \hat{x}_n P(n|x) (n - NP_{1|x}) \right)^2}{NP_{1|x}(1 - P_{1|x})} \right] + E[b_{\hat{x}}(x)^2] \\ &= \int_x P(x) \left[\frac{\left(\sum_{n=0}^N \hat{x}_n P(n|x) (n - NP_{1|x}) \right)^2}{NP_{1|x}(1 - P_{1|x})} \right] dx + \int_x P(x) b_{\hat{x}}(x)^2 dx. \end{aligned} \quad (6.132)$$

The expression on the RHS of Inequality (6.132) is a lower bound on the MSE distortion. It is a quantity smaller than, or equal to, the MMSE distortion that results from the optimal nonlinear decoding. Note that it has two components, one which we will refer to as the *average error variance* component, and the other being the mean square bias. We will call this expression the *Average Information Bound* (AIB) and denote it as

$$\text{AIB} = E \left[\frac{\left(\sum_{n=0}^N \hat{x}_n P(n|x) (n - NP_{1|x}) \right)^2}{NP_{1|x}(1 - P_{1|x})} \right] + E[b_{\hat{x}}(x)^2]. \quad (6.133)$$

The AIB is plotted from numerical calculations in Fig. 6.20 for the cases of matched Gaussian, Laplacian and logistic signal and noise. When compared with the MMSE distortion as shown in Fig. 6.11, the MMSE distortion looks to be very close to the AIB. To verify this, Fig. 6.21 shows the percentage difference between the MMSE distortion and the AIB. Although the percentage difference is quite small—always less than eight percent—it is clear that the AIB is always smaller than the MMSE distortion, as expected. Notice that for $N = 1$, the difference between the MMSE distortion and the AIB is constant for all σ and appears to be very nearly zero.

Fig. 6.21 also shows that as N increases, the peak percentage difference increases with N , however the rate of increase appears to slow for larger N . For larger σ , the percentage difference is increasing with N for small N , but is decreasing with N for large N . These observations indicate that for sufficiently large N , the MMSE distortion might converge towards the AIB for all σ .

Note that the difference between the MMSE distortion and the AIB is solely due to the difference between the average error variance term—that is, the Fisher information term in the information bound—and the actual average error variance, since the mean square bias term is identical for both the MMSE distortion and the AIB.




Figure 6.20. Average information bound. This figure shows plots of the average information bound against increasing σ for various values of N , and the cases of matched Gaussian, Laplacian and logistic signal and noise. When compared with the MMSE distortion as shown in Fig. 6.11, the MMSE distortion is very close to the AIB.




Figure 6.21. Percentage difference between MMSE distortion and average information bound. This figure shows plots of the percentage difference between the MMSE distortion, and the average information bound against increasing σ , for various values of N . As N increases, the peak percentage difference increases with N , however the rate of increase appears to slow as N gets larger. For larger σ , the percentage difference is increasing with N for small N , but is decreasing with N for large N . These observations indicate that for sufficiently large N , the MMSE might converge towards the AIB for all σ .

Tradeoff between error variance and bias

To illustrate the effects of the two different components of Eqn. (6.133), the mean square bias is plotted in Fig. 6.22, and the average error variance in Fig. 6.23. The mean square bias can be seen to decrease with increasing N for all σ . It also decreases with increasing σ for small σ , but reaches a minimum value before increasing again as σ increases. For large N , the increase in the mean square bias with increasing σ greater than the minimising σ , is very slow. By contrast, apart from very small values of N , the average error variance decreases with increasing N but increases with increasing σ .

This behaviour is somewhat analogous to the mutual information for the SSR model, as discussed in Chapter 4. There, we saw that the mutual information has a maximum for nonzero noise, but consists of two components—the output entropy, and the average conditional output entropy. We saw that the average conditional output entropy always increases with increasing σ , but the output entropy increases for small σ , reaches a maximum value at $\sigma = 1$, and then decreases again with increasing σ . Thus, the optimal value of σ for the mutual information occurs when the slope of the output entropy with respect to σ equals the slope of the average conditional output entropy, thus indicating a tradeoff between output entropy and conditional output entropy. Here we have the minimising value of σ for the AIB being given when the slope of the average error variance with respect to σ is equal to the negative of the slope of the mean square bias. This, again, illustrates that the optimal MSE distortion is obtained for σ that gives the best tradeoff between a small error variance and a small bias.

Also of interest is the fact that the mean square bias seems to get very small for large N and σ larger than unity. This might indicate that for sufficiently large N , the mean square bias is negligible with respect to the average error variance, and the AIB becomes a function of only the average error variance. We will return such large N behaviour of the MSE distortion in Chapter 7.

We now consider in detail previously proposed output SNR measures for the SSR model.




Figure 6.22. Mean square bias term. Plot of the mean square bias against increasing σ for various values of N . The mean square bias can be seen to decrease with increasing N , for all σ , and have a minimum value at some value of σ near unity. As N increases, the rate of increase of the mean square bias for σ larger than the minimising σ is very slow.




Figure 6.23. Average error variance term. Plot of the average error variance term, against increasing σ for various values of N . Apart from very small values of N , the average error variance decreases with increasing N , but increases with σ .

6.8 Output Signal-to-Noise Ratio

Previously in this chapter, we have measured the performance of the SSR stochastic quantiser by the SQNR measure. Although SQNR is commonly used in quantisation theory, we saw in Chapter 3 that other SNR measures are often used in the SR literature, since often the output is not a quantisation of the input, nor is the input signal a sequence of samples drawn from a stationary probability distribution.

However, a constant problem in using SNR to measure a nonlinear system's output performance, is the question of how to separate the output signal from output noise. In linear systems, where the signal and noise are uncorrelated, noise power can be easily separated from signal power due to both signals being amplified or attenuated by the same ratio, so that the output signal and noise power sums to the total power. Furthermore, the output noise power can be measured as the power of the output when no signal is present at the input.

As we saw in Section 3.4 of Chapter 3, in nonlinear systems the separation of noise from signal at the output of a nonlinear system is not so straightforward, as the input signal can interact with the input noise to yield extra output noise that is not present in the absence of a signal. In other words, if the signal and noise are correlated, then their powers do not necessarily sum.

One solution to this problem is to suitably define the noise component of the output power, and then let the signal component be the total output power less the noise power. Alternatively, the signal component of the output power might be defined, and the noise power is the total power less the signal power. If this approach is taken, then the signal and noise powers sum, as carried out in Section 3.4 of Chapter 3. However, if the input signal and output noise are correlated, then does such a definition of an output signal really relate well to the actual input signal?

An approach which avoids this problem is that taken with the SQNR measure, which defines the output signal power as the *input* signal power. The output noise power is considered to be the power of the error signal, that is, the variance of the error. Note that with these definitions, the output signal power and output noise power do not add up to the total output signal power, since the signal, x , and the error, ϵ , are correlated for a given value of x .

The remainder of this section discusses this question of how to separate signal and noise in a nonlinear system, and argues that previous approaches taken for the SSR

model—or equivalent systems—are inferior to the SQNR measure, as they do not take into account the correlation between output noise and input signal, and the associated fact that a small variance in the output for a given input can also result in a large mean error—that is, bias—for that input, and conversely, a small mean error may have a large variance.

6.8.1 SNR in the DIMUS model

Recall that the Digital Multibeam Steering (DIMUS) sonar array is very similar to the SSR model. One approach in the literature to measuring the SNR at the output of a DIMUS array relies on a method for separating signal and noise where the noise is considered to be the output response in the absence of a signal.

Recall that Eqn. (6.53) is identical to the expression obtained in Remley (1966) for the output power, P_T , of the DIMUS sonar array for a Gaussian signal and Gaussian noise. The output of the DIMUS array is equivalent to linearly decoding the SSR model with $c = N$. For the output noise power, Remley (1966) uses the output variance when no signal is present. For the SSR model, this is equivalent to letting $\sigma = \infty$, and therefore $\rho_i = 0$ in Eqn. (6.53), and the output noise power, $P_N = N$. The output signal power is taken to be $P_S = P_T - P_N$

With these definitions, the output SNR is

$$\text{SNR} = \frac{P_T - P_N}{P_N} = \frac{2(N-1)}{\pi} \arcsin\left(\frac{1}{\sigma^2 + 1}\right), \quad (6.134)$$

which—by realising that Remley (1966)'s input SNR variable, R_{in} , is the input SNR, $1/\sigma^2$ —is identical to Eqn. (19) for the SNR in Remley (1966).

Note that Eqn. (6.134) is strictly decreasing with increasing σ , and therefore SR does not occur in this SNR measure. Furthermore, when $\sigma = 0$, that is, in the absence of input noise, the SNR attains its maximum value of $N - 1$. This does not make a lot of sense, as just as with the SSR model, at $\sigma = 0$ the output can only have two states. For DIMUS these are $\pm N$, with equal probability, and hence the total output power is large. This SNR measure does not take into account the fact that the noise power at $\sigma = \infty$ is significantly different from that of small values of σ , due to the need to consider the output noise to be correlated with the input signal, and this correlation varying as a function of σ .

6.8 Output Signal-to-Noise Ratio

Remley (1966) however, also points out that the case of most interest for DIMUS is when N is large, and R_{in} is small, i.e. σ is large. For sufficiently large σ , we have

$$\text{SNR} \simeq \frac{2(N-1)}{\pi(\sigma^2+1)}. \quad (6.135)$$

As noted by Remley (1966), if N is also large, and $\sigma^2 \gg 1$, $\text{SNR}_0 \simeq \frac{2N}{\pi\sigma^2}$.

6.8.2 Stocks' SNR Measure for SSR

The final section of Stocks (2001a) considers SSR from an SNR point of view for the first time. The method of output signal and noise separation relies on defining the signal component of the overall output for every possible value of the input and can be described in the following manner. It does not define a decoding for the output, but simply considers the SNR of the encoded output, y .

For a given value of the input signal, x , the corresponding output is the random variable, $w = y|x$. Stocks (2001a) notes that the ensemble average of w is $E[w] = E[y|x]$, which can be considered to be the signal component of the output for the given input value, x . This conditional ensemble average has an uncertainty, or noise, that can be described by the variance of w , that is, $\text{var}[w] = \text{var}[y|x]$.

Now, for the SSR model, the conditional variance of the output is $\text{var}[y|x] = NP_{1|x}(1 - P_{1|x})$. Stocks (2001a) takes this to be the noise power for a given x , and hence the *average* noise power can be defined as the expected value of the *conditional* noise power, that is $E[\text{var}[y|x]] = \frac{N}{2} - NE[P_{1|x}^2]$. Note that since $\text{var}[y|x] = \text{var}[\epsilon|x]$, we have already come across this term in this chapter, and labelled it as the average error variance. Recall from Eqn. (6.6) that the average error variance is only one component of the MSE distortion, the other component being the mean square bias.

Thus, Stocks (2001a) defines the average error variance as being the output noise power, N_p , and the total output power as the variance of the output, y , which is $\text{var}[y] = N(N-1)E[P_{1|x}^2] - \frac{N(N-2)}{4}$. The difference between this total output power, and the noise power is considered to be the output *signal* power, S_p .

Thus, with these definitions, the average output noise power for the undecoded SSR model is

$$N_p = E[\text{var}[y|x]] = \frac{N}{2} - NE[P_{1|x}^2], \quad (6.136)$$

and the average output signal power is

$$S_p = \text{var}[y] - \text{E}[\text{var}[y|x]] = N^2 \left(\text{E}[P_{1|x}^2] - \frac{1}{4} \right) = \frac{N^2 \rho_o}{4}, \quad (6.137)$$

where ρ_o is the correlation coefficient between the output of any two threshold devices, as given by Eqn. (D.20) in Appendix D.

Hence, the output SNR is

$$\text{SNR}_o = \frac{S_p}{N_p} = \frac{N^2 \left(\text{E}[P_{1|x}^2] - \frac{1}{4} \right)}{\frac{N}{2} - N \text{E}[P_{1|x}^2]} = \frac{N(4\text{E}[P_{1|x}^2] - 1)}{2 - 4\text{E}[P_{1|x}^2]} = \frac{0.5N\rho_o}{1 - 2\text{E}[P_{1|x}^2]}. \quad (6.138)$$

As shown in Stocks (2001a), unlike for the mutual information—or, as seen in this Chapter, the SQNR—the SNR measure of Eqn. (6.138) monotonically decreases with σ and hence no SR effect is seen in this measure. The signal power, S_p , also decreases monotonically with increasing σ if $\theta = 0$, since, as we saw earlier, ρ_o decreases with increasing σ . Stocks (2001a) notes that the reason for this is that, in the noiseless case, the output is deterministic, and hence has zero noise power with these definitions. This implies an infinite output SNR for zero noise. For non-zero noise that is still small, the output is still nearly always in states 0 and N , and therefore has little variance for a given x , and thus very little noise power, and a very high output SNR. As the input noise increases, the average output noise power also increases and therefore the output SNR decreases.

It is clear that this definition of SNR is deficient for this situation. Although the ideal situation is for the output to have no variance given the input, the ideal situation is also for the bias to be zero, that is, $\text{E}[w] = \text{E}[y|x] = x$. This is clearly not the case when the average conditional variance is minimised, which occurs at $\sigma = 0$. Defining the signal power in this manner ignores the fact that the ideal estimator of the input is one that finds the best tradeoff between bias and variance as illustrated by Eqns. (6.4) and (6.110).

Another problem with this approach is the assumption that the average output noise power and the average output signal power should sum to the total output power. The reason that the output signal and noise power does not sum to the total power is that the input signal and the output noise are correlated.

These problems have also been recognised by Martorell *et al.* (2005), who attempt to overcome them by modifying the noise power definition of Stocks (2001a) to take into

6.8 Output Signal-to-Noise Ratio

account the error between the expected value of the output and some “desired function,” as well as the variance. An example given by Martorell *et al.* (2005) for a “desired function,” $q(x)$, is $q(x) = x$, so that the error considered is $E[y|x] - x$, which we recognise as being the *bias* for a given x . Thus, if the output is decoded so that it has a mean of zero, the modified noise power formula given by Martorell *et al.* (2005) is exactly the MSE distortion discussed in this chapter. However, the signal power of Martorell *et al.* (2005) is not modified from that used in Stocks (2001a), and the resultant SNR measure is neither that of Stocks (2001a), nor the SQNR. The SNR measure of Martorell *et al.* (2005) could perhaps be modified to allow the output signal power to be the same as the input signal power, and therefore be equivalent to the SQNR in the case where the desired function is x . Therefore, if the desired function is not x —for example if a nonzero bias is considered ideal—the SNR measure of Martorell *et al.* (2005) would then be a generalisation of the SQNR.

One final note, is that for Gaussian signal and noise we have $E[P_{1|x}^2]$ given by Eqn. (6.50) and therefore the SNR is

$$\text{SNR}_o = \frac{N \arcsin(\rho_i)}{\left(\frac{\pi}{2} - \arcsin(\rho_i)\right)} = \frac{N \arcsin(\rho_i)}{\arccos(\rho_i)}. \quad (6.139)$$

For sufficiently large σ , the SNR measure can be approximated by

$$\text{SNR}_o \simeq \frac{2N}{\pi(\sigma^2 + 1) - 2}. \quad (6.140)$$

Thus, Eqn. (6.140) is very similar to the DIMUS SNR formula for large σ , given by Eqn. (6.135). In fact, for large N and $\sigma \gg 1$, they are identical. This fact indicates that both the DIMUS SNR formula, and the formula from Stocks (2001a) may be valid for sufficiently large σ and N . We conjecture that under these conditions, the output noise becomes almost uncorrelated with the input signal, and therefore the total output power is the sum of the average output noise power, and the average output signal power.

Furthermore, recall from Eqn. (6.16) that if a continuously valued noisy signal is averaged over N *iid* realisations, where the noise has variance σ_η^2 , then the SNR is N/σ_η^2 . Thus, the scaling of this SNR with N and noise intensity is the same as the scaling in the DIMUS SNR of Eqn. (6.135) and Stock’s SNR of Eqn. (6.140). This indicates that for large N and σ , the SSR model’s output behaves like a continuously valued signal.

6.8.3 Gammaitoni's Dithering Formula

As briefly mentioned in Chapter 2, a similar situation occurs in Gammaitoni (1995a), which proposes a measure of comparing the average output of a quantiser with a linear response. In Gammaitoni (1995a), the ideal output response is considered to be a linear transformation of the input signal. Hence, a measure of the quantisers performance is given by

$$D_G = \sqrt{\int_0^1 (\mathbb{E}[y|x] - x)^2 dx}, \quad (6.141)$$

where a smaller D_G means a better performance.

Eqn. (6.141) assumes all values of the input signal should be uniformly weighted. Generalising Eqn. (6.141) to the case considered throughout this chapter of signals being a random sequence of samples drawn from the distribution with PDF, $P(x)$, gives

$$D_G = \sqrt{\int_x (\mathbb{E}[y|x] - x)^2 P(x) dx}, \quad (6.142)$$

which, from Eqn. (6.112) means that D_G^2 is the mean square bias, $\mathbb{E}[b_y(x)^2]$.

Thus, the formula given in Gammaitoni (1995a) measures only the mean square bias in a quantiser, and ignores the average conditional error variance. This is the opposite to the SNR measure proposed in Stocks (2001a), which ignores the bias, and measures only the average conditional error variance. The MSE distortion—or equivalently, the SQNR measure—used by engineers to measure a quantiser's performance, which takes the variance *and* bias into account, gives a more complete picture of a quantiser's performance.

6.9 Chapter Summary

The initial section of this chapter, Section 6.1, introduces the concept of decoding the SSR model, first introduced in Chapter 4. It defines an error between the input signal and the decoded output, as well as introducing appropriate measures of the average performance of that error, including the mean square error distortion, signal to quantisation noise ratio, and correlation coefficient.

Section 6.2 shows how the signal-to-noise ratio of a signal can be reduced by averaging N independently noisy realisations of that signal.

Sections 6.3 through 6.6 then examine linear and nonlinear decoding schemes. Linear decoding schemes are ones for which the spacing between reconstruction points is constant for all points. Nonlinear ones are not necessarily constantly spaced. Standard theory is introduced to show that both optimal—in the sense of minimum mean square error distortion—linear and optimal nonlinear decoding schemes exist, and can be calculated for the SSR model. Such calculations always show that the mean square error distortion is minimised for some nonzero value of σ for the SSR model. Comparisons are made between each decoding, which verify that the optimal decoding is the nonlinear MMSE decoding.

Section 6.7 considers the SSR decoding from the point of view of estimation theory, and using the biased-estimate generalisation of the Cramer-Rao bound, finds an expression for a lower bound on the achievable mean square error distortion.

Finally, Section 6.8 discusses an SNR measure defined in Stocks (2001a), as well as a measure of a quantiser's performance defined in Gammaitoni (1995a), and argues that both these measures do not take into account both of the error variance, and the bias, which are both necessary components of the standard mean square error distortion measure.

6.9.1 Original Contributions for Chapter 6

This chapter included the following original contributions:

- The application of the concept of decoding to the SSR model, in order to provide a signal that approximates, or reconstructs, the input signal. Such a reconstruction can be measured with the mean square error distortion—as conventionally is the case for quantisers—as well as the correlation coefficient, and signal to quantisation noise ratio.
- Analytical derivations of expressions for the mean square error distortion, signal to quantisation noise ratio, and correlation coefficient for linear decoding schemes of the SSR model, in terms of the function $P_{1|x}$.
- Analytical derivations of expressions for the correlation coefficient for a linear decoding of the SSR model in terms of N and σ for the specific cases of Gaussian signal and noise, uniform signal and noise, and Laplacian signal and noise.

- Plots for the MSE distortion in the case of matched-moments and Wiener linear decoding schemes, and MAP and MMSE nonlinear decoding schemes, and the associated reconstruction points, and verification that MMSE decoding is the optimal decoding.
- Analytical derivation of an expression for the MMSE distortion and optimal reconstruction points for the SSR model for the case of uniform signal and noise and noise intensity, $\sigma \leq 1$.
- The application of the Cramer-Rao and information bounds to find a new lower bound, named here as the *average information bound*, on the MSE distortion for decoding of the SSR model. It is shown that a linear decoding for the SSR model is a biased efficient estimator, and that MMSE decoding is biased but does not meet the average information bound with equality. Analysis of the average information bound confirms that the value of noise intensity, σ , which minimises the MSE distortion means finding the best tradeoff between average error variance and mean square bias. This tradeoff is analogous to the tradeoff between output entropy and average conditional output entropy required to maximise the mutual information.
- Arguments that the SNR measure developed for the DIMUS sonar array in Remley (1966) is not valid for small σ , as it ignores the fact that the output noise will be correlated with the input signal, and varies with σ .
- Arguments are given that (i) the SNR measure used in Stocks (2001a) decreases monotonically with σ , because the definition of noise power used does not take bias into account, and therefore is incomplete; (ii) the measure used to describe quantisation given in Gammaitoni (1995a) is also incomplete, because it only takes into account mean square bias, and not error variance. The MSE distortion and SQNR, as used by engineers in quantisation theory, do take both of these factors into account.

6.9.2 Further Work

Possible future work and open questions arising from this chapter might include:

- Application of mean square error distortion decoding to one sided signal and noise distributions, such as the Rayleigh, or exponential distributions.

6.9 Chapter Summary

- Application of mean square error distortion decoding to mixed signal and noise distributions, for example, a uniform signal subject to Gaussian noise.
- Consideration of distortion measures other than the mean square error distortion, such as the absolute error distortion.

This concludes Chapter 6, which studies the array of threshold devices in which SSR occurs as a quantisation model, and discusses various output decoding schemes.

Chapter 7 now examines decoding schemes and the MSE distortion for the SSR model under the assumption of a large number of threshold devices.

Chapter 7

Suprathreshold Stochastic Resonance: Large N Decoding

THE aim of this Chapter is to find asymptotic large N approximations to the mean square error distortion for the suprathreshold stochastic resonance model. In particular, we are interested in how the distortion varies with noise intensity and how it scales with the number of threshold devices. That is, does the distortion become asymptotically small for large N ?

7.1 Introduction

Chapter 6 develops the idea of treating the SSR model as a lossy source coding or quantisation model. We saw that such a treatment requires specification of reproduction points corresponding to each of the $N + 1$ discrete output states. Once specified, an approximate reconstruction of the input signal can be made from a decoding, and the average error between this approximation and the original signal subsequently measured by the Mean Square Error (MSE) distortion. We also saw in Chapter 5 that asymptotic approximations to the output probability mass function, $P_y(n)$, output entropy, average conditional output entropy, and mutual information can be found if the number of threshold devices, N , in the SSR model is allowed to become very large. The aim of this Chapter is to again allow N to become very large, and develop asymptotic approximations to the MSE distortion for the cases of optimal linear and optimal nonlinear decodings of the SSR model.

7.1.1 Chapter Structure

This Chapter has three main Sections. We begin in Section 7.2 by letting N become large in the formulas derived in Chapter 6, and analysing the result. Next, Section 7.3 takes the same approach from the estimation theory perspective. Finally, we give a brief discussion in Section 7.4 on the concept of *stochastic resonance without tuning*, based on the results obtained in this Chapter. Here, *tuning* refers to ‘tuning to the optimal noise intensity’, that is, finding the optimal noise level. The term is used in an analogy with frequency resonance, where tuning to the right frequency gives the resonant response.

7.2 Mean Square Error Distortion for Large N

Recall that Chapter 6 considers two different classes of decoding, and their application to the SSR model, namely linear decoding, and nonlinear decoding. Linear decoding refers to the situation where all reproduction points are linearly spaced, so that the decoding can be written in the form $\hat{y} = ay + b$, where a and b are constants for given signal and noise distributions, but may vary with the ratio of noise standard deviation to signal standard deviation, σ . Nonlinear decoding refers to the situation where the decoded output of the SSR model varies nonlinearly for each possible value of y . We

saw that the optimal linear decoding is known as the Wiener decoding, and the optimal nonlinear decoding, which we refer to as Minimum Mean Square Error (MMSE) distortion decoding, has reproduction points given by $\hat{x}_n = E[x|y = n]$, $n = 0, \dots, N$. This section examines the large N asymptotic behaviour of both Wiener and MMSE distortion decoding for the SSR model. As in previous Chapters, we assume the signal PDF, $P(x)$, and the noise PDF, $R(\eta)$, are even functions with zero means and that the threshold values are all $\theta = 0$.

7.2.1 Wiener Linear Decoding

For zero-mean input signals and a linear decoding of the form $\hat{y} = ay + b$, we impose the condition that $E[\hat{y}] = 0$. Then $b = -aE[y]$, and as discussed in Chapter 6, the optimal value of a is given by $a = \frac{E[xy]}{\text{var}[y]}$. Since for SSR we have $E[y] = N/2$, this means that the optimal linear reconstruction points are

$$\hat{y}(n) = \frac{E[xy]}{\text{var}[y]} \left(y - \frac{N}{2} \right), \quad (7.1)$$

and the Wiener MSE distortion can be expressed in terms of the linear correlation coefficient, ρ_{xy} , as

$$\text{MSE} = E[x^2](1 - \rho_{xy}^2). \quad (7.2)$$

Recall also that the correlation coefficient is identical for all possible linear decodings. In addition, for SSR, when the signal and noise PDFs are both even valued functions about a mean of zero, the correlation coefficient can be expressed in terms of the probability that any given threshold device is 'on', $P_{1|x}$, as

$$\rho_{x\hat{y}} = \rho_{xy} = \frac{\sqrt{NE}[xP_{1|x}]}{\sqrt{E[x^2] \left((N-1)E[P_{1|x}^2] - \frac{(N-2)}{4} \right)}}. \quad (7.3)$$

Letting N approach infinity in Eqn. (7.3) gives the large N correlation coefficient as

$$\begin{aligned} \lim_{N \rightarrow \infty} \rho_{xy} &= \frac{E[xP_{1|x}]}{\sqrt{E[x^2] \left(E[P_{1|x}^2] - \frac{1}{4} \right)}} \\ &= \frac{2E[xP_{1|x}]}{\sqrt{E[x^2]\rho_o}}, \end{aligned} \quad (7.4)$$

7.2 Mean Square Error Distortion for Large N

where ρ_o is the correlation coefficient between the outputs of any two threshold devices in the SSR model, as given by Eqn. (D.20) in Appendix D. By substitution of Eqn. (7.4) into Eqn. (7.2), the large N Wiener linear decoding MSE distortion is

$$\lim_{N \rightarrow \infty} \text{MSE} = E[x^2] - \frac{4E[xP_{1|x}]^2}{\rho_o}. \quad (7.5)$$

Note that the large N limits given by Eqns. (7.4) and (7.5) are independent of N . Hence, as N increases, the correlation coefficient and MSE distortion get closer and closer to these expressions. This means that the MSE distortion does not necessarily approach zero for large N .

We now present exact results for the large N correlation coefficient and MSE distortion for specific matched signal and noise distributions. These expressions are found by letting N approach infinity in the corresponding equations stated in Chapter 6.

Gaussian signal and noise

For Gaussian signal and noise and large N , the linear decoding correlation coefficient is, from Eqn. (6.59),

$$\begin{aligned} \lim_{N \rightarrow \infty} \rho_{xy} &= \frac{1}{\sqrt{(\sigma^2 + 1) \arcsin\left(\frac{1}{\sigma^2 + 1}\right)}} \\ &= \sqrt{\frac{\rho_i}{\arcsin(\rho_i)}}, \end{aligned} \quad (7.6)$$

where $\sigma = \sigma_y/\sigma_x$ is the ratio of noise standard deviation to signal standard deviation, and ρ_i is the correlation coefficient between the inputs to any two devices in the SSR model, as defined in Eqn. (D.19) in Appendix 6. Using Eqns. (7.2) and (7.6) we also have the Wiener decoding MSE distortion as

$$\lim_{N \rightarrow \infty} \text{MSE} = E[x^2] \left(1 - \frac{\rho_i}{\arcsin(\rho_i)}\right). \quad (7.7)$$

For large σ we have $\rho_i \rightarrow \sin(\rho_i)$, and therefore the correlation coefficient is strictly increasing with σ for large N , and asymptotically approaches unity for large σ , while the MSE distortion is strictly decreasing with σ , and asymptotically approaches zero for large σ .

Uniform signal and noise

For uniform signal and noise and large N , the linear decoding correlation coefficient is, from Eqn. (6.64),

$$\lim_{N \rightarrow \infty} \rho_{x,y} = \begin{cases} \frac{3-\sigma^2}{2\sqrt{3-2\sigma}} & (\sigma \leq 1), \\ 1 & (\sigma \geq 1), \end{cases} \quad (7.8)$$

and the Wiener decoding MSE distortion is therefore

$$\lim_{N \rightarrow \infty} \text{MSE} = \begin{cases} E[x^2] \left(\frac{(\sigma-1)^3(\sigma+3)}{8\sigma-12} \right) & (\sigma \leq 1), \\ 0 & (\sigma \geq 1). \end{cases} \quad (7.9)$$

Thus, for uniform signal and noise, and infinite N , the correlation coefficient is exactly unity and the MSE distortion is exactly zero for $\sigma \geq 1$.

Laplacian signal and noise

For Laplacian signal and noise and large N , the linear decoding correlation coefficient is, from Eqn. (6.68),

$$\lim_{N \rightarrow \infty} \rho_{xy} = \frac{(2\sigma + 1)\sqrt{(\sigma + 1)(\sigma + 2)}}{2(\sigma + 1)^2}, \quad (7.10)$$

and the Wiener decoding MSE distortion is therefore

$$\lim_{N \rightarrow \infty} \text{MSE} = E[x^2] \left(\frac{3\sigma + 2}{4(\sigma + 1)^3} \right). \quad (7.11)$$

For large σ , the linear decoding correlation coefficient asymptotically approaches unity, and the Wiener MSE distortion asymptotically approaches zero.

Logistic signal and noise

As discussed in Chapter 6, we do not have an exact expression for the correlation coefficient or Wiener MSE distortion for logistic signal and noise. However, since the quantities $E[xP_{1|x}]$ and ρ_o can be calculated numerically, the large N correlation coefficient and Wiener MSE distortion can be found numerically using Eqns. (7.4) and (7.5).

We now graphically illustrate the behaviour of the above formulas and briefly discuss their main features.

Results

Fig. 7.1 shows the large N Wiener decoding MSE distortion and linear correlation coefficient for the four different matched signal and noise situations considered. It is clear that in all cases the MSE distortion is strictly decreasing with increasing σ and the correlation coefficient is strictly increasing towards unity with increasing σ . This behaviour is discussed further in Section 7.4.

We now consider the case of optimal MMSE distortion decoding.

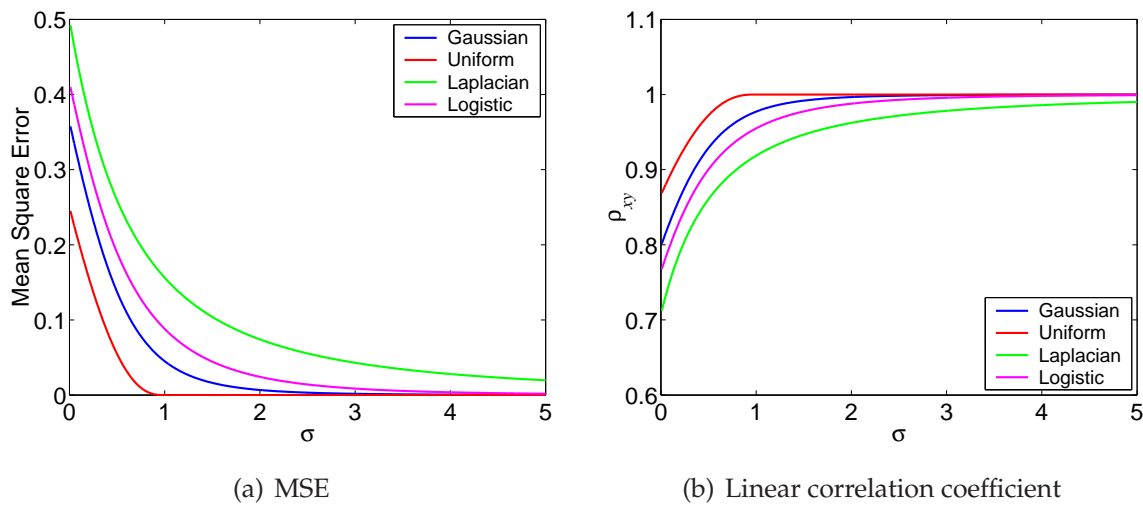


Figure 7.1. MSE distortion and correlation coefficient for Wiener decoding and large N .

Fig. 7.1(a) shows the large N Wiener decoding MSE distortion and Fig. 7.1(b) shows the large N linear correlation coefficient, for the cases of matched Gaussian, uniform, Laplacian and logistic signal and noise. For small noise intensity, σ , the MSE distortion is clearly nonzero. However, as σ increases, the MSE distortion decreases asymptotically towards zero. Likewise, the correlation coefficient approaches unity for increasing σ .

7.2.2 MMSE Optimal Decoding

As discussed in Chapter 6, and as shown in Section D.5 of Appendix D, the reproduction points that give the minimum possible MSE distortion are

$$\hat{x}_n = E[x|n] = \int_x xP(x|n)dx = \frac{1}{P_y(n)} \int_x xP(n|x)P(x)dx, \quad n = 0, \dots, N. \quad (7.12)$$

The MSE distortion that results from this decoding is

$$\begin{aligned}
 \text{MMSE} &= E[x^2] - E[\hat{x}_n^2] \\
 &= E[x^2] - \sum_{n=0}^N E[x|n]^2 P_y(n) \\
 &= E[x^2] - \sum_{n=0}^N \frac{1}{P_y(n)} \left(\int_x x P(n|x) P(x) dx \right)^2.
 \end{aligned} \tag{7.13}$$

Convergence to maximum likelihood decoding

Recall from Chapter 5 that for matched signal and noise distributions, and PDFs that are even functions about a mean of zero, the following large N approximation holds,

$$P_y(n) \simeq \frac{Q(\frac{n}{N})}{N}, \quad n = 0, \dots, N. \tag{7.14}$$

As derived in Chapter 4, $Q(\tau)$ is a PDF defined with support $\tau \in [0, 1]$, and for $\theta = 0$ and $P(x)$ and $R(\eta)$ even, and the support of P contained in the support of R , is given by

$$Q(\tau) = \frac{P(x)}{R(x)} \Big|_{x=F_R^{-1}(\tau)}, \tag{7.15}$$

where $F_R^{-1}(\cdot)$ is the Inverse Cumulative Distribution Function (ICDF) of the noise. For finite N , Eqn. (7.14) becomes an approximation, and is most accurate near $\sigma = 1$, and exact at $\sigma = 1$. We also saw in Chapter 5 that

$$P(\bar{y}) = \frac{Q(\frac{\bar{y}}{N})}{N}, \quad \bar{y} \in [0, 1], \tag{7.16}$$

is the exact PDF of the continuously valued random variable, $\bar{y}(x)$, describing the expected value of the SSR encoding, y , given the input, x .

We also used in Chapter 5 a large N approximation to the transition probabilities, which from Eqns. (5.53) and (5.54), can be written as

$$P(n|x) \simeq \frac{\delta\left(P_{1|x} - \left(\frac{n}{N}\right)\right)}{N}. \tag{7.17}$$

Substituting Eqns. (7.17) and (7.14) into Eqn. (7.12) gives

$$\hat{x}_n \simeq \frac{1}{Q(\frac{n}{N})} \int_x x \delta\left(P_{1|x} - \left(\frac{n}{N}\right)\right) P(x) dx, \quad n = 0, \dots, N. \tag{7.18}$$

7.2 Mean Square Error Distortion for Large N

Making a change of variable to $\tau = P_{1|x} = F_R(x)$ in Eqn. (7.18) gives a large N approximation of \hat{x}_n as

$$\begin{aligned}\hat{x}_n &\simeq \frac{1}{Q\left(\frac{n}{N}\right)} \int_{\tau=0}^{\tau=1} F_R^{-1}(\tau) \delta\left(\tau - \left(\frac{n}{N}\right)\right) Q(\tau) d\tau \\ &= \frac{1}{Q\left(\frac{n}{N}\right)} F_R^{-1}\left(\frac{n}{N}\right) Q\left(\frac{n}{N}\right) \\ &= F_R^{-1}\left(\frac{n}{N}\right) \quad n = 0, \dots, N.\end{aligned}\tag{7.19}$$

Recall from Chapter 6 that this is exactly the maximum likelihood decoding, and that therefore the MMSE decoding converges with large N to the maximum likelihood decoding. This result is well known in estimation theory; the maximum likelihood estimate achieves minimum variance as the number of samples approaches infinity (Lehmann and Casella 1998). Thus, the main result here is that the MMSE distortion reconstruction points converge to the maximum likelihood points, for large N .

However, unlike the optimal linear Wiener reconstruction points, which depend on σ_x , the maximum likelihood decoding reconstruction points are completely independent of the signal PDF. This is an indication that the result above may not always be accurate. However, provided that $R(\eta)$ has infinite support, we might expect that it is only the reconstruction points corresponding to n near 0 and N that are significantly different from the maximum likelihood points, and that should depend on the variance of the signal distribution. We now find evidence for this assertion.

Large N MMSE

Substituting the maximum likelihood decoding reconstruction points into Eqn. (7.13) gives the MMSE distortion for large N as

$$\text{MMSE} \simeq E[x^2] - \sum_{n=0}^N P_y(n) F_R^{-1}\left(\frac{n}{N}\right)^2.\tag{7.20}$$

Substituting Eqn. (7.14) into Eqn. (7.20) gives

$$\text{MMSE} \simeq E[x^2] - \frac{1}{N} \sum_{n=0}^N Q\left(\frac{n}{N}\right) F_R^{-1}\left(\frac{n}{N}\right)^2.\tag{7.21}$$

Suppose we replace the summation in Eqn. (7.21) by an integral as

$$\text{MMSE} \simeq E[x^2] - \frac{1}{N} \int_{n=0}^N Q\left(\frac{n}{N}\right) F_R^{-1}\left(\frac{n}{N}\right)^2 dn.\tag{7.22}$$

With a change of variable from n to $s = \frac{n}{N}$, Eqn. (7.22) becomes a limiting expression,

$$\lim_{N \rightarrow \infty} \text{MMSE} \simeq E[x^2] - \int_{s=0}^{s=1} Q(s) F_R^{-1}(s)^2 ds. \quad (7.23)$$

With a further change of variable from s to $x = F_R^{-1}(s)$, we have

$$\begin{aligned} \lim_{N \rightarrow \infty} \text{MMSE} &\simeq E[x^2] - \int_{x=-\infty}^{x=\infty} P(x) x^2 dx \\ &= 0. \end{aligned} \quad (7.24)$$

Thus, provided the summation in Eqn. (7.21) can be approximated by the integral, the MMSE approaches zero for large N . Note that we would expect the integral to more closely approximate the summation as N increases.

As stated in Section C.2 of Appendix C, a more accurate approximation of a summation by an integral is given by the Euler-Maclaurin summation formula (Spiegel and Liu 1999). This formula has remainder terms, which are significant if the 0-th and N -th terms are significant. For SSR, since $P_y(0)$ and $P_y(N)$ are significant for $\sigma \leq 0$, it is difficult to justify the approximation given in Eqn. (7.22). Hence, we do not expect Eqn. (7.24) to hold for $\sigma \leq 1$. This confirms our previous discussion regarding the maximum likelihood reconstruction points not being dependent on the signal distribution.

In order to further investigate this, we now examine letting N become large in our known exact results for MMSE decoding for uniform signal and noise.

Exact result for uniform signal and noise

Recall that Chapter 6 presents exact analytical expressions for the MMSE distortion decoding reproduction points, and the MMSE distortion for uniform signal and noise and $\sigma \leq 1$. From Eqn. (6.96), the reproduction points for large N are

$$\hat{x}(n) \simeq \frac{\sigma_\eta}{2} \left(2 \left(\frac{n}{N} \right) - 1 \right), \quad n = 1, \dots, N, \quad (7.25)$$

and from Eqn. (6.97), $\lim_{N \rightarrow \infty} \hat{x}(N) = -\hat{x}(0) = \frac{\sigma_x}{4} (1 + \sigma)$. From Eqn. (6.100), the MMSE distortion reduces to

$$\lim_{N \rightarrow \infty} \text{MMSE} = \frac{\sigma_x^2}{48} (1 - \sigma)^3. \quad (7.26)$$

Thus, for large N , the MMSE distortion is strictly decreasing for $\sigma \leq 1$, and is zero at $\sigma = 1$. We saw for a linear decoding that the MSE distortion for large N is exactly

7.3 Large N Estimation Perspective

zero for $\sigma \geq 1$ for infinite N , since it scales with $1/N$. The result of Eqn. (7.26) therefore means that no decoding scheme can make the MSE distortion asymptotically approach zero for $\sigma < 1$ for uniform signal and noise.

From Table 4.1 in Chapter 4, the ICDF of the noise is given by $F_R^{-1}(w) = \frac{\sigma_n}{2}(2w - 1)$. Therefore, for $n = 1, \dots, N - 1$ the MMSE reproduction points, for large N , given by Eqn. (7.25) are exactly the maximum likelihood reproduction points given by $F_R^{-1}(\frac{n}{N})$. However, for $n = 0$ and $n = N$, the MMSE and maximum likelihood reproduction points do not coincide, apart from at $\sigma = 1$. This fact illustrates why the MMSE decoding is said to asymptotically converge to the maximum likelihood decoding, since it only converges in the event of an infinite number of observations. In such a situation, the 0-th and N -th reproduction values become insignificant, and can be ignored in calculations of the MSE distortion, regardless of the probability of the 0-th and N -th state occurring, provided $\sigma \neq 0$.

We now consider the large N distortion performance of the SSR model from the point of view of estimation theory.

7.3 Large N Estimation Perspective

Recall that Chapter 6 discusses decoding of the SSR model from the point of view of estimation theory. This discussion includes the introduction of the definitions of *bias* and *Fisher information*, and their application to the concept of the *Cramer-Rao bound*, which gives a lower bound on the MSE distortion for an unbiased estimator, and the *information bound*, which gives a lower bound on the MSE distortion of biased estimators. Furthermore, Chapter 6 uses such estimation theory to show that the MSE distortion performance of the SSR model depends strongly on the tradeoff between bias and error variance. This section extends this discussion to the case of a large number of threshold devices in the SSR model, which is equivalent to the case of a large number of independently noisy observations of a parameter in estimation theory. We begin by introducing some estimation terminology that applies in such large N situations, before illustrating the asymptotic behaviour of the SSR model from this perspective.

7.3.1 Asymptotic Large N Theory

Suppose one can make a measurement, z , of some parameter, θ . If there is some random or systematic error in the measurement, then the error of any given measurement is $\epsilon(z) = z - \theta$.

Asymptotically unbiased estimators

An estimator is called *asymptotically unbiased* if in the limit of an infinite number of observations, the estimate becomes unbiased, that is

$$\lim_{N \rightarrow \infty} E[\epsilon(z)|\theta] = 0. \quad (7.27)$$

Consistent estimators

An estimator is called *consistent* if its MSE distortion approaches zero in the limit of an infinite number of observations, that is

$$\lim_{N \rightarrow \infty} E[\epsilon^2(z)|\theta] = D(\theta) = 0, \quad (7.28)$$

where $D(\cdot)$ is our notation for the conditional MSE distortion.

Mean asymptotic square error

Suppose for large N , an estimator is asymptotically unbiased. Then since the bias is zero, the mean square bias is also zero, and the derivative of $E[z|\theta]$ is unity. Thus, from Eqn. (6.110) in Chapter 6, the MSE distortion for a given θ is lower bounded by the Cramer-Rao bound, rather than the information bound, so that

$$D(\theta) = \text{var}[\epsilon(z)|\theta] \geq \frac{1}{J(\theta)}. \quad (7.29)$$

When averaged over all possible values of θ , the result is a lower bound on the MSE distortion as

$$\text{MSE} \geq E \left[\frac{1}{J(\theta)} \right] = \int_{\theta} P(\theta) \frac{1}{J(\theta)} d\theta. \quad (7.30)$$

Asymptotic efficiency

If the Cramer-Rao bound is met with equality, then the estimator, z , is said to be *efficient*. If efficiency is met asymptotically with increasing N , then the estimator is said to

7.3 Large N Estimation Perspective

be *asymptotically efficient*. When this occurs for all θ , the lower bound on the MSE distortion given by Inequality (7.30) becomes an equality, which we call the Mean Asymptotic Square Error (MASE), that is,

$$\text{MASE} = \mathbb{E} \left[\frac{1}{J(\theta)} \right] = \int_{\theta} P(\theta) \frac{1}{J(\theta)} d\theta, \quad (7.31)$$

which is simply the expected value of the reciprocal of the Fisher information. Note that such a quantity is discussed by Bethge *et al.* (2002) and Bethge (2003), in the context of neural population coding.

For situations where the asymptotic behaviour of an estimator's bias or efficiency is unknown, Eqn. (7.31) can still be useful, as it provides a lower bound on the MSE distortion when only the Fisher information is known, without the need for specifying a decoding.

7.3.2 Application to SSR

When using the SSR output encoding, y , as an estimate for the input, x , it does not make sense to consider asymptotic unbiasedness or consistency, as $\mathbb{E}[y|x]$ ranges between 0 and N , whereas we assume that the signal mean is zero. However, once the output is decoded, it is possible that asymptotic unbiasedness and consistency can occur. The purpose of this Section is to examine this possibility.

The Fisher information for SSR is given in Chapter 5 by Eqn. (5.69). Assuming that for large N there exists an asymptotically unbiased and efficient estimator, the MASE is then

$$\text{MASE} = \frac{1}{N} \int_x P(x) \frac{P_{1|x}(1 - P_{1|x})}{R^2(x)} dx. \quad (7.32)$$

If the integral in Eqn. (7.32) does not diverge, then the MASE exists and decreases with $\frac{1}{N}$, and we have a consistent estimator for all x . However, it is possible that the integral diverges. For example, in the case of Logistic signal and noise, we have $R(x) = P_{1|x}(1 - P_{1|x})/b_{\eta}$ and therefore the MASE can be written as

$$\text{MASE} = \frac{b_{\eta}}{N} \int_x \frac{P(x)}{R(x)} dx, \quad (7.33)$$

which diverges for $\sigma \leq 1$.

In the event that the MASE diverges, then it is likely that an asymptotically unbiased and efficient estimator does not exist. The remainder of this section examines firstly

the question of whether an asymptotically unbiased estimator exists for the SSR model, and secondly the question of efficiency.

Linear decoding

We have already seen that for a linear decoding, as N approaches infinity and σ approaches infinity, the MSE distortion approaches zero. Therefore, under these conditions, the Wiener linear decoding is a consistent estimator.

The large N Wiener decoding reconstruction points are given by

$$\hat{y}(n) \simeq c \left(2 \left(\frac{n}{N} \right) - 1 \right), \quad (7.34)$$

where $c = 2E[xP_{1|x}]/\rho_o$.

We have also, by taking the conditional expectation of Eqn. (7.1), that

$$E[\hat{y}|x] = c \left(2P_{1|x} - 1 \right). \quad (7.35)$$

Therefore, the bias is asymptotically zero for a linear decoding, iff

$$c \left(2P_{1|x} - 1 \right) - x = 0. \quad (7.36)$$

This condition means that $P_{1|x} = \frac{x}{2c} + \frac{1}{2}$, which holds for the case of uniform noise when $\sigma_\eta = 2c$. However, in general, the bias for a linear decoding is not asymptotically equal to zero for all x in the SSR model.

Nonlinear decoding

Rearranging the large N reconstruction points of Eqn. (7.19) gives $y(n) = NP_{1|x=\hat{x}_n}$. Since the average transfer function of the SSR model is $\bar{y}(x) = E[y|x] = NP_{1|x}$, then this implies that for large N the bias of the MMSE decoding asymptotically approaches zero for all x . This is in agreement with the known result that the maximum likelihood decoding is asymptotically unbiased (Lehmann and Casella 1998).

Consequently,

$$\lim_{N \rightarrow \infty} E[\hat{x}|x] = \sum_{n=0}^N F_R^{-1} \left(\frac{n}{N} \right) P(n|x) = x. \quad (7.37)$$

Since these arguments indicate that the large N bias is asymptotically zero for the maximum likelihood decoding, then how is it possible for the MASE to diverge, when we

7.4 Discussion on Stochastic Resonance Without Tuning

know from Chapter 6 that the MMSE distortion is smaller than the Wiener MSE distortion, for which large N expressions were presented in Section 7.2.1? The answer to this problem is related to the fact that for $\sigma < 1$, the most probable output states of the SSR model are states 0 and N . Recall also from Chapter 5 that the average transfer function PDF, $Q(\tau)$, is infinite at zero and unity.

The maximum likelihood reconstruction points that correspond to these states are $F_R^{-1}(0)$ and $F_R^{-1}(1)$, which for noise PDFs with infinite support are at $\pm\infty$. With such large values, if $P_y(0)$ and $P_y(N)$ do not approach zero, then the MSE distortion will be very large. This appears to be the case when the MASE diverges.

Also note that for the MMSE reconstruction points at $n = 0$ and $n = N$ to be equal to $\pm\infty$, N truly needs to be infinite. If it is not, then there is some large error between the MMSE points and the maximum likelihood points, and if $P_y(0)$ and $P_y(N)$ are not small, a significant difference between the MSE distortion obtained with the MMSE reconstruction points, and maximum likelihood reconstruction points will occur.

These issues mean that no meaningful analytical or numerical results can be obtained for the MASE for small σ , say $\sigma < 1$, and that no efficient unbiased estimator exists under these conditions. The best we can do is numerically calculate the MMSE reconstruction points and MMSE distortion for specific signal and noise distributions for as large an N as is practicable. Such results show the same qualitative behaviour as Fig. 7.1.

7.4 Discussion on Stochastic Resonance Without Tuning

We have seen that for large N the correlation coefficient increases with increasing σ and approaches unity as σ approaches infinity. Likewise, the MSE distortion decreases with increasing σ towards zero. Thus, there is a broadening of the system response, meaning near optimal performance occurs for a broad range of noise intensities. Furthermore, other than the case of uniform signal and noise, for infinite N the optimal noise intensity is at infinity, a strongly counterintuitive notion, even given an understanding of stochastic resonance. However, such a result has previously been published under the title of *stochastic resonance without tuning* (Collins *et al.* 1995).

The model examined in Collins *et al.* (1995) is quite similar to the SSR model, in that the outputs of N independently noisy individual neuron models are summed to give

an overall system output in response to aperiodic stimuli. By contrast to SSR however, Collins' stimuli are always subthreshold. Collins *et al.* (1995) use the correlation coefficient to show that SR occurs in this system, and emphasise the fact that as the number of elements, N , increases, that the peak of the correlation coefficient with noise intensity broadens, in the same manner as the results presented here. Note that the interpretations of Collins *et al.* (1995) have been criticised or explained in terms of well-known statistical results in several papers (Noest 1995, Petracchi 2000, Greenwood *et al.* 2003). The general theme is that the reported result should come as no surprise to those familiar with the averaging of independently noisy data. As is shown in Chapter 6, if a signal is averaged without thresholding, the MSE distortion is reduced by a factor of $\frac{1}{N}$. This means that regardless of the original SNR, the SNR of the averaged signal can be made arbitrarily small by increasing N . This is essentially the same result as that given in Collins *et al.* (1995), and here for large N and σ .

The difference between SSR and noise reduction by averaging, is that quantisation of the noisy input signal is performed in the SSR model, and this adds a component to the error between an input signal and the resultant output approximation. For sufficiently large noise, the quantisation loss can be overcome by averaging. However for small noise intensities, there is no decoding scheme that can overcome the mismatch of SSR quantisation. To explain this further, we saw in Chapter 5 that for small noise, the most probable output states are states close to zero and N . By contrast, for signals with infinite support PDFs, the most likely values are close to the mean. The results presented here indicate that even infinite N is not sufficient to overcome this mismatch between the shape of the output probability mass function, and the input PDF. By contrast as the noise intensity increases, this mismatch lessens, and at the optimum value of noise, the MSE distortion can be asymptotically reduced to zero by increasing N . However, for large σ , using large values of N is also very inefficient, because nearly all values of the SSR model's output, y , become highly improbable, other than those close to $y = N/2$. This has some analogies with *rate-distortion* theory (Berger and Gibson 1998). Information theorists understand that decreasing the average distortion in a quantisation scheme can be achieved by increasing the *rate*—here, *rate* can be interpreted to mean the number of output states, N . We will further discuss rate-distortion theory in Chapter 9.

7.5 Chapter Summary

Section 7.1 briefly introduces the aims of this Chapter, which are to examine the large N performance of the SSR model from the point of view of signal quantisation and reconstruction. Hence, we re-examine the results presented in Chapter 6, with the goal of finding large N limiting expressions.

Section 7.2 then presents large N asymptotic expressions for the Wiener decoding—that is, optimal linear decoding—reconstruction points, MSE distortion and correlation coefficient. Exact analytical formulas for these expressions are found in terms of σ for the cases of matched Gaussian, uniform and Laplacian signal and noise. Section 7.2 then goes on to consider MMSE nonlinear decoding, and finds an approximation that states that the large N reconstruction points are the maximum likelihood reconstruction points, and points out that with these points, an expression for the large N MMSE distortion is available. Verification of this result is given for the case of uniform signal and noise, and $\sigma \leq 1$, for which we have an exact expression for the large N reconstruction points. It is however, also pointed out that such an approximation breaks down for SSR and small σ , due to the approximation being more inaccurate for large and small n , since large and small n are also the most probable states.

Section 7.3 introduces large N asymptotic estimation theory, including the ideas of asymptotic unbiasedness and consistent estimators, before presenting an expression giving a lower bound in terms of Fisher information for the MSE distortion of an asymptotically unbiased estimator. The Fisher information for the SSR model is then shown to be such that the integral component of this lower bound does not always converge, meaning that asymptotically small MSE distortions are not achievable for small σ in the SSR model, even for infinite N .

Finally, Section 7.4 discusses the results presented in Section 7.2 which show that the optimal noise intensity appears to be infinite for infinite N , and why this result is somewhat misleading.

7.5.1 Original Contributions for Chapter 7

This chapter included the following original contributions on the large N decoding of the SSR model:

- Derivation of large N expressions for the optimal linear mean square error distortion, reconstruction points and linear correlation coefficient.
- Discussion of the MMSE distortion for large N , and the mean asymptotic square error, and how these expressions cannot be made asymptotically small for small noise intensities in the SSR model, even for infinite N .
- Discussion of the fact that for large N , the correlation coefficient and MSE distortion appear to be optimised for infinite noise intensity.

7.5.2 Further Work

Possible future work and open questions arising from this chapter might include:

- More rigorous analysis and justification of results showing that the MSE distortion cannot be made arbitrarily small for small noise intensities.

This concludes Chapter 7, which studies the decoding of the SSR model for large N . Chapter 8 introduces a new variation to the SSR model, by allowing the threshold values in each device in the model to vary independently.

Chapter 8

Optimal Stochastic Quantisation

AS described and illustrated in Chapters 4 to 7, a form of stochastic resonance called suprathreshold stochastic resonance can occur in a model system where more than one identical threshold device receives the same signal, but is subject to independent additive noise. In this Chapter we relax the constraint in this model that each threshold must have the same value, and aim to find the set of threshold values that either maximises the mutual information, or minimises the mean square error distortion, for a range of noise intensities. Such a task is a stochastic optimal quantisation problem. For sufficiently large noise, we find that the optimal quantisation is achieved when all thresholds have the same value. In other words, the suprathreshold stochastic resonance model provides an optimal quantisation for small input signal-to-noise ratios.

8.1 Introduction

The previous four chapters consider a form of stochastic resonance, known as Suprathreshold Stochastic Resonance (SSR), which occurs in an array of identical noisy threshold devices. The noise at the input to each threshold device is independent and additive, and this causes a randomisation of effective threshold values, so that all thresholds have unique, but random, effective values. Chapter 4 discusses and extends Stocks' result (Stocks 2000a) that the mutual information between the SSR model's input and output signals is maximised for some non-zero value of noise intensity. Chapter 6 considers how to reconstruct an approximation of the input signal by decoding the SSR model's output signal. It is shown that the Mean Square Error (MSE) distortion between the original signal and the reconstruction is minimised for some nonzero value of noise. Chapters 5 and 7 examine the large N behaviour of the mutual information and MSE distortion, and show that stochastic resonant behaviour persists for large N , and can be described by asymptotic formulas.

In this chapter, the constraint in the SSR model that all thresholds are identical is discarded, and we examine how to set the threshold values to optimise the performance of the resultant array of threshold devices for a range of noise intensities. Other than this modification, all other conditions are identical to the SSR model. The motivation for this modification is that—in the absence of noise—it is straightforward to show that setting all thresholds in such a model to identical values, as is the case for the SSR model, is not an optimal situation. To illustrate this, as discussed in Section 8.4, the maximum mutual information—obtained with *distributed* fixed threshold values—in the absence of noise is $\log_2(N + 1)$ bits per sample. In contrast, for SSR, we saw in Chapter 5 that the maximum mutual information is of the order of $0.5 \log_2(N)$.

Given that in the SSR case the mutual information is only 1 bit per sample in the absence of noise, it is clear that for small noise SSR is very far from optimal. The unresolved question is whether the SSR situation of all-identical thresholds is optimal for any range of noise intensities. The obvious way to address this question is to attempt to find the optimal thresholds as the noise intensity increases from zero. This is the focus of this chapter, and can be described as an optimal stochastic quantisation problem.

This Chapter contains original work on such optimal stochastic quantisation problems that has been published, in part, in the open literature (McDonnell *et al.* 2002c, McDonnell and Abbott 2004a, McDonnell *et al.* 2005c, McDonnell *et al.* 2005d, McDonnell *et al.* 2005b, McDonnell *et al.* 2006).

8.1.1 Optimal Quantisation Literature Review

Deterministic optimal quantisation

Theoretical results on optimal quantisation tend to focus on the conventional situation where quantisation is performed by fixed deterministic thresholds, and dithering is not considered. A comprehensive reference for such results is the textbook, Gersho and Gray (1992). Such research has usually been published in the electronic engineering literature, either in the field of information theory or communications theory.

Optimal stochastic quantisation

There have been very few previous studies on the problem of optimal quantisation in the presence of *independent* threshold noise. There has been a large amount of prior research into dithering¹³, however without considering how to optimise the thresholds. Other studies into noisy quantisation either do not consider input signal-to-noise ratios that are as large as those required for SSR occur, do not consider optimal quantisation in the presence of *independent* threshold noise or do not consider optimal quantisation as the input signal-to-noise ratio decreases from very large to very small values. For completeness, listed below are some references that, while at first glance might consider stochastic quantisation in the same manner as we do here, are actually about different topics:

- The existing work perhaps most closely related to this Chapter considers the topics of “random reference” correlation (Castanie *et al.* 1974) and “random reference quantizing” (Castanie 1979, Castanie 1984). In particular, in Castanie (1979) and Castanie (1984), the “transition points”—that is, threshold values—in a quantisation operation are independent random variables. However, unlike here, it is also assumed that the set of threshold values remains unique and ordered. In our notation, this means that the random variables corresponding to the i -th and $(i + 1)$ -th threshold values are such that $\theta_i + \eta_i < \theta_{i+1} + \eta_{i+1} \forall i$. We do

¹³Note that unlike the SSR effect, or what we refer to as “stochastic quantisation,” conventional dithering involves adding random or pseudo-random noise signals to an input signal prior to quantisation—whereas for SSR, the noise is *independent* for each threshold. Furthermore, as discussed in Section 2.3.3 of Chapter 2, dither signals are usually assumed to have a small dynamic range when compared to the input signal. All thresholds in a quantiser are subjected to the same dither signal, rather than independent noise.

not impose such restrictions here, and therefore there is no ordering of threshold values. Furthermore, no result presented in Castanie *et al.* (1974), Castanie (1979) and Castanie (1984) considers a quantiser's performance as a function of the threshold noise intensity.

- The “stochastic quantization” and “randomized quantisation” referred to in Berndt and Jentschel (2001) and Berndt and Jentschel (2002) does not refer to the case of independently randomised thresholds.
- The term “random quantization” used in Bucklew (1981) and Zador (1982) refers to random reconstruction points, rather than random thresholds, while the term “randomized quantizer” in Zamir and Feder (1992) refers to conventional dithering.
- There is a large body of research that considers noisy source coding, where a noisy signal needs to be compressed or coded, for example, Ayanoglu (1990). However generally, the noise is considered to be added to the signal prior to arriving at a quantiser. This has the effect of making all thresholds the same random variable, rather than being independent random variables, just as with dithering.
- The term “randomized quantizers” is also used in the context of statistical signal detection theory, but refers to a random choice between a number of deterministic quantisers (Tsitsiklis 1993, Blum 1995, Warren and Willett 1999).

Stochastic resonance, dithering and quantisation

In the SR literature, quantisation by a multithreshold system has only really been considered in the context of dithering (Gammaitoni 1995b, Wannamaker *et al.* 2000a), and not from the point of view of optimising the threshold values. Furthermore, the work on multithreshold systems contained in Gammaitoni (1995b) restricts the input signal to be entirely subthreshold, which in the language of quantization theory means that the input signal is always smaller than the quantiser bin-size. The noise signal is also not independent for each threshold. We point out that if an attempt had been made to optimise the threshold values in Gammaitoni (1995b), then the signal would no longer be subthreshold, and SR would not occur, due to the lack of noise independence.

By contrast, quantisation by a single threshold has been discussed in many papers, as pointed out elsewhere in this thesis, for example Bulsara and Zador (1996).

However, the extension of the SSR model to an optimal quantisation problem is briefly explored, for the first time, in Stocks (2000b). This paper firstly discusses optimal quantisation in the absence of noise, and points out that an exact solution to the optimal noiseless thresholds exists for maximum mutual information. It also comments on the fact that it seems there have been no previous studies on optimal quantisation when all thresholds are subject to internal noise. Stocks (2000b) does not attempt to solve this problem by finding the optimal thresholds in the presence of noise, but does plot the mutual information for a range of σ , i.e. the ratio of noise standard deviation to signal standard deviation, and various N , when the thresholds are set to maximise the mutual information in the absence of noise. These results show that the mutual information decreases monotonically with increasing σ , and hence SR does not occur. However, they also show that for sufficiently large σ , that the mutual information obtained with the SSR situation of all thresholds equal to the signal mean is greater than the mutual information obtained with the optimal noiseless thresholds. For small σ , SSR is very far from optimal.

Similar approaches examining independently noisy and distributed thresholds, as an extension of the SSR model were subsequently undertaken in McDonnell *et al.* (2002a) and also Rousseau and Chapeau-Blondeau (2005), without attempting to find the optimal thresholds values.

A very interesting observation made in Stocks (2000b) is that although better performance than SSR in the absence of noise—or very small noise—is obtained by a quantiser with optimal thresholds, the SSR situation may be more robust to non-stationary signal distributions, since if the signal distribution changes, the optimal noiseless thresholds will change.

Later, Stocks (2001b) analyses the simple case of $N = 3$ and Gaussian signal and noise, by finding the optimal thresholds for a range of values of σ . Stocks (2001b) states the assumption that by symmetry, one would expect that one of three threshold values would be zero, say $\theta_2 = 0$, and that $\theta_1 = -\theta_3$. With these assumptions, maximising the mutual information depends only on a single variable, that is, θ_1 . The mutual information can be calculated for various values of θ_1 to find the optimal value. Stocks (2001b) does not plot the optimal value of θ_1 against σ , but does plot the mutual information corresponding to the optimal θ_1 , which decreases monotonically with σ .

8.2 Optimal Quantisation Model

Neither Stocks (2000b) nor Stocks (2001b) discuss how the mutual information is calculated for the case of non-identical threshold values, however as is discussed in Section 8.2.3, numerical calculations are reasonably straightforward, especially for small N such as $N = 3$.

Thus, the sparse nature of previous work on optimal quantisation in the presence of independent threshold noise demonstrates that the area is wide open for study, and this Chapter now comprehensively investigates this problem.

8.1.2 Chapter Structure

Section 8.2 describes the stochastic quantisation model, and the differences between this model and the SSR model. It also mathematically describes the optimal stochastic quantisation problem as an unconstrained optimisation problem. Of particular importance is the outline of a method for recursively generating the transition probabilities for arbitrary thresholds and signal and noise distributions. Section 8.3 then describes relevant solution methods for optimisation problems like that discussed in Section 8.2. Next, Sections 8.4 and 8.5 present the results of finding the optimal threshold values for various numbers of threshold devices, signal and noise distributions, and the two different objectives of maximising the mutual information, and minimising the MSE distortion. Section 8.6 then discusses the results of Sections 8.4 and 8.5 and explains some of the key features of these results. Section 8.7 briefly examines how the optimality of SSR depends on N . The chapter is then summarised in the concluding section.

8.2 Optimal Quantisation Model

A schematic model of the system we examine in this Chapter is given by Fig. 4.1 in Chapter 4. Note that when we use this model in the SSR configuration, we set all thresholds to the same value, θ . Fig. 4.1 can be seen to be more general than SSR, since it explicitly shows that each threshold device may have different threshold values, labelled $\theta_1, \dots, \theta_N$.

As with the SSR model, the system consists of N threshold devices, which all receive the same signal, x . This signal is a sequence of *iid* samples drawn from a continuously valued distribution with Probability Density Function (PDF), $P(x)$. The i -th threshold device is subject to continuously valued *iid* additive noise, η_i , ($i = 1, \dots, N$) with PDF

$R(\eta)$. Each noise signal is also independent of the input signal. The output from each comparator, y_i , is unity if the input signal plus the noise, η_i is greater than the threshold, θ_i , of that device and zero otherwise. The outputs from each threshold, y_i , are summed to give the overall output signal, y . Hence, y is a discrete signal which can have integer values between 0 and N .

Unlike Chapter 4, we need to label the output of each of the i threshold devices independently in terms of the i -th threshold value, for a given x , as

$$y_i(x) = \begin{cases} 1 & \text{if } x + \eta_i \geq \theta_i, \\ 0 & \text{otherwise.} \end{cases} \quad (8.1)$$

The overall output of the array is still $y(x) = \sum_{i=1}^N y_i(x)$, which can be expressed in terms of the signum (sign) function as

$$y(x) = \frac{1}{2} \sum_{i=1}^N \text{sign}[x + \eta_i - \theta_i] + \frac{N}{2}. \quad (8.2)$$

As is the case in Chapter 4, the joint input-output PDF can be written as

$$P(x, y) = P(y = n|x)P(x) \quad (8.3)$$

$$= P(x|y = n)P_y(n), \quad (8.4)$$

where $P(y = n|x)$ are the *transition probabilities*, and $P_y(n)$ is the output probability mass function, which can be written as

$$P_y(n) = \int_{-\infty}^{\infty} P(y = n|x)P(x)dx, \quad n = 0, \dots, N. \quad (8.5)$$

We will always assume knowledge of $P(x)$, and from this point forward abbreviate our notation for the transition probabilities to $P(y = n|x) = P(n|x)$. Furthermore, in this Chapter the PDFs studied, i.e. $P(x)$ and $R(x)$, are always even functions about a mean of zero.

Recall from Table 4.1 in Chapter 4 that for the distributions considered, the variance of the signal is a function of σ_x and the variance of the noise is a function of σ_η . It was shown in Chapter 4 that for such distributions, when $\theta = 0$, the mutual information is always a function of the ratio $\sigma = \sigma_\eta/\sigma_x$. In this chapter we will again parameterise the noise intensity with this same inverse signal-to-noise ratio, σ . However, we also briefly pointed out in Chapter 4 that if $\theta \neq 0$, the mutual information will also be a function of θ/σ_x as well as σ and θ independently. We might therefore expect that for

8.2 Optimal Quantisation Model

arbitrary thresholds, the mutual information will also depend on these factors. This makes sense, as σ_x is generally a measure of the ‘width’ of the PDF $P(x)$. Therefore, if both σ_η and σ are fixed, then σ_x must change. If this happens, then we would intuitively expect the optimal thresholds to also vary with θ_i/σ_x . Thus, we arbitrarily set $\sigma_x = 1$ in the remainder of this Chapter. Therefore, results plotted against σ are equivalent to results plotted against σ_η , and results given for optimal thresholds and optimal reconstruction points can be taken to be optimal values of θ_i/σ_x and \hat{x}_n/σ_x .

To allow for non-identical thresholds, we generalise the notation used in Stocks (2000a) by letting $P_{1|x,i}$ be the probability of device i being ‘on’—that is, the probability that the sum of the input signal and noise exceeds the threshold value, θ_i —given the input signal x . Then

$$P_{1|x,i} = \int_{\theta_i-x}^{\infty} R(\eta) d\eta = 1 - F_R(\theta_i - x) \quad i = 1, \dots, N, \quad (8.6)$$

where $F_R(\cdot)$ is the Cumulative Distribution Function (CDF) of the noise. Since we assume $R(\eta)$ is an even function of η then

$$P_{1|x,i} = F_R(x - \theta_i). \quad (8.7)$$

Given a noise density, $R(\eta)$, and threshold value, θ_i , $P_{1|x,i}$ can be calculated exactly for any value of x from Eqn. (8.7).

8.2.1 Moment Generating Function

The *Moment Generating Function* (MGF) of a random variable, X , is defined (Yates and Goodman 2005) as

$$\phi_X(s) = E[\exp(Xs)], \quad (8.8)$$

where s is real-valued. The main use of the MGF is to derive the moments of a random variable, since it can be shown that the p -th moment of X is

$$E[X^p] = \left. \frac{d^p \phi_X(s)}{ds^p} \right|_{s=0}. \quad (8.9)$$

Now consider the conditional output of the array of threshold devices,

$$y(x) = \sum_{i=1}^N y_i(x). \quad (8.10)$$

Each $y_i(x)$ in the summation are independent random variables. It can be shown in such a situation that the MGF of $y(x)$ is the product of the MGFs of each $y_i(x)$ (Yates and Goodman 2005) and hence,

$$\phi_y(s) = \prod_{i=1}^N \phi_{y_i}(x). \quad (8.11)$$

However, we also have the probability mass function of y given x as $P(n|x)$. Therefore, from Eqn. (8.8),

$$\begin{aligned} \phi_y(s) &= \text{E}[\exp(ys)] \\ &= \sum_{n=0}^N P(n|x) \exp(ns). \end{aligned} \quad (8.12)$$

We also have for each individual $y_i(x)$ that

$$\begin{aligned} \phi_{y_i}(s) &= (1 - P_{1|x,i}) \exp(0s) + P_{1|x,i} \exp(1s) \\ &= 1 - P_{1|x,i} + P_{1|x,i} \exp(s). \end{aligned} \quad (8.13)$$

Substituting Eqn. (8.13) into Eqn. (8.11) and equating with Eqn. (8.12) gives a relationship between the set of $\{P(n|x)\}$ and the set of $\{P_{1|x,i}\}$ as

$$\phi_y(s) = \prod_{i=1}^N (1 - P_{1|x,i} + P_{1|x,i} \exp(s)) = \sum_{n=0}^N P(n|x) \exp(ns), \quad (8.14)$$

which holds for arbitrary threshold values.

Note that instead of defining the MGF in terms of $\exp(s)$, it is sometimes defined in terms of z^{-1} , where $z = \exp(-s)$. Using this approach, $P(n|x)$ can be described as the coefficient of z^{-n} in the power series expansion of $\prod_{i=1}^N [1 - P_{1|x,i} + z^{-1}P_{1|x,i}]$. That is,

$$\prod_{i=1}^N [1 - P_{1|x,i} + z^{-1}P_{1|x,i}] = \sum_{n=0}^N P(n|x) z^{-n}. \quad (8.15)$$

8.2.2 Conditional Output Moments

Since each $y_i(x)$ is independent, the expected value of y given x is the sum of the expected values of each y_i given x ,

$$\text{E}[y|x] = \sum_{i=1}^N P_{1|x,i}. \quad (8.16)$$

8.2 Optimal Quantisation Model

For the special case of SSR, each $P_{1|x,i} = P_{1|x}$ and $E[y|x] = NP_{1|x}$, which agrees with the result for SSR given in Chapter 6.

Similarly, an expression for the variance of y given x can be derived as the sum of the N individual variances. This fact is due to the random variable y given x being the sum of the N individual threshold outputs, y_i , each of which is an independent random variable for a given x (Yates and Goodman 2005). Thus, the covariance between y_i and y_j , $i \neq j$ is zero and

$$\text{var}[y|x] = \sum_{i=1}^N P_{1|x,i}(1 - P_{1|x,i}). \quad (8.17)$$

We saw in Chapter 6 that the variance of y given x for a binomial distribution is $NP_{1|x}(1 - P_{1|x})$, which also precisely agrees with Eqn. (8.17) for the SSR case.

These results for the conditional mean and variance can be easily verified from the expression for the MGF of Eqn. (8.14) and the relation of Eqn. (8.9).

8.2.3 Transition Probabilities

Recall from Chapter 4 that the transition probabilities for the SSR model are given by the binomial distribution, as in Eqn. (4.9). Although the previous Section derives a formula for the MGF of the conditional output distribution, this formula is not useful for specifying the distribution itself for arbitrary thresholds. However, it is possible to find $P(n|x)$ numerically.

This requires making use of the probability that any given threshold device is ‘on,’ given x . Assuming $P_{1|x,i}$ has been calculated for desired values of x , a convenient way of numerically calculating the probabilities $P(n|x)$ for a given number of threshold devices, N , is as follows.

Let $T_j^k(x)$ denote the probability that j of the thresholds $1, \dots, k$ are ‘on,’ given x . Then $T_0^1(x) = 1 - P_{1|x,1}$ and $T_1^1(x) = P_{1|x,1}$ and we have a set of recursive formulas,

$$\begin{aligned} T_0^{k+1}(x) &= (1 - P_{1|x,k+1})T_0^k(x), \\ T_j^{k+1}(x) &= P_{1|x,k+1}T_{j-1}^k(x) + (1 - P_{1|x,k+1})T_j^k(x) \quad j = 1, \dots, k, \\ T_{k+1}^{k+1}(x) &= P_{1|x,k+1}T_k^k(x). \end{aligned} \quad (8.18)$$

The recursion ends when $k + 1 = N$ and we have $P(n|x)$ given by $T_j^N(x) = T_n^N(x)$. We can rewrite Eqn. (8.18) in matrix form as

$$\begin{bmatrix} T_0^{k+1}(x) \\ T_1^{k+1}(x) \\ T_2^{k+1}(x) \\ \vdots \\ T_k^{k+1}(x) \\ T_{k+1}^{k+1}(x) \end{bmatrix} = \begin{bmatrix} 0 & T_0^k(x) \\ T_0^k(x) & T_1^k(x) \\ T_1^k(x) & T_2^k(x) \\ \vdots & \vdots \\ T_{k-1}^k(x) & T_k^k(x) \\ T_k^k(x) & 0 \end{bmatrix} \begin{bmatrix} P_{1|x,k+1} \\ 1 - P_{1|x,k+1} \end{bmatrix}. \quad (8.19)$$

Thus, to derive the transition probabilities, $P(n|x)$, for a given x and any arbitrary threshold settings and noise PDF, it suffices to firstly calculate the set of $P_{1|x,i}$ from Eqn. (8.6), and then to apply the recursive formula of Eqn. (8.19).

Note that this approach also works well in the special case of SSR where all thresholds are zero. Although in this case $P(n|x)$ can be calculated from the binomial formula, for large N using the binomial formula directly requires calculations of very large numbers, since $\binom{N}{n}$ can be very large. Such large numbers cannot be accurately represented in a computer. The alternative approach of using the recursive formulas of Eqn. (8.19) has the benefit that $P(n|x)$ can be found with excellent accuracy for very large values of N , since no such large numbers need to be calculated. In fact, since each T_j^k represents a probability, all numbers used in the calculation of each $P(n|x)$ are numbers between zero and unity.

The order of complexity of a software implementation of the recursive formula of Eqn. (8.19) is $O(N^2)$, for any given x . Thus, if N is doubled, then the run time of an implementation will be approximately quadrupled. More specifically, if $P_{1|x,i}$ and $1 - P_{1|x,i}$ are pre-computed, there will be $0.5N(N - 1)$ additions and $(N + 2)(N - 1)$ multiplications to find the transition probabilities.

Appendix B gives expressions for a number of probability distributions, and for each of these gives expressions for $P_{1|x}$ obtained directly from Eqn. (8.6) for threshold value, θ .

We are now ready to mathematically describe our optimal stochastic quantisation problem.

8.2.4 Optimisation Problem Formulation

We denote a vector of threshold values, $(\theta_1, \theta_2, \dots, \theta_N)$ as $\boldsymbol{\theta}$. Note that for the purpose of optimisation, the order of the values of the vector components is not important. The aim is to find the vector of thresholds, $\boldsymbol{\theta}^*$, that maximises some *objective function*, $f(\sigma, \theta_1, \theta_2, \dots, \theta_N) = f(\sigma, \boldsymbol{\theta})$, as the noise parameter, σ varies. The optimal value of the objective function for a given σ is then $f(\sigma, \boldsymbol{\theta}^*)$. Such a problem can be formulated for a given value of σ as the nonlinear—since the objective function will be nonlinear—optimisation problem

$$\begin{aligned} \text{Find:} \quad & \max_{\boldsymbol{\theta}} f(\sigma, \boldsymbol{\theta}) \\ \text{subject to:} \quad & \boldsymbol{\theta} \in \mathbb{R}^N. \end{aligned} \tag{8.20}$$

Such a problem can be solved by standard optimisation techniques, as will be described in Section 8.3.

8.3 Optimisation Solution Algorithms

8.3.1 Unconstrained Multidimensional Optimisation

The optimisation problem of (8.20) simply describes an N dimensional unconstrained function optimisation. The function $f(\sigma, \boldsymbol{\theta})$ is, as we will see, a nonlinear function of $\boldsymbol{\theta}$. Solving such a problem is amenable to many standard techniques (Nocedal and Wright 1999), most of which rely on calculations of the *gradient* of the objective function with respect to the free variables, which in this case is the vector $\boldsymbol{\theta}$,

$$\nabla f(\sigma, \boldsymbol{\theta}) = \left(\frac{df}{d\theta_1}, \frac{df}{d\theta_2}, \dots, \frac{df}{d\theta_N} \right). \tag{8.21}$$

Extrema—maxima, minima or saddle points—occur when the gradient vector is equal to $\mathbf{0}$.

Standard optimisation techniques usually start with some arbitrary initial solution, say $\boldsymbol{\theta}_0$, and then apply some method of finding a new solution, $\boldsymbol{\theta}_1$, from $\boldsymbol{\theta}_0$, that increases the value of the objective function. This process is repeated in an iterative fashion, until no further increase in the objective function can be found. Finding a new vector that increases the value of the objective function is often achieved by finding the gradient vector at a current solution, $\boldsymbol{\theta}_k$, and then performing a one dimensional search for the maximum in the direction of the gradient.

The method used to obtain the results presented in this Chapter is a method called the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm. This algorithm falls into the sub-category of optimisation algorithms known as *quasi-Newton methods* (Press *et al.* 1992). The BFGS method was found to be superior in speed and convergence to other similar methods such as the Fletcher-Reeves-Polak-Ribiere conjugate gradient method, and the Powell conjugate gradient method. Further information on these algorithms can be found in Nocedal and Wright (1999) or Press *et al.* (1992).

Note that although this is not relevant for an implementation of such an optimisation procedure, changing the threshold vector will induce a new set of transition probabilities, $P(n|x)$, which will give a new value of the objective function. Thus, it is equivalent to write the optimisation problem of (8.20) in terms of an optimisation over the $N + 1$ correlated functions of x that describe the transition probabilities. Such formulations are often solved using the calculus of variations.

For example, our optimisation problem is similar to previous work on clustering and neural coding problems solved using a combination of the calculus of variations and a method known as deterministic annealing (Rose *et al.* 1990, Rose *et al.* 1992, Rose 1994, Rose 1998, Tishby *et al.* 1999, Dimitrov and Miller 2001, Dimitrov *et al.* 2003). In particular, the formulation reached in Dimitrov and Miller (2001) can be expressed in a fashion identical to Problem (8.20) with one exception. Here, we have *structural constraints* on how the transition probabilities, $P(n|x)$, are formed, since—as expressed by Eqn. (8.19)—each $P(n|x)$ is a function of the set of probabilities, $P_{1|x,i}$. Due to this difference, the solution method used in Dimitrov and Miller (2001) to find the optimal conditional distribution, $\{P(n|x)\}$, cannot be used here, and instead we concentrate on optimizing the only free variable, the vector of threshold values, θ .

8.3.2 Dealing with Local Optima

One of the major problems with standard deterministic optimisation methods such as the BFGS method, is that they do not cope well when applied to objective functions that possess more than one local minima or maxima. If the objective function is known to be convex, this is not an issue. However, if it is not convex, then these standard methods will converge towards a local optimum, but give no guarantees about whether or not this optimum is the global optimum. The local optimum found depends on the initial solution used. Hence, one possibility for trying to find a good local optimum

8.3 Optimisation Solution Algorithms

is to run a standard optimisation procedure many times, using many different initial solutions, and then pick the solution found with the largest value—if the problem is a maximisation—of the objective function.

Numerical experiments show that for the objective functions considered in this chapter, there are many local maxima. In particular, it seems that the SSR situation of all thresholds zero is always a local optimum. This means that the gradient of the mutual information with respect to θ is always the zero vector, $\mathbf{0}$, at $\theta = \mathbf{0}$.

However, we will see that there is a special structure to the local optima, which make it theoretically possible to find all local optima. In particular, numerical experiments indicate that there are at most two locally optimal solutions corresponding to each possible ordered clustering of N thresholds to unique values. For example, in the case of $N = 3$, there are four locally optimal situations—1) the case of all thresholds equal to the same value; 2) the case of all three threshold values unique; 3) the case of two thresholds being equal to the same positive value, and the third threshold being negative; and 4) the case of one threshold being positive, and the other two thresholds being equal to the same negative value. Note that the third and fourth cases are effectively equivalent, as the optimal threshold values in the third case are the negative of the optimal threshold values in the fourth case, and give the same function maximum.

Thus, in order to find the globally optimal solution, all possible ‘clusters’ of N thresholds can be trialed. The dimension of the optimisation problem in each case is the number of unique threshold values. The locally optimal thresholds only change very slightly for a small change in σ . Hence, for each possible local solution, once the optimal thresholds are found for a particular value of σ , these values can be used as the initial values for a new value of σ a very small increment larger.

This procedure works well for the results presented in this Chapter where N is no greater than five, however, the number of local optima increases combinatorially with N , and it is not practical to consider all local optima for N any larger than 5. As mentioned, one way of dealing with this is to simply try a number of random initial solutions, and pick the best. Another way is to employ random search optimisation algorithms such as simulated annealing (Kirkpatrick *et al.* 1983) or genetic algorithms (Beasley *et al.* 1993). Unlike deterministic methods, such algorithms are known to be able to ‘climb’ out of local optima, and provide a greater chance of finding a global optimum. However, for the small N results presented in this Chapter, it was found that the deterministic combination of the BFGS and initial solution methods outlined above

is far superior to random search methods in runtime and for ensuring all local-optima are tracked.

Having now described the solution methods we use, the next two Sections give results for the optimal stochastic quantisation for maximum mutual information and minimum MSE distortion.

8.4 Optimal Quantisation for Mutual Information

Recall from Chapter 4 that the mutual information between the input and output signals of the array of threshold devices in Fig. 4.1 is given by

$$\begin{aligned} I(x, y) &= H(y) - H(y|x) \\ &= - \sum_{n=0}^N P_y(n) \log_2 P_y(n) - \left(- \int_{-\infty}^{\infty} P(x) \sum_{n=0}^N P(n|x) \log_2 P(n|x) dx \right), \end{aligned} \quad (8.22)$$

where $H(y)$ is the entropy of the output signal, and $H(y|x)$ is the average conditional output entropy. Given the procedure described in Section 8.2.3 for numerically calculating the transition probabilities for any value of x , the mutual information can be calculated for arbitrary thresholds and $P(x)$ by numerical integration of Eqn. (8.22). Note that this integration requires calculations of $P(n|x)$ from Eqn. (8.19) for many values of x .

Given the aim of maximising the mutual information, the optimisation problem is

$$\begin{aligned} \text{Find:} \quad & \max_{\boldsymbol{\theta}} I(x, y) \\ \text{subject to:} \quad & \boldsymbol{\theta} \in \mathbb{R}^N. \end{aligned} \quad (8.23)$$

8.4.1 Absence of Noise

As pointed out in Stocks (2000b), an exact solution to Problem (8.23) can be found in the absence of noise, that is, when $\sigma = 0$. In this case, the average conditional output entropy, $H(y|x)$, is zero, and therefore the mutual information is the output entropy, $H(y)$. Maximising the mutual information reduces to maximising the output entropy. It is well known that for a discrete random variable this occurs when all states are equally likely (Cover and Thomas 1991), and hence $I(x, y) = \log_2(N + 1)$. Thus, in

8.4 Optimal Quantisation for Mutual Information

the absence of noise, we require that the output probability mass function is

$$P_y(n) = \frac{1}{N+1} \quad n = 0, \dots, N. \quad (8.24)$$

If we let $\theta_n > \theta_{n-1}$, then—since there is no threshold noise—the output, y , is in state n iff $x \in [\theta_{n-1}, \theta_n]$. This means that there is no ambiguity about what output state corresponds to a given value of x , and that therefore $P(n|x) = 1 \forall x \in [\theta_{n-1}, \theta_n]$. More precisely, since $P_y(n) = \int_x P(x)P(n|x)dx$, this requires that

$$\begin{aligned} \int_{-\infty}^{\theta_1} P(x)dx &= \frac{1}{N+1} \\ \int_{\theta_i}^{\theta_{i+1}} P(x)dx &= \frac{1}{N+1}, \quad i = 1, \dots, n-1 \\ \text{and } \int_{\theta_N}^{\infty} P(x)dx &= \frac{1}{N+1}. \end{aligned} \quad (8.25)$$

This implies that

$$F_x(\theta_i) = \frac{i}{N+1}, \quad i = 1, \dots, N, \quad (8.26)$$

where $F_x(\cdot)$ is the CDF of the input signal, and that therefore

$$\theta_i = F_x^{-1}\left(\frac{i}{N+1}\right), \quad i = 1, \dots, N, \quad (8.27)$$

where $F_x^{-1}(\cdot)$ is the Inverse Cumulative Distribution Function (ICDF) of the input signal.

Hence, a simple formula for the optimal thresholds for mutual information is available in the absence of noise and shows that all N optimal thresholds are uniquely valued.

Optimal decoding for maximised $I(x, y)$ in the absence of noise

If the thresholds are set to maximise the noiseless mutual information, then from Eqn. (8.40) the Minimum Mean Square Error (MMSE) reconstruction points are given by

$$\begin{aligned} \hat{x}_0 &= (N+1) \int_{-\infty}^{\theta_1} xP(x)dx, \\ \hat{x}_n &= (N+1) \int_{\theta_n}^{\theta_{n+1}} xP(x)dx \quad n = 1, \dots, N-1 \\ \hat{x}_N &= (N+1) \int_{\theta_N}^{\infty} xP(x)dx, \end{aligned} \quad (8.28)$$

since all output states are equally probable. Substituting Eqn. (8.27) into Eqn. (8.28) gives the exact optimal MMSE reconstruction points for the thresholds that maximise the noiseless information.

We can simplify this substitution by making a change of variable. Let $x = F_x^{-1}\left(\frac{\tau}{N+1}\right)$. Then $\tau = (N+1)F_x(x)$ and $d\tau = (N+1)P(x)dx$. Carrying this out in Eqn. (8.28), and using Eqn. (8.26) gives

$$\hat{x}_n = \int_{\tau=n}^{\tau=n+1} F_x^{-1}\left(\frac{\tau}{N+1}\right) d\tau \quad n = 0, \dots, N. \quad (8.29)$$

The integral in Eqn. (8.29) can be solved exactly in some cases. For example, for a uniformly distributed signal we get

$$\hat{x}_n = \sigma_x \left(\frac{2n+1}{2N+2} - \frac{1}{2} \right) \quad n = 0, \dots, N, \quad (8.30)$$

and for a logistically distributed signal we get

$$\hat{x}_n = \frac{\sqrt{3}\sigma_x}{\pi} \left((n+1) \ln(n+1) - n \ln(n) + (N-n) \ln(N-n) - (N+1-n) \ln(N+1-n) \right), \quad n = 0, \dots, N. \quad (8.31)$$

8.4.2 Results in the Presence of Noise

Results for $N = 2, \dots, 5$

We now present the results of solving Problem (8.23) for nonzero σ , and various matched signal and noise distributions. The optimal thresholds as found by the deterministic algorithm explained in Section 8.3, are plotted for $N = 2, 3, 4$ and 5 , and the optimal reconstruction points for the thresholds of the $N = 5$ case. Discussion of these results is left for Section 8.6.

Results for $N = 15$

Fig. 8.6 shows the optimal thresholds and mutual information for $N = 15$ and Gaussian signal and noise. The same qualitative behaviour as for each $N \leq 5$ case can be seen. However, we cannot be completely sure that the thresholds shown provide the globally optimal solution, as the number of local optima is now very large. However, provided only near optimal thresholds are required, we can find an approximation to the optimal thresholds for reasonably large N by applying the BFGS method for various initial conditions and can still be reasonably confident that the threshold values shown provide values of the mutual information very close to the optimal solution.

The next Section looks at optimal quantisation for the objective of minimising the MSE distortion.

8.4 Optimal Quantisation for Mutual Information

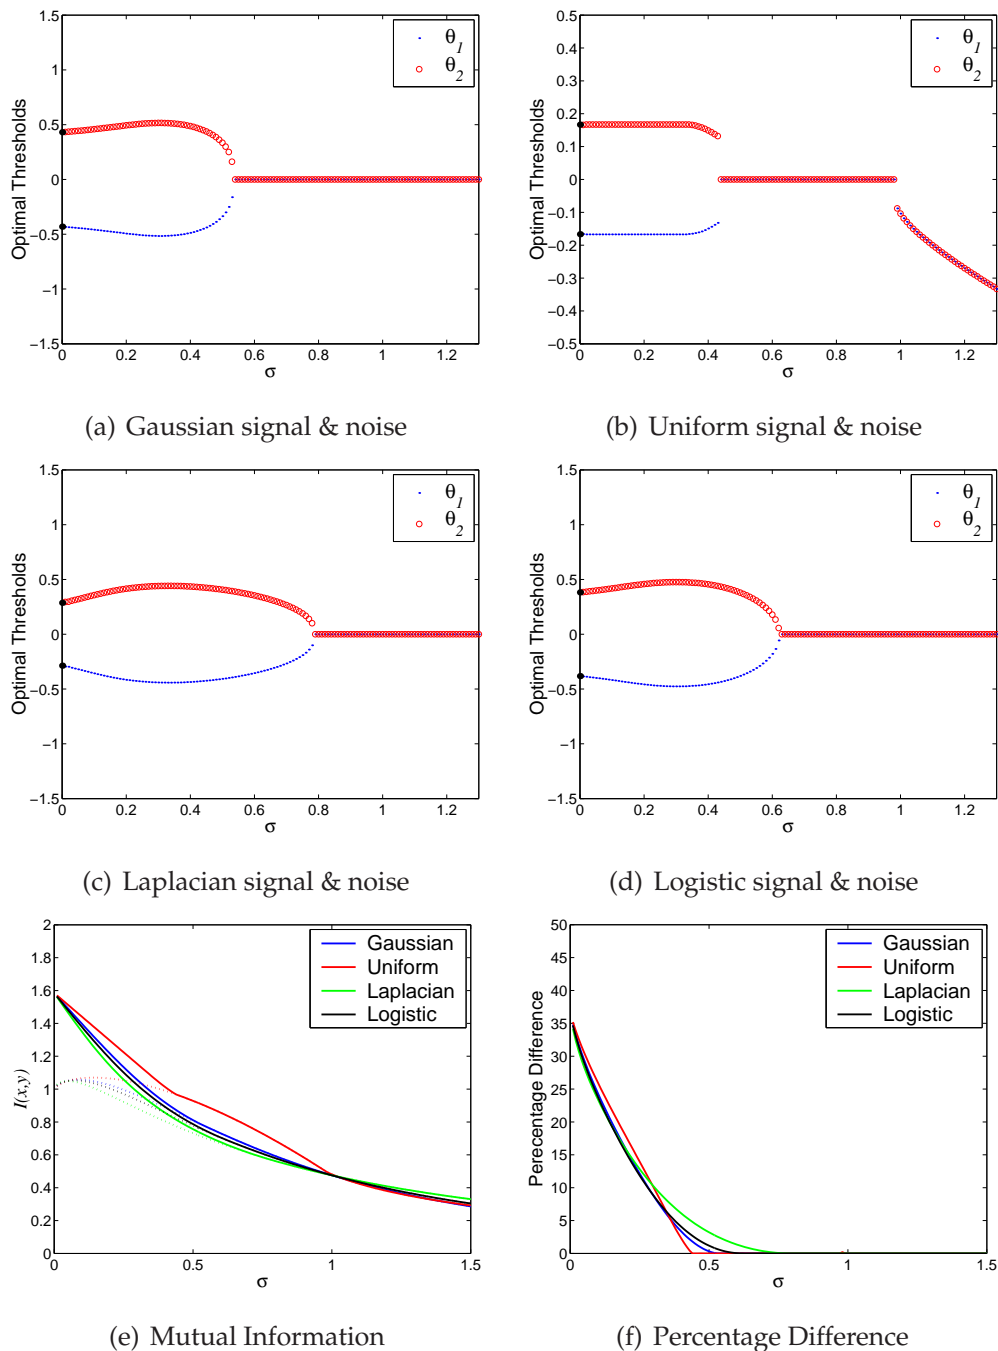


Figure 8.1. Optimal thresholds for mutual information, $N = 2$. Figs. 8.1(a), 8.1(b), 8.1(c) and 8.1(d) show plots of the optimal thresholds for four different matched signal and noise pairs against increasing noise intensity, σ . The optimal noiseless threshold values calculated from Eqn. (8.27) are shown by large black dots. Fig. 8.1(e) shows the mutual information obtained with these optimal threshold settings against increasing σ , as well as the mutual information for SSR shown with dotted lines. Fig. 8.1(f) shows the percentage difference between the mutual information obtained by optimally setting the thresholds, and that for SSR, against increasing σ . For σ greater than some critical value, σ_c , the SSR situation can be seen to be optimal.

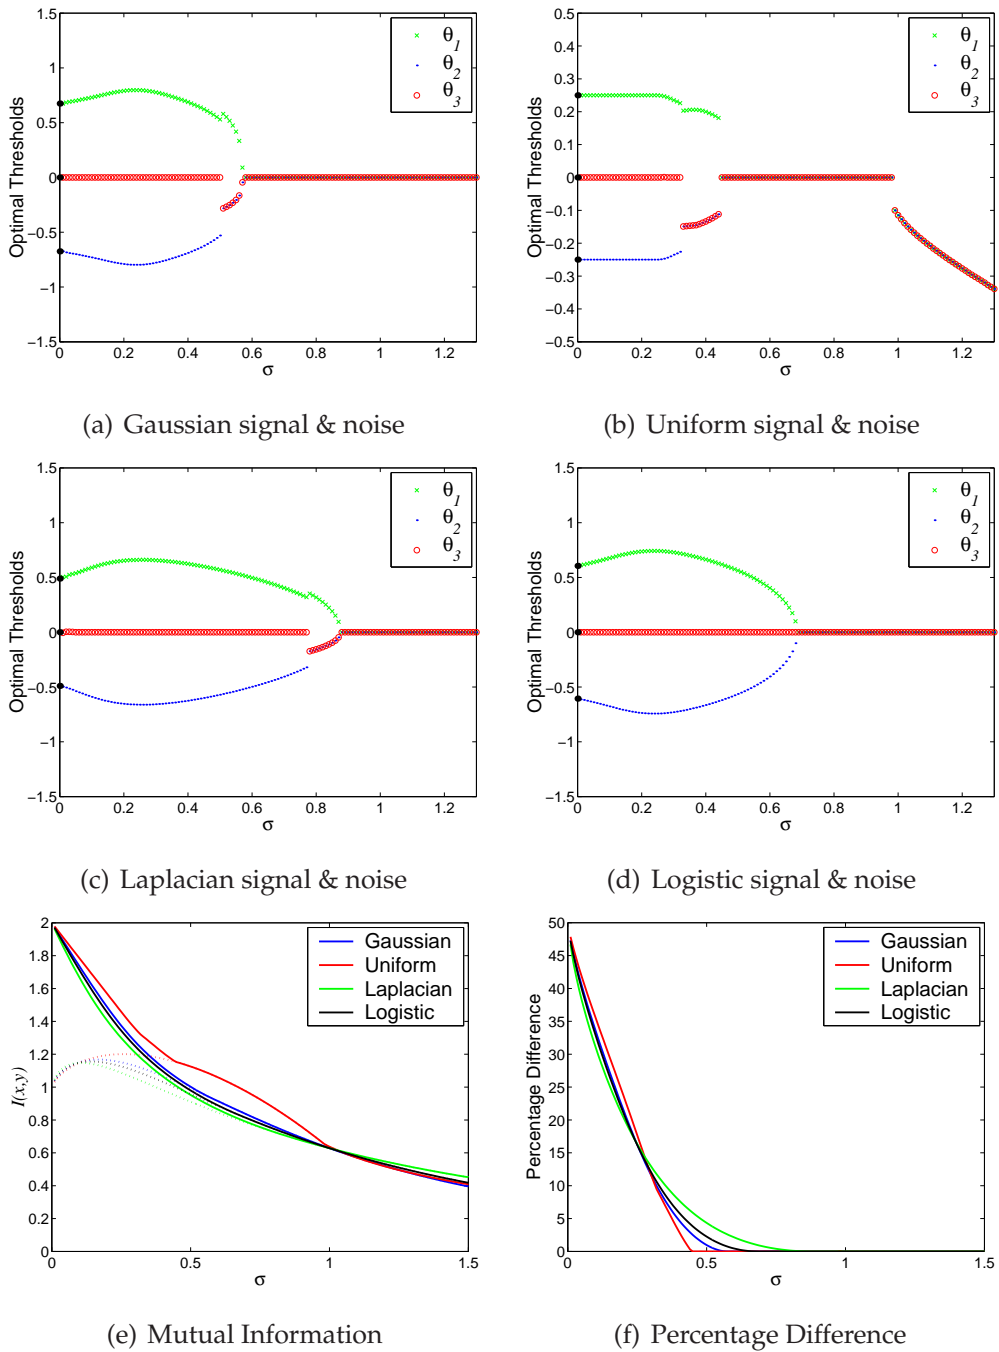


Figure 8.2. Optimal thresholds for mutual information, $N = 3$. Figs. 8.2(a), 8.2(b), 8.2(c) and 8.2(d) show plots of the optimal thresholds for four different matched signal and noise pairs against increasing noise intensity, σ . The optimal noiseless threshold values calculated from Eqn. (8.27) are shown by large black dots. Fig. 8.2(e) shows the mutual information obtained with these optimal threshold settings against increasing σ , as well as the mutual information for SSR shown with dotted lines. Fig. 8.2(f) shows the percentage difference between the mutual information obtained by optimally setting the thresholds, and that for SSR, against increasing σ . For σ greater than some critical value, σ_c , the SSR situation can be seen to be optimal.

8.4 Optimal Quantisation for Mutual Information

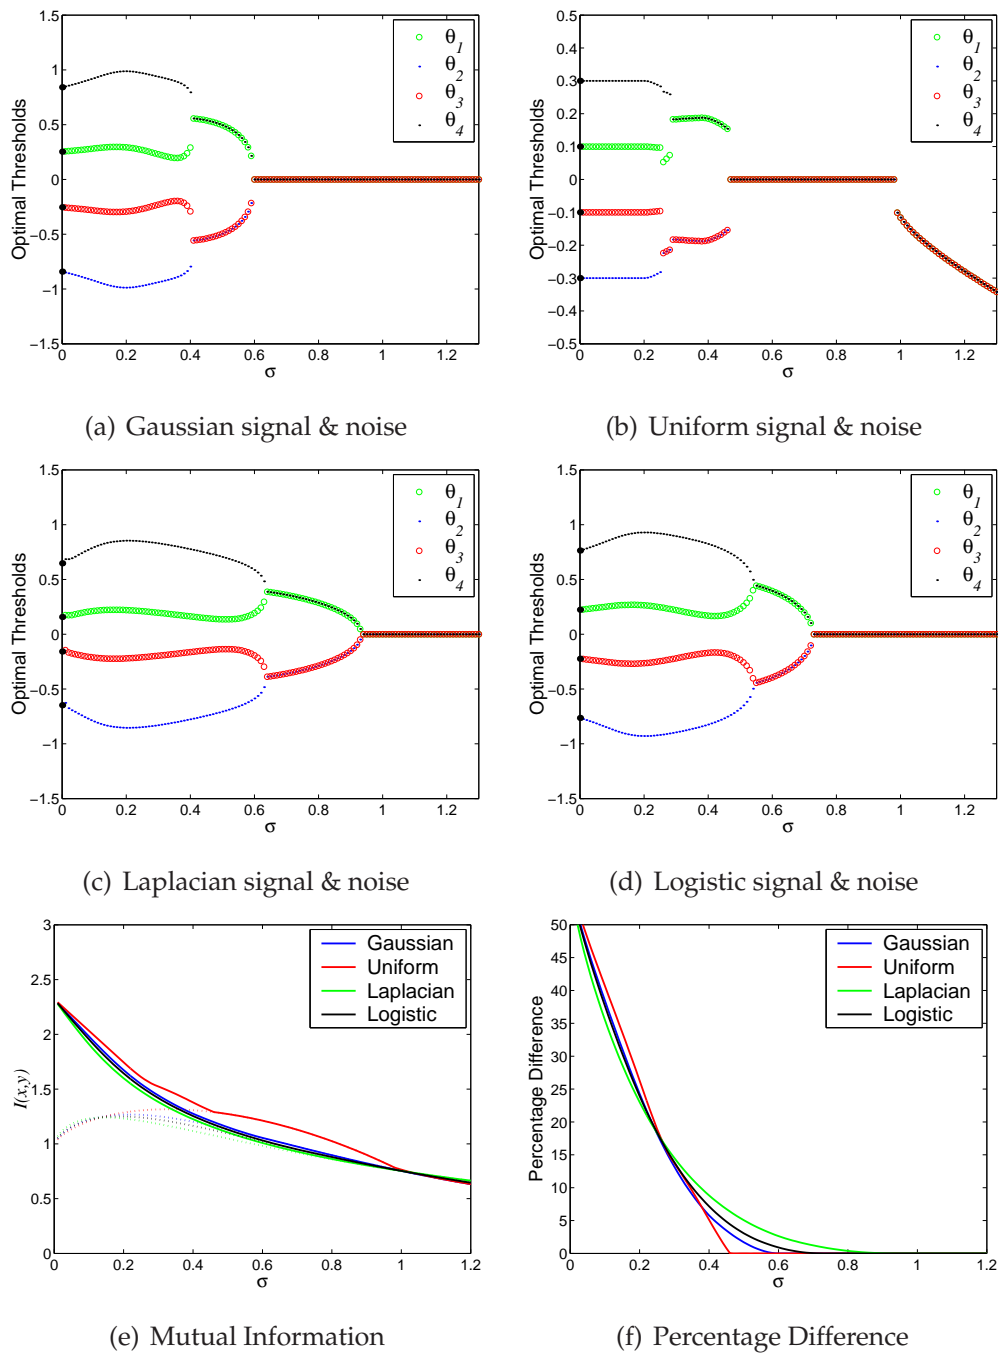


Figure 8.3. Optimal thresholds for mutual information, $N = 4$. Figs. 8.3(a), 8.3(b), 8.3(c) and 8.3(d) show plots of the optimal thresholds for four different matched signal and noise pairs against increasing noise intensity, σ . The optimal noiseless threshold values calculated from Eqn. (8.27) are shown by large black dots. Fig. 8.3(e) shows the mutual information obtained with these optimal threshold settings against increasing σ , as well as the mutual information for SSR shown with dotted lines. Fig. 8.3(f) shows the percentage difference between the mutual information obtained by optimally setting the thresholds, and that for SSR, against increasing σ . For σ greater than some critical value, σ_c , the SSR situation can be seen to be optimal.

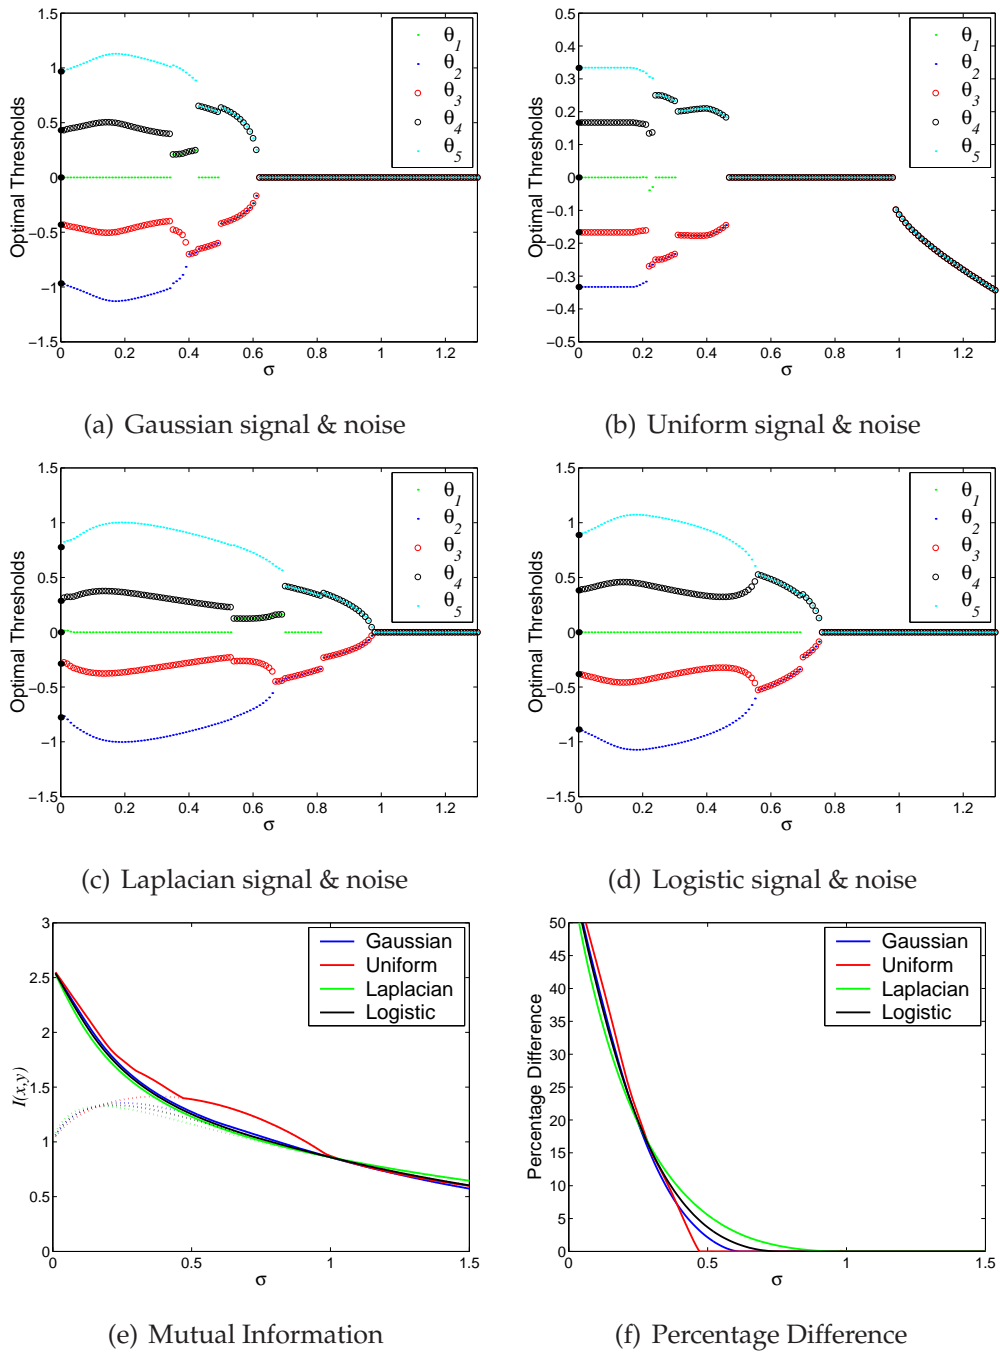





Figure 8.4. Optimal thresholds for mutual information, $N = 5$. Figs. 8.4(a), 8.4(b), 8.4(c) and 8.4(d) show plots of the optimal thresholds for four different matched signal and noise pairs against increasing noise intensity, σ . The optimal noiseless threshold values calculated from Eqn. (8.27) are shown by large black dots. Fig. 8.4(e) shows the mutual information obtained with these optimal threshold settings against increasing σ , as well as the mutual information for SSR shown with dotted lines. Fig. 8.4(f) shows the percentage difference between the mutual information obtained by optimally setting the thresholds, and that for SSR, against increasing σ . For σ greater than some critical value, σ_c , the SSR situation can be seen to be optimal.




(a) Gaussian signal & noise



(b) Uniform signal & noise



(c) Laplacian signal & noise



(d) Logistic signal & noise

Figure 8.5. Optimal reconstruction points for mutual information, $N = 5$. Figs. 8.5(a), 8.5(b), 8.5(c) and 8.5(d) show plots of the optimal reconstruction points for four different matched signal and noise pairs against increasing noise intensity, σ . The optimal noiseless reconstruction points calculated from Eqn. (8.29) are shown by large black dots.

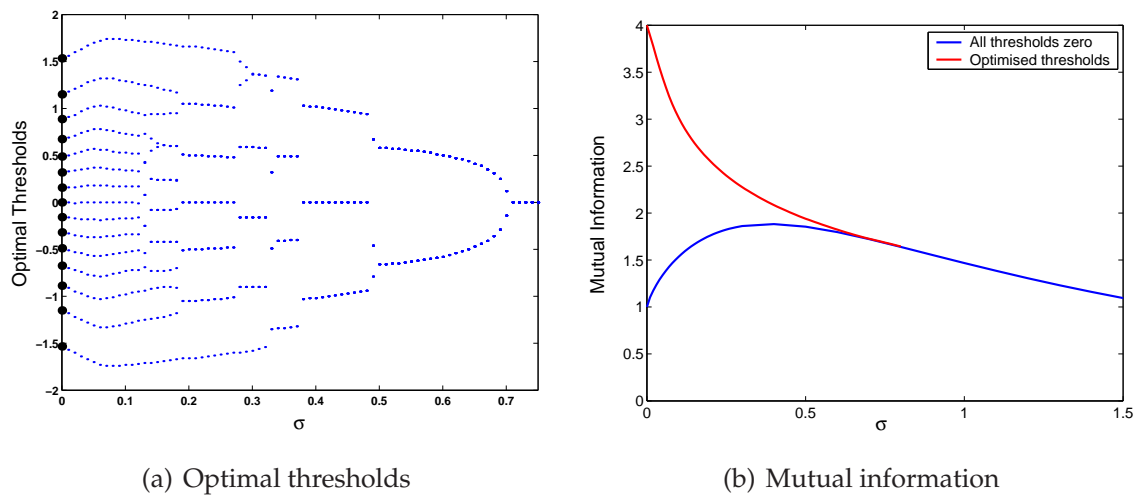


Figure 8.6. Optimal thresholds for mutual information, Gaussian signal and noise, $N = 15$.

Fig. 8.6(a) shows the optimal thresholds for Gaussian signal and noise against increasing noise intensity, σ . The optimal noiseless threshold values calculated from Eqn. (8.27) are shown by large black dots. Fig. 8.6(b) shows the mutual information for the thresholds shown in Fig. 8.6(a), as well as the mutual information obtained in the SSR case. As with the $N = 2, \dots, 5$ cases, for σ greater than some critical value, σ_c , the SSR situation can be seen to be optimal.

8.5 Optimal Quantisation for MSE Distortion

8.5.1 Linear Decoding

Recall from Chapter 6 that for zero-mean input signals and a linear decoding of the form $\hat{y} = ay + b$, that if the condition $E[\hat{y}] = 0$ is imposed, then $b = -aE[y]$. Note that while for SSR the expected value of the output is $E[y] = N/2$, this will certainly not necessarily be the case for arbitrary thresholds.

As discussed in Chapter 6, with such a linear decoding, the optimal value of a is given by $a = \frac{E[xy]}{\text{var}[y]}$, so that the optimal linear reconstruction points are

$$\hat{y} = \frac{E[xy]}{\text{var}[y]} (y - E[y]). \quad (8.32)$$

Such a decoding is the optimal linear decoding, and is known as the Wiener decoding. The MSE distortion with such a decoding can be written in terms of the correlation coefficient of a linear decoding, ρ_{xy} , as

$$\text{MSE} = E[x^2](1 - \rho_{xy}^2), \quad (8.33)$$

where the correlation coefficient can be expressed as

$$\rho_{xy} = \frac{E[xy]}{\sqrt{E[x^2]\text{var}[y]}}. \quad (8.34)$$

Minimising the linear decoding MSE is equivalent to maximising the linear decoding correlation coefficient. Since we also assume knowledge of the input signal PDF, and therefore of its mean square value, $E[x^2]$, minimising the linear decoding MSE distortion is equivalent to solving the optimisation problem,

$$\begin{aligned} \text{Find:} \quad & \max_{\theta} \frac{E[xy]}{\sqrt{\text{var}[y]}} \\ \text{subject to:} \quad & \theta \in \mathbb{R}^N. \end{aligned} \quad (8.35)$$

8.5.2 Nonlinear Decoding

Recall from Chapter 6 that the optimal MSE decoding is the nonlinear decoding given by $\hat{x}_n = E[x|n]$. This decoding results in the minimum possible MSE distortion for given transition probabilities, which is given by

$$\text{MMSE} = E[x^2] - E[\hat{x}^2]. \quad (8.36)$$

Thus, since $E[x^2]$ is known, if the aim is to find the thresholds that minimise the MSE distortion resulting from using the MMSE decoding, the optimisation problem is

$$\begin{aligned} \text{Find:} \quad & \max_{\boldsymbol{\theta}} E[\hat{x}^2] \\ \text{subject to:} \quad & \boldsymbol{\theta} \in \mathbb{R}^N. \end{aligned} \quad (8.37)$$

8.5.3 Absence of Noise

Unlike maximising the mutual information, in general no simple formula exists for the threshold values that minimise the MSE distortion in the absence of noise for an arbitrary signal PDF. However, the optimal noiseless thresholds can be easily found numerically for a given $P(x)$, and the reconstruction points shown to all be unique.

In the absence of noise, assume all thresholds are unique, and that therefore a given value of the output, $y = n$, can only be achieved by values of x that lie between consecutive thresholds, say θ_{i-1} and θ_i . The optimal MMSE reconstruction points are given by

$$\hat{x}_n = E_x[x|n] = \int_x xP(x|n)dx, \quad (8.38)$$

and thus, since $P(n|x)$ is unity for $x \in [\theta_n, \theta_{n+1}]$, then

$$\begin{aligned} P_y(0) &= \int_{-\infty}^{\theta_1} P(x)dx, \\ P_y(n) &= \int_{\theta_n}^{\theta_{n+1}} P(x)dx \quad n = 1, \dots, N-1 \\ P_y(N) &= \int_{\theta_N}^{\infty} P(x)dx, \end{aligned} \quad (8.39)$$

and

$$\begin{aligned} \hat{x}_0 &= \frac{1}{P_y(0)} \int_{-\infty}^{\theta_1} xP(x)dx, \\ \hat{x}_n &= \frac{1}{P_y(n)} \int_{\theta_n}^{\theta_{n+1}} xP(x)dx \quad n = 1, \dots, N-1 \\ \hat{x}_N &= \frac{1}{P_y(N)} \int_{\theta_N}^{\infty} xP(x)dx. \end{aligned} \quad (8.40)$$

Therefore, each MMSE reconstruction point is the centroid of the corresponding partition of the input PDF. The MMSE distortion is

$$\text{MMSE} = E[x^2] - \sum_{n=0}^N \frac{1}{P_y(n)} \left(\int_{\theta_n}^{\theta_{n+1}} xP(x)dx \right)^2. \quad (8.41)$$

8.5 Optimal Quantisation for MSE Distortion

In the absence of noise, finding the optimal thresholds for the MMSE distortion can be achieved by a simple iterative procedure known as the *Lloyd Method I algorithm* (Lloyd 1982), which is a commonly used technique for finding the optimal quantisation for a given source PDF. A second algorithm known as the Lloyd Method II algorithm is also called the Lloyd-Max algorithm, due its rediscovery by Max (1960)¹⁴. See also Gersho and Gray (1992) and Gray and Neuhoff (1998) for more details. The Lloyd Method I algorithm begins with an initial guess for the reconstruction points, and then finds the optimal thresholds for those points, which for MSE distortion, and any given decoding, can be shown to simply be the midpoints of the reconstruction points. Given these new thresholds, a new set of reconstruction points are found from Eqn. (8.40), and the new MMSE distortion from Eqn. (8.41). This new MMSE distortion can be shown to smaller than the previous MMSE distortion. This iteration is repeated until the MMSE distortion no longer decreases, at which point the optimal noiseless thresholds and reconstruction points have been found.

The Lloyd method I algorithm can also be used for the case of finding the optimal thresholds for a linear decoding. The only difference is that instead of calculating the new optimal reconstruction points at each iteration from Eqn. (8.40), the optimal linear reconstruction points are calculated from Eqn. (8.32), and the resulting MSE distortion from Eqn. (8.33). The update of the optimal thresholds at each iteration remains as the midpoint between the current reconstruction points.

8.5.4 Results in the Presence of Noise

We now present the results of solving Problems (8.35) and (8.37) for nonzero σ and various matched signal and noise distributions. Due to the results having the same qualitative behaviour of those found in the case of maximising the mutual information, we present optimal thresholds only for the case of $N = 5$ for minimised linear decoding MSE distortion, and $N = 2$ and $N = 5$ for minimised MMSE distortion, and the optimal reconstruction points for $N = 5$ and MMSE distortion. Since the uniform distribution has a different mean square value than the other distributions for the same value of σ , we plot the SQNR rather than the MSE distortion.

¹⁴Note that both of Lloyd's algorithms were first described in an unpublished Bell Laboratories technical report in 1957 (Gersho and Gray 1992), and not published in the open literature until 1982 (Lloyd 1982).

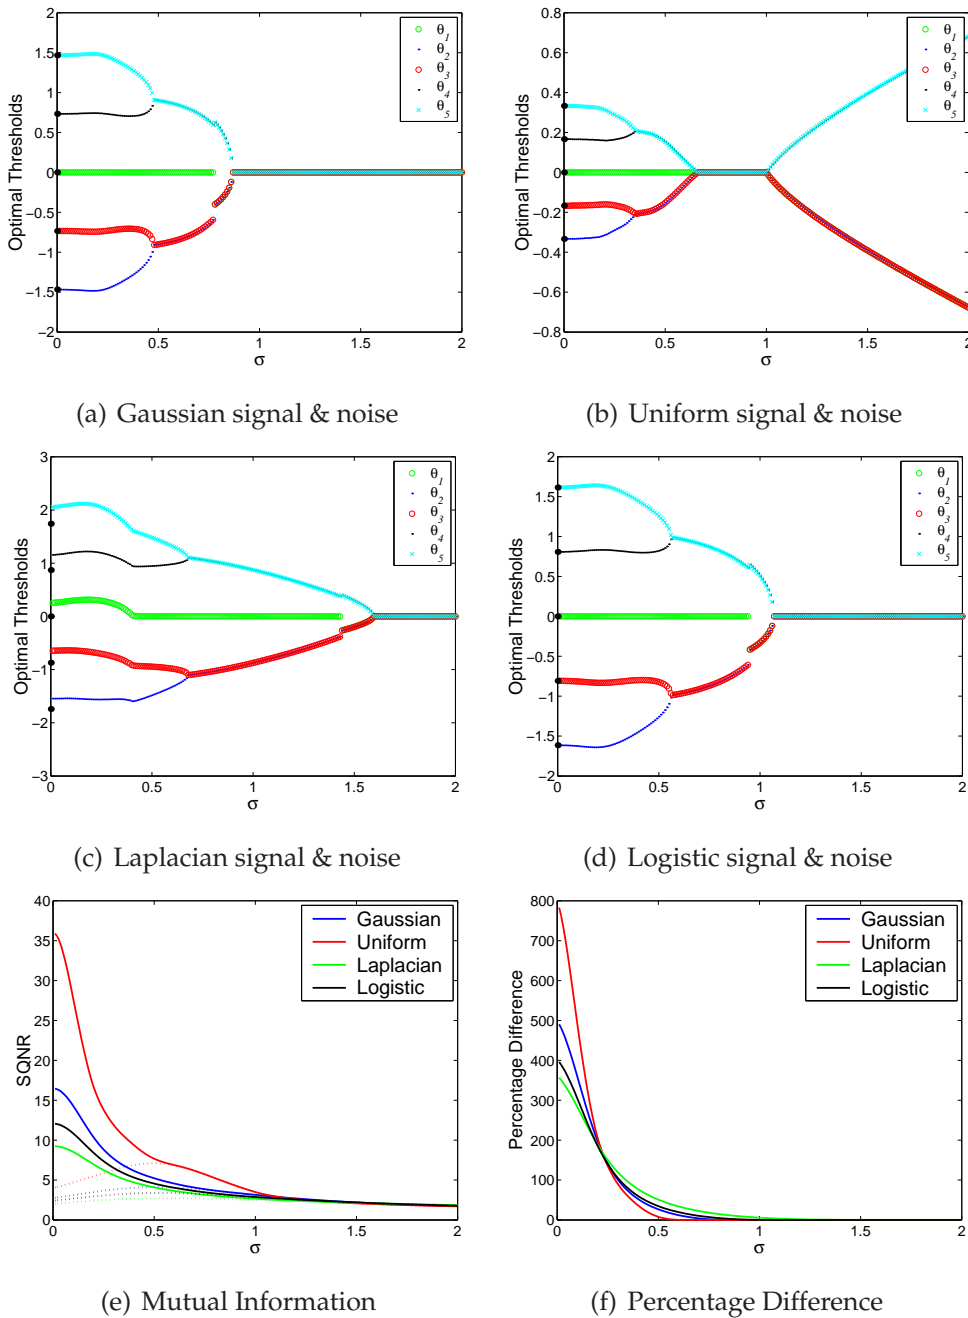


Figure 8.7. Optimal thresholds for linear decoding MSE distortion, $N = 5$. Figs. 8.7(a), 8.7(b), 8.7(c) and 8.7(d) show plots of the optimal thresholds for four different matched signal and noise pairs against increasing increasing noise intensity, σ , and the objective of minimised linear decoding MSE distortion. The optimal noiseless thresholds, as calculated by the Lloyd Method I algorithm, are shown with large black dots. Fig. 8.7(e) shows the SQNR obtained with these optimal threshold settings against increasing σ , as well as the linear decoding SQNR for SSR, which is shown with dotted lines. Fig. 8.7(f) shows the percentage difference between the MSE distortion obtained by optimally setting the thresholds, and that for SSR, against increasing σ . For σ greater than some critical value, σ_c , the SSR situation can be seen to be optimal.

8.5 Optimal Quantisation for MSE Distortion

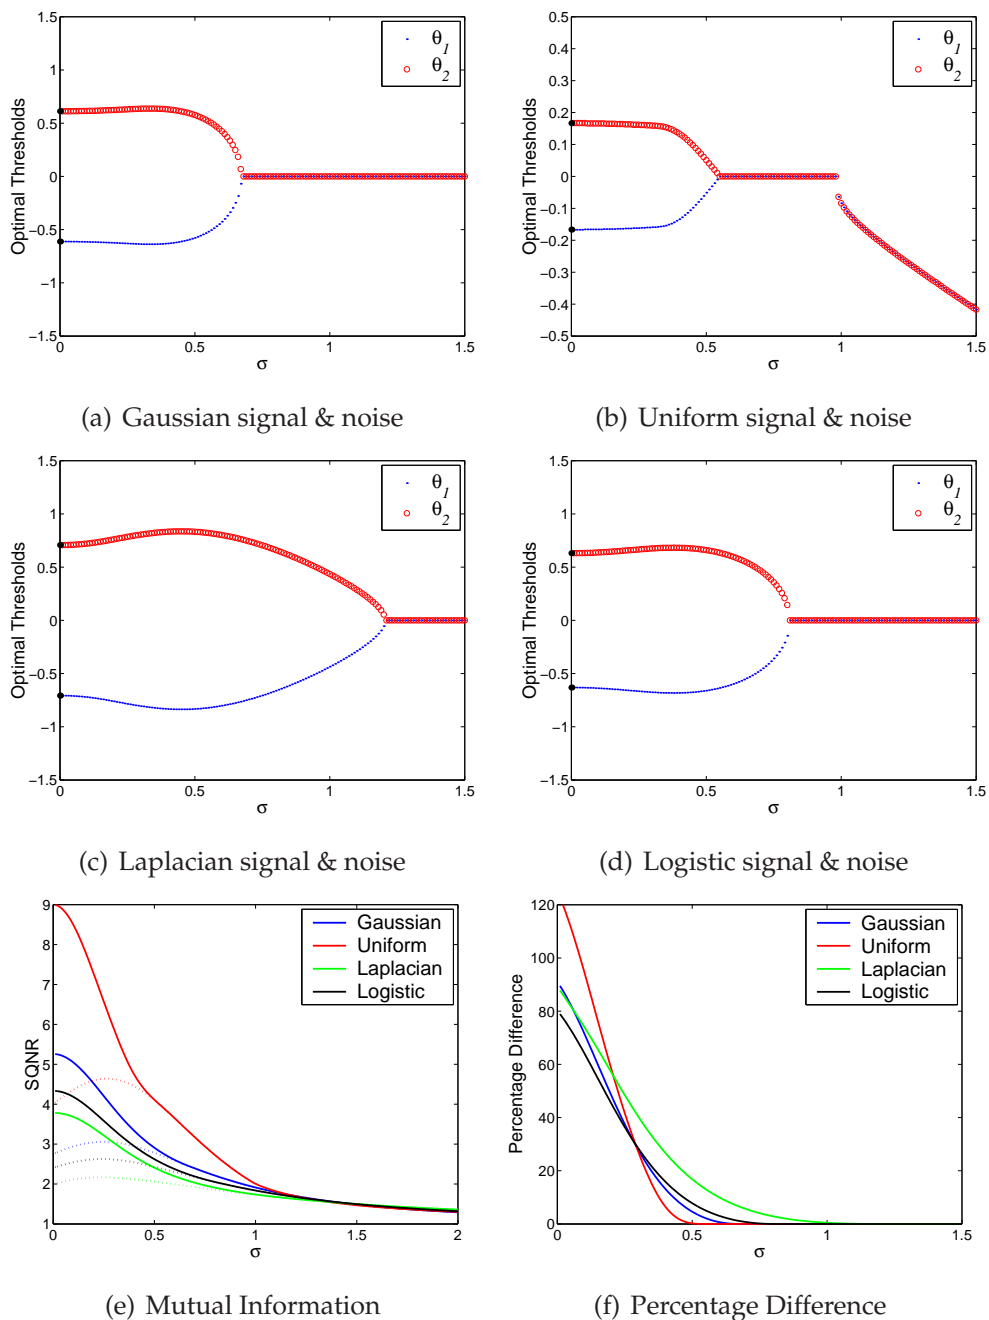


Figure 8.8. Optimal thresholds for MMSE distortion, $N = 2$. Figs. 8.8(a), 8.8(b), 8.8(c) and 8.8(d) show plots of the optimal thresholds for four different matched signal and noise pairs against increasing noise intensity, σ , and the objective of minimised MMSE distortion. The optimal noiseless thresholds, as calculated by the Lloyd Method I algorithm, are shown with large black dots. Fig. 8.8(e) shows the SQNR obtained with these optimal threshold settings against increasing σ , as well as the SQNR for SSR, which is shown with dotted lines. Fig. 8.8(f) shows the percentage difference between the MMSE distortion obtained by optimally setting the thresholds, and that for SSR, against increasing σ . For σ greater than some critical value, σ_c , the SSR situation can be seen to be optimal.

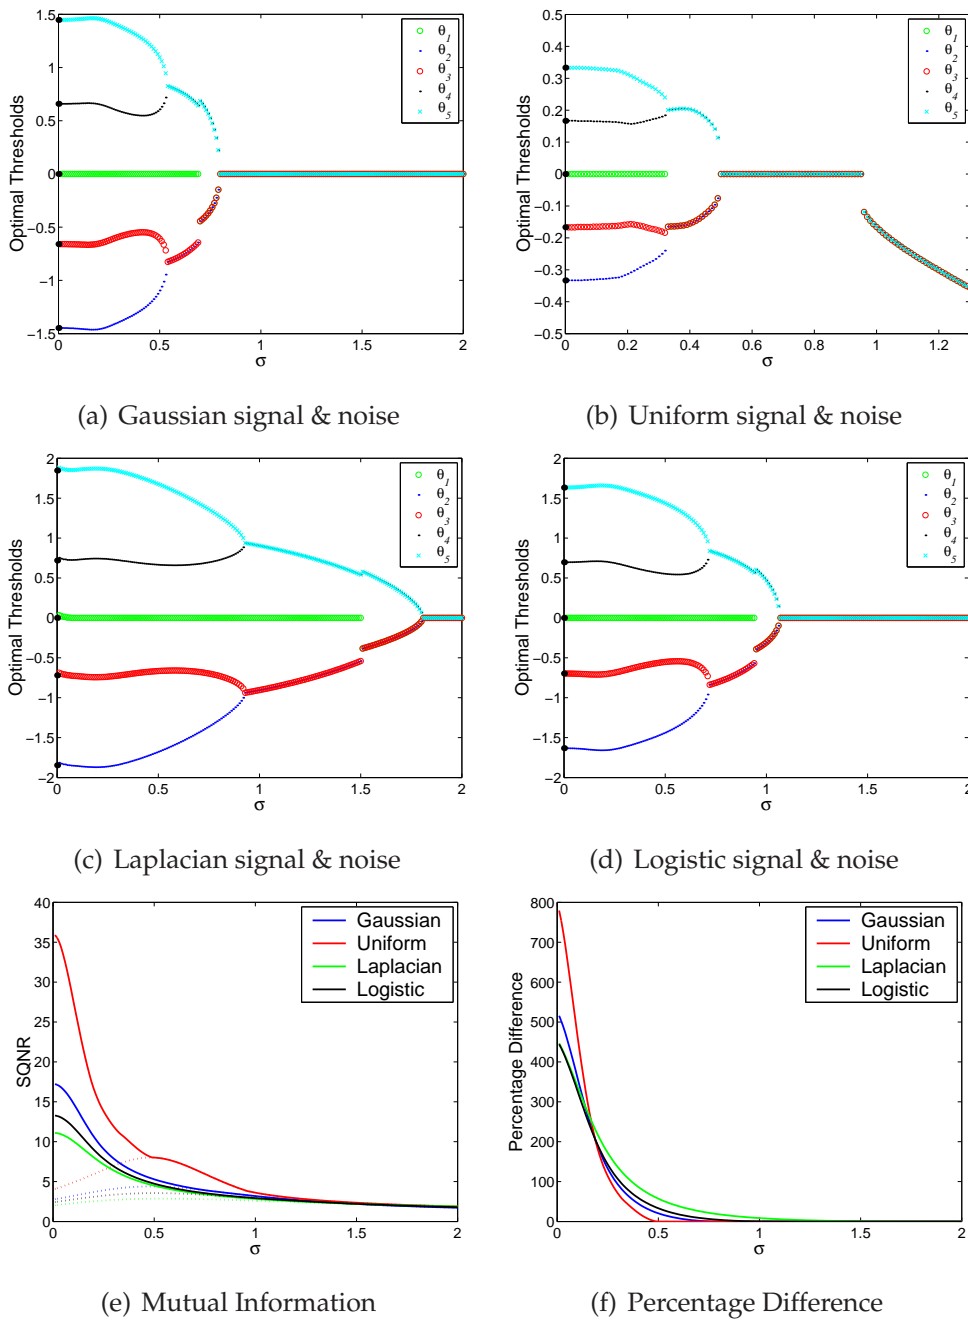


Figure 8.9. Optimal thresholds for MMSE distortion, $N = 5$. Figs. 8.9(a), 8.9(b), 8.9(c) and 8.9(d) show plots of the optimal thresholds for four different matched signal and noise pairs against increasing noise intensity, σ , and the objective of minimised MMSE distortion. The optimal noiseless thresholds, as calculated by the Lloyd Method I algorithm, are shown with large black dots. Fig. 8.9(e) shows the SQNR obtained with these optimal threshold settings against increasing σ , as well as the SQNR for SSR, which is shown with dotted lines. Fig. 8.9(f) shows the percentage difference between the MMSE distortion obtained by optimally setting the thresholds, and that for SSR, against increasing σ . For σ greater than some critical value, σ_c , the SSR situation can be seen to be optimal.

8.5 Optimal Quantisation for MSE Distortion

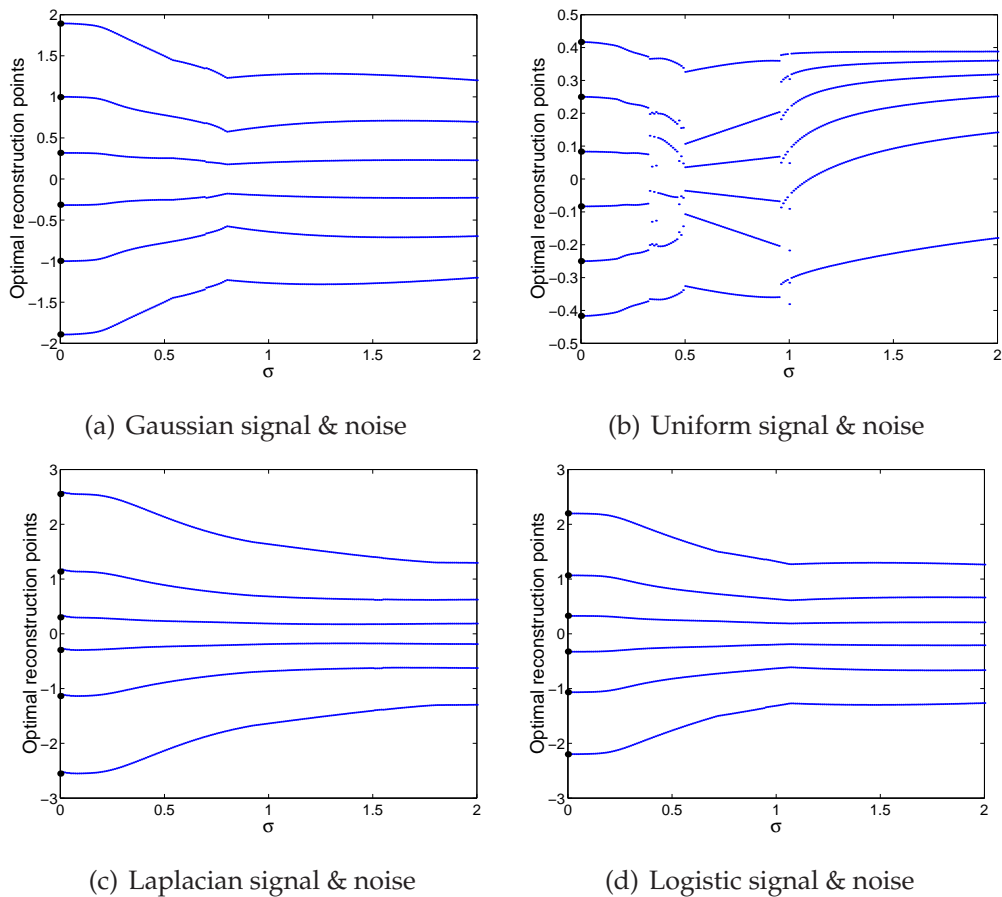


Figure 8.10. Optimal reconstruction points for minimised MMSE, $N = 5$. Figs. 8.10(a), 8.10(b), 8.10(c) and 8.10(d) show plots of the optimal reconstruction points for four different matched signal and noise pairs against increasing noise intensity, σ . The optimal noiseless reconstruction points, as calculated after applying the Lloyd Method I algorithm from Eqn. (8.40) are shown with large black dots.

8.6 Discussion of Results

8.6.1 Observations

Firstly, we see from the figures above that the mutual information, or MSE distortion, obtained with the optimal thresholds is strictly monotonic with increasing σ . This means that no SR effect is seen for optimised thresholds. For small σ , optimising the thresholds gives a substantial increase in performance. However, as σ increases, the difference between the SSR case and optimised thresholds decreases, until for sufficiently large σ , SSR is optimal. While the fact that the performance decreases with increasing σ means that there is no advantage to be gained by increasing the noise level in such an optimised system, the fact that SSR is optimal for sufficiently large noise intensity is still highly significant. Unlike a single threshold device, where the optimal threshold value is always at the signal mean regardless of the noise level and therefore SR can never occur, this result shows that when the noise is large, changing the thresholds from a situation where SR can occur gains no advantage.

Secondly, inspection of all the figures showing optimal threshold values against σ in the previous two Sections indicates several common features of the optimal threshold configuration. Firstly, for very small noise, the optimal thresholds and reconstruction points are consistent with the optimal noiseless values. There does not appear to be a discontinuity in the optimal thresholds as the noise intensity increases from zero to some small nonzero value.

The most striking feature is the fact that *bifurcations* are present. Consider firstly the simplest case of maximised mutual information and $N = 2$. For each signal and noise pair, for σ between zero and some critical value greater than zero, σ_c , the optimal placement of the two thresholds are at $\pm A$, where $A > 0$. However, for $\sigma > \sigma_c$, the optimal thresholds both have the same value. Apart from the uniform case, this value is always the signal mean of zero, which is simply the SSR situation. For the uniform case, we see that for $\sigma \in [\sigma_c, 1]$, we have the SSR situation. However for $\sigma > 1$, both thresholds remain identical, but via a discontinuous bifurcation this identical value is no longer zero.

The behaviour for small σ and large σ seen in the $N = 2$ case persists for larger N . That is, for sufficiently small σ , the optimal threshold values are all unique, and for sufficiently large σ —with the exception of the case of uniform signal and noise, $N =$

8.6 Discussion of Results

5 and minimised linear decoding MSE distortion—the optimal threshold values are identical for each threshold.

Most importantly, apart from the special situation of the uniform case, it is evident in all cases that above a certain value of σ the SSR situation is optimal. That is, the optimal quantisation for large noise is to set all thresholds to the signal mean.

For $N > 2$, we see that there are also regions of σ where some fraction of the optimal thresholds tend to cluster to particular identical values. We will refer to the number of thresholds with the same value as the *size* of a cluster, and the actual threshold value of those thresholds as the *value* of a cluster.

This tendency of the optimal thresholds to form clusters at identical values leads to regions of asymmetry about the x -axis, since if, for example, $N = 5$ and there are two clusters of size three and two, then the value of the cluster of size two is larger in magnitude than the value of the cluster of size three. Note that in such regions of asymmetry, there are two globally optimal threshold vectors. The second global solution is simply the negative of the set of thresholds in the first global solution, that is, $f(\boldsymbol{\theta}^*) = f(-\boldsymbol{\theta}^*)$. This result stems from the fact that both the signal and noise PDFs are even functions.

We can also see that it is quite common for bifurcations to occur, so that the number of clusters suddenly decreases with increasing σ . Sometimes a continuous bifurcation occurs as more than one cluster converges to the same value, as σ increases, to form a larger, merged, cluster. On other occasions a discontinuous bifurcation occurs, and two clusters with completely different values merge to form a larger cluster with a value somewhere between the two values of the two merging clusters. It does not appear possible for the number of clusters to increase with increasing σ , other than, again, for the case of uniform signal and noise, and minimised linear decoding MSE distortion.

However, further bifurcations can occur within a region of σ with k clusters, where the order of the size of the clusters changes with respect to the values of those clusters. For example, for Gaussian signal and noise and $N = 5$, Fig. 8.4(a) shows that there are three distinct clusters at $\sigma = 0.42$, and three at $\sigma = 0.45$. However there is a bifurcation between $\sigma = 0.42$ and $\sigma = 0.43$, since for $\sigma = 0.42$ the optimal solution is to have clusters of size 2, 2, 1, in that order from smallest cluster values to largest, while for $\sigma = 0.43$, the optimal solution is to have clusters of size 2, 1, 2.

Fig. 8.11 further illustrates the behaviour of the clusters for this case, for six different values of σ . Each sub-panel shows the size of each cluster on the y -axis, and the value of each cluster on the x -axis. For noise greater than the final bifurcation point, we have the SSR region occurring, that is, the optimal solution is for all thresholds to be equal to the signal mean of zero.




Figure 8.11. Optimal mutual information threshold clusters for various σ , $N = 5$. Panel plot showing the number of thresholds in each cluster for various values of σ in the case of Gaussian signal and noise and $N = 5$, and maximised mutual information.

The bifurcational structure is quite surprising, but appears to be fundamental to the problem type, since the same qualitative behaviour occurs whether we are maximising the mutual information, or minimising the MSE distortion. In Chapter 9, we will see the same behaviour for constrained mutual information maximisation. We can also see that the pattern is qualitatively the same for each signal and noise pair considered. Furthermore, numerical experiments find very similar patterns appear for mixed signal and noise distributions, such as a uniform signal subject to Gaussian noise.

However, there are some anomalies that make it difficult to generalise too much. For example consider the case of logistic signal and noise, and $N = 3$ shown in Fig. 8.2(d). In this case, there is only one bifurcation, since at some value of σ between 0.68 and

0.69, the optimal threshold solution changes from all three thresholds being unique, to all three thresholds being zero, without an intermediate region with a cluster of size two, and a cluster of size one. Furthermore, while in the Gaussian case of $N = 5$ there are five bifurcations, one of which occurs within a region of three clusters, in the logistic case of $N = 5$, there are only three bifurcations, since there is no region where there are four clusters.

A clue to the reason for the bifurcations comes from inspection of the plots of the optimal reconstruction points corresponding to the optimal thresholds. Despite the bifurcational structure in the optimal thresholds, in all cases but the uniform case, the optimal reconstruction points appear to change very smoothly with σ . This indicates that at values of σ where discontinuous bifurcations occur in the optimal thresholds, that more than one local optimum gives the same optimal reconstruction points.

8.6.2 Mathematical Description of Optimal Thresholds

Quantiser point density function

We now mathematically describe the observations made above. For the purposes of optimisation, the ordering of the optimal threshold vector, θ^* , is not important. However, to simplify the mathematical description, we now introduce an ordered sequence notation for the optimal thresholds. Specifically, we label the i -th optimal threshold value as θ_i^* , so that the sequence $(\theta_i^*)_{i=1}^N$ is non-decreasing. As we saw in Eqn. (8.27), in the absence of noise, it is straightforward to show for the goal of maximum mutual information that each optimal threshold is given by $\theta_i^* = F_x^{-1}\left(\frac{i}{N+1}\right)$.

We now introduce a concept used in the theoretical analysis of high resolution quantisers in information theory—that of a quantiser *point density function*, $\lambda(x)$, defined over the same support variable as the source PDF, $P(x)$ (Gray and Neuhoff 1998). The point density function has the property that $\int_x \lambda(x)dx = 1$, and usually is only used in the context where the number of thresholds is very large. In this situation, the point density function gives the density of thresholds across the support of the signal PDF.

For any given N , and some nonzero value of σ , we observe from the plots of optimal thresholds that our empirically optimal threshold sequence, $(\theta_i^*)_{i=1}^N$, can have at most $k(\sigma)$ unique values, where $1 \leq k \leq N$. When bifurcations occur as σ increases, $k(\sigma)$ may either decrease or—in the situation where the ordering of clusters change—remain constant.

We now denote $v(j, \sigma)$ as the fraction of the total thresholds in the j -th cluster, at noise intensity σ , where $j \in \{1, \dots, k(\sigma)\}$, so that $\sum_{j=1}^{k(\sigma)} v(j, \sigma) = 1$. Thus, $v(j, \sigma)$ is the size of the j -th cluster divided by N .

Denote the value of the j -th cluster as Θ_j , so that the size of the cluster at $x = \Theta_j$ is $Nv(j, \sigma)$. As with the ordered optimal threshold sequence, we can define an ordered sequence of cluster values as $(\Theta_j)_{j=1}^{k(\sigma)}$. Unlike the optimal threshold sequence, this sequence is strictly increasing.

We are now able to write a point density function as a function of σ to describe our empirically optimal threshold configuration. This is

$$\lambda(x, \sigma) = \sum_{j=1}^{k(\sigma)} v(j, \sigma) \delta(x - \Theta_j), \quad (8.42)$$

where $\delta(\cdot)$ is the delta function. We also note that $\int_{x=-\infty}^a \lambda(x, \sigma) dx$ is the fraction of thresholds with values less than or equal to a , and that $\int_{x=-\infty}^{\infty} \lambda(x, \sigma) dx = 1$.

For the special case of $\sigma = 0$ and maximised mutual information, we can use Eqn. (8.27) to write the analytically optimal point density function as

$$\lambda(x, 0) = \sum_{j=1}^N v(j, 0) \delta(x - \Theta_j) = \sum_{i=1}^N \frac{1}{N} \delta \left(x - F_x^{-1} \left(\frac{i}{N+1} \right) \right). \quad (8.43)$$

As an example for nonzero σ , consider the case of $N = 5$ and Gaussian signal and noise shown for maximised mutual information and various σ in Fig. 8.11. The value of k for each σ and the corresponding cluster sizes, $Nv(j, \sigma)$, and the approximate cluster values, $(\Theta_j)_{j=1}^{k(\sigma)}$, are shown in Table 8.1.

So far, we have point density functions consisting only of singularities. In high resolution quantisation theory, point density functions are generally continuous functions, analogous to PDFs. Due to the small N considered here, our point density functions are analogous to discrete probability mass functions, rather than PDFs. Discussion of the behaviour of this description of the optimal thresholds for large N is left for future work.

Conditional output moments

Using the notation introduced above, we are able to rewrite our previous expressions for the conditional mean and variance of the output encoding, y , for arbitrary thresholds.

8.6 Discussion of Results

Table 8.1. Parameters for point density function, $N = 5$, Gaussian signal and noise, and maximised mutual information.

σ	$k(\sigma)$	$\{Nv(j, \sigma)\}$	$(\Theta_j)_{j=1}^{k(\sigma)}$
0.3	5	$\{1, 1, 1, 1, 1\}$	$(-1.0410, -0.4164, 0.0, 0.4164, 1.0410)$
0.39	4	$\{1, 1, 2, 1\}$	$(-0.8075, -0.5967, 0.2288, 0.9489)$
0.42	3	$\{2, 2, 1\}$	$(-0.6783, 0.2549, 0.8770)$
0.45	3	$\{1, 2, 1\}$	$(-0.6319, 0.0, 0.6319)$
0.55	2	$\{3, 2\}$	$(-0.3592, 0.5513)$
0.7	1	$\{5\}$	(0.0)

From Eqn. (8.16), the expected value of y given x is

$$E[y|x] = \sum_{i=1}^N P_{1|x,i} = N \sum_{j=1}^k v_j P_{1|x,j} = N \sum_{j=1}^k v_j F_R(x - \Theta_j). \quad (8.44)$$

and from Eqn. (8.17), the variance of y given x is

$$\begin{aligned} \text{var}[y|x] &= \sum_{i=1}^N P_{1|x,i}(1 - P_{1|x,i}) = N \sum_{j=1}^k v_j P_{1|x,j}(1 - P_{1|x,j}) \\ &= N \sum_{j=1}^k v_j F_R(x - \Theta_j)(1 - F_R(x - \Theta_j)). \end{aligned} \quad (8.45)$$

We will use these results in Section 8.6.4.

8.6.3 Local Maxima

Partitions of integers

A *partition* of a positive integer is a way of writing that integer as a sum of smaller positive integers. The order of the integers in the sum is not considered, and conventionally a partition is written in the order of largest to smallest integers in the sum. For example, there are 7 partitions of the integer 5. These are $\{5\}$, $\{4, 1\}$, $\{3, 2\}$, $\{3, 1, 1\}$, $\{2, 2, 1\}$, $\{2, 1, 1, 1\}$, and $\{1, 1, 1, 1, 1\}$. The theory of partitions of integers is a rich area of number theory, and was of interest to such well known number theorists as Ramanujan, Hardy and Littlewood (Andrews 1976). The number of partitions of the integer N increases very rapidly with N .

If ordering is taken into account, then extra partitions are possible, namely $\{1, 3, 1\}$, $\{2, 1, 2\}$, and $\{1, 2, 1, 1\}$. The reverse of all the ten partitions listed are also feasible. Hence, there are 20 possible ordered partitions of the integer 5.

Description of local maxima

Explaining the presence of discontinuous bifurcations in the optimal threshold figures at first seems very difficult. However, most of the discontinuous bifurcations are actually due to the presence of many locally optimal threshold configurations. In fact, numerical experiments find that for every value of σ , there is at least one locally optimal solution—that is a set of threshold values giving a gradient vector of zero—corresponding to every possible *partition* of N . For each partition¹⁵, there are as many locally optimal solutions as there are unique orderings of that partition. For small σ , all of these local optima are unique. As σ increases, more and more of these local optima bifurcate continuously to be equal to other local optima. For example, a local optimum corresponding to $k = 3$ clusters, with $\{Nv(j, \sigma)\} = \{2, 2, 1\}$ might have Θ_2 and Θ_3 converge to the same value with increasing σ . At the point of this convergence, a bifurcation occurs, and the local optimum becomes one consisting of $k = 2$ clusters, with $\{Nv(j, \sigma)\} = \{2, 3\}$.

Again, the exception to this rule of thumb seems to be the uniform case, in which bifurcations can occur discontinuously in locally optimal solutions.

To illustrate this effect, in the simplest possible case of $N = 3$, Fig. 8.12 shows the optimal thresholds for all three locally optimal solutions for maximum mutual information, for all four matched signal and noise cases. Notice that apart from the uniform case when $\sigma > 1$, the SSR situation is always a local optimum. Fig. 8.13 shows the mutual information achieved by each of the three locally optimal solutions. Notice that for small σ , there is a significant difference between the mutual information in each case. However, as σ increases, the difference between each case decreases.

Notice in Fig. 8.12 that both the $\{1, 1, 1\}$ and $\{2, 1\}$ cases change smoothly with increasing σ until they both converge to the SSR case for sufficiently large noise. This is in contrast to the globally optimal thresholds for Gaussian signal and noise shown in Fig. 8.2(a), where there is a discontinuous bifurcation. This is due to the optimal solution switching from being the $\{1, 1, 1\}$ case to the $\{2, 1\}$ case. Fig. 8.14 shows the difference in mutual information between these two cases, in the region of σ where the bifurcation occurs. It is clear that the difference between each case is very small, and therefore the fact that the optimal thresholds correspond to the $\{2, 1\}$ case rather than

¹⁵Here, however, due to considering only PDFs that are even functions, there will be some symmetry such that we can ignore all partitions that are the reverse order, and in the example of the integer 5, we only need consider 10 possible partitions.

8.6 Discussion of Results

the $\{1, 1, 1\}$ case is really only of academic interest. The more important fact is that, for sufficiently large σ , having all thresholds identical is optimal—since this implies a reduction in complexity for the specification of the optimal quantiser. This is because only one threshold value is required, rather than N , which is the case for no noise.

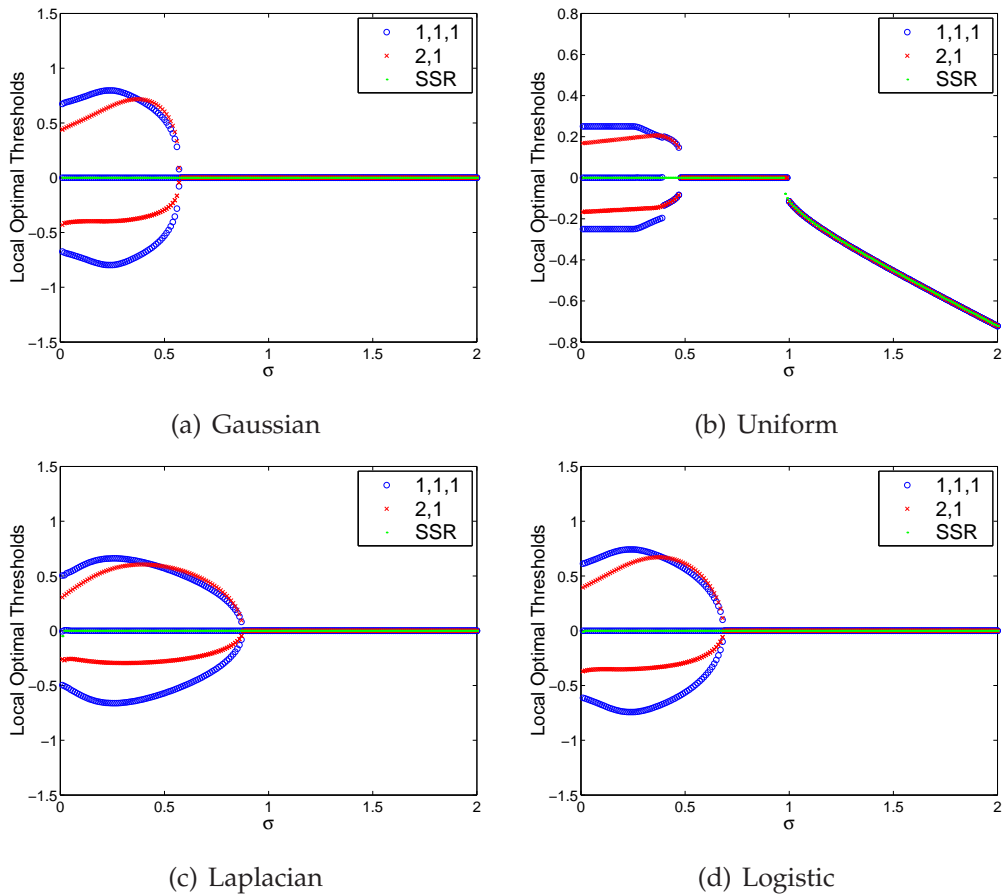


Figure 8.12. Locally optimal thresholds for maximum mutual information and $N = 3$. This figure shows the three locally optimal solutions for $N = 3$ and maximised mutual information. Apart from the uniform case when $\sigma > 1$, the SSR situation is always a local optimum. For the uniform case when $\sigma > 1$, although SSR is no longer a local optimum, the optimum is for all thresholds to have the same value.

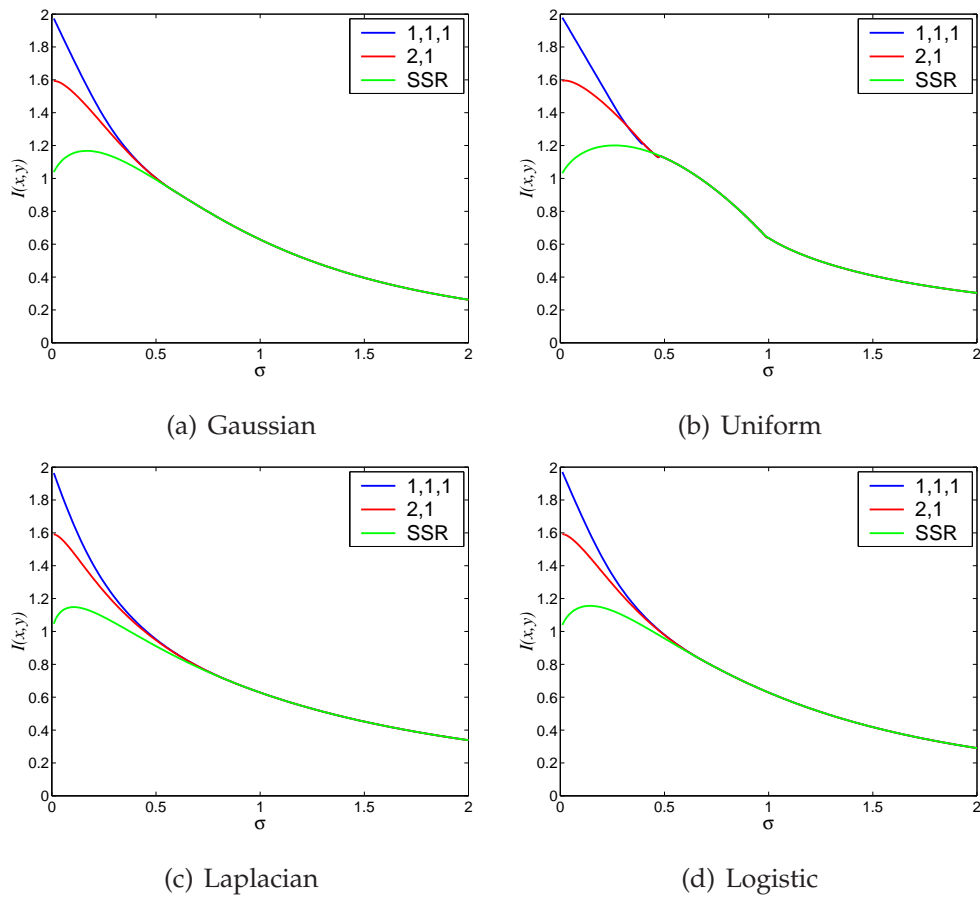


Figure 8.13. Locally optimal mutual information, $N = 3$. This figure shows the mutual information for the three locally optimal solutions for $N = 3$. It is clear that for small σ , all thresholds unique gives a much larger mutual information than the other cases. However, for sufficiently large σ , the differences in mutual information become smaller and smaller, until all three local solutions converge to the same solution corresponding to SSR.




Figure 8.14. Difference in mutual information between local optima. This plot shows the difference in mutual information between the locally optimal solutions with $v = \{1, 1, 1\}$ and $v = \{2, 1\}$ for Gaussian signal and noise and $N = 3$. Clearly, $v = \{2, 1\}$ is larger than $v = \{1, 1, 1\}$ only for $\sigma \in [0.51, 0.57]$, and is only larger by an amount of the order of 0.5×10^{-4} bits per sample, which is very small. However, this changeover from $v = \{1, 1, 1\}$ being optimal to $v = \{2, 1\}$ being optimal, is the reason that a *discontinuous* bifurcation appears in the *globally* optimal solution shown in Fig. 8.2(a). For $\sigma \geq 0.57$, the SSR situation is optimal, and both the $v = \{1, 1, 1\}$ and $v = \{2, 1\}$ situation bifurcate at this point to the SSR case.

8.6.4 An Estimation Perspective

Fisher information

Recall from Chapter 6 that the Fisher information for the SSR model can be calculated and used in a formula giving a lower bound on the MSE distortion. This result can also be applied to the arbitrary threshold model considered in this Chapter.

The output of the array of threshold devices, y , provides a biased estimate of the input, x . Recall from Chapter 6 that the *information bound* states that

$$\text{var}[y|x] \geq \frac{\left(\frac{d}{dx}E[y|x]\right)^2}{J(x)}, \quad (8.46)$$

where $J(x)$ is the Fisher information, which for arbitrary thresholds is given by

$$J(x) = \sum_{n=0}^N \frac{\left(\frac{dP(n|x)}{dx}\right)^2}{P(n|x)}. \quad (8.47)$$

Now, from Eqn. (8.16), we have for the arbitrary threshold model that

$$\begin{aligned} \frac{d}{dx}E[y|x] &= \sum_{i=1}^N \frac{dP_{1|x,i}}{dx} \\ &= \sum_{i=1}^N \frac{dF_R(x - \theta_i)}{dx} \\ &= \sum_{i=1}^N R(x - \theta_i) \\ &= N \sum_{j=1}^k v_j R(x - \Theta_j). \end{aligned} \quad (8.48)$$

Substituting Eqns. (8.48) and (8.45) into Inequality (8.46) and rearranging gives an inequality for the Fisher information as

$$\begin{aligned} J(x) &\geq \frac{\left(\frac{d}{dx}E[y|x]\right)^2}{\text{var}[y|x]} \\ &= \frac{\left(\sum_{i=1}^N R(x - \theta_i)\right)^2}{\sum_{i=1}^N P_{1|x,i}(1 - P_{1|x,i})} \\ &= \frac{N \left(\sum_{j=1}^k v_j R(x - \Theta_j)\right)^2}{\sum_{j=1}^k v_j P_{1|x,j}(1 - P_{1|x,j})}. \end{aligned} \quad (8.49)$$

8.6 Discussion of Results

For the special case of SSR, we have $k = 1$, $v_k = 1$ and $\Theta_j = 0$, which when substituted into Inequality (8.49) gives a RHS that is exactly the Fisher information for SSR expressed by Eqn. (6.119) in Section 6.7.2 of Chapter 6. The Fisher information for SSR is derived from first principles in Section D.6 of Appendix D.

Thus, for SSR, Inequality (8.49) becomes an equality and, as discussed in Chapter 6, the SSR Fisher information meets the information bound with equality, in the case of no decoding. However, for arbitrary thresholds, a derivation of the Fisher information from first principles using Eqn. (8.47) is not a trivial task—since we do not have an analytic expression for $P(n|x)$ —and therefore we do not know whether the information bound is also met with equality in the general case.

However, the Fisher information for arbitrary thresholds and signal and noise distributions can be calculated numerically from Eqn. (8.47), and compared with numerical calculations of the RHS of Inequality (8.49). Experiments with such calculations indicate that the bound does not hold exactly, apart from the SSR situation, but that, even for small N , the bound has a maximum error when compared to the exact Fisher information, in the order of one percent. It is possible this error is attributable to numerical errors, and that the bound does hold with equality. In any case, we are able to state

$$J(x) \simeq \frac{\left(\sum_{i=1}^N R(x - \theta_i)\right)^2}{\sum_{i=1}^N P_{1|x,i}(1 - P_{1|x,i})} = \frac{N \left(\sum_{j=1}^k v_j R(x - \Theta_j)\right)^2}{\sum_{j=1}^k v_j P_{1|x,j}(1 - P_{1|x,j})}. \quad (8.50)$$

Future work may be able to justify making Inequality (8.49) a strict equality under certain conditions, using Brunel and Nadal (1998) as a starting reference.

Average Information Bound

From inspection of Eqn. (8.47), the Fisher information for x is unchanged if the output y is decoded to \hat{y} , since $J(x)$ depends only on $P(n|x)$. However, the bias changes for a decoding, as does the conditional variance. Thus, the information bound for the decoding, \hat{y} , is

$$\text{var}[\hat{y}|x] \geq \frac{\left(\frac{d}{dx} \mathbb{E}[\hat{y}|x]\right)^2}{J(x)}. \quad (8.51)$$

Substituting Eqn. (8.50) into Inequality (8.51) gives

$$\begin{aligned} \text{var}[\hat{y}|x] &\geq \frac{\sum_{i=1}^N P_{1|x,i}(1 - P_{1|x,i}) \left(\frac{d}{dx} E[\hat{y}|x]\right)^2}{\left(\sum_{i=1}^N R(x - \theta_i)\right)^2} \\ &= \frac{\left(\sum_{i=1}^N P_{1|x,i}(1 - P_{1|x,i})\right) \left(\sum_{n=0}^N \hat{y}_n \frac{d}{dx} P(n|x)\right)^2}{\left(\sum_{i=1}^N R(x - \theta_i)\right)^2}. \end{aligned} \quad (8.52)$$

Thus, a lower bound on the conditional MSE distortion is

$$D(x) \geq \frac{\left(\sum_{i=1}^N P_{1|x,i}(1 - P_{1|x,i})\right) \left(\sum_{n=0}^N \hat{x}_n \frac{d}{dx} P(n|x)\right)^2}{\left(\sum_{i=1}^N R(x - \theta_i)\right)^2} + b_{\hat{x}}(x)^2, \quad (8.53)$$

where the decoding is the optimal decoding, $\hat{x}_n = E[x|n]$, and $b_{\hat{x}}(x) = E[\hat{x}|x] - x$ is the bias of the decoding.

Multiplying the RHS of Inequality (8.53) by $P(x)$ and then integrating over all x gives the Average Information Bound (AIB) introduced in Chapter 6. The AIB is a lower bound on the MMSE distortion, and is therefore

$$\text{AIB} = \int_x P(x) \left(\frac{\left(\sum_{i=1}^N P_{1|x,i}(1 - P_{1|x,i})\right) \left(\sum_{n=0}^N \hat{x}_n \frac{d}{dx} P(n|x)\right)^2}{\left(\sum_{i=1}^N R(x - \theta_i)\right)^2} \right) dx + E[b_{\hat{x}}(x)^2], \quad (8.54)$$

where $E[b_{\hat{x}}(x)^2]$ is the mean square bias.

As carried out in Chapter 6, it is instructive to numerically calculate the average information bound, and its two components, the mean square bias, and the average error variance.

Fig. 8.15 shows the AIB, and its components, the mean square bias, and the average error variance, for the case of Gaussian signal and noise, and $N = 3$, for all three locally optimal solutions when the thresholds are optimised to minimise the MMSE distortion. Fig. 8.15(b) shows that the AIB is less than two percent different from the actual optimal MMSE distortion in each case. As we saw earlier, it is clear that optimising the thresholds results in a decrease in the MMSE distortion when compared to SSR. Figs. 8.15(c) and 8.15(d) show that this decrease is a result of optimising the thresholds to offset a small increase in the average conditional variance against a larger decrease in the mean square bias. Thus, optimising the thresholds means optimising the trade-off between bias and variance. The end result of this is a MMSE distortion and AIB that are strictly increasing with increasing σ .

8.6 Discussion of Results

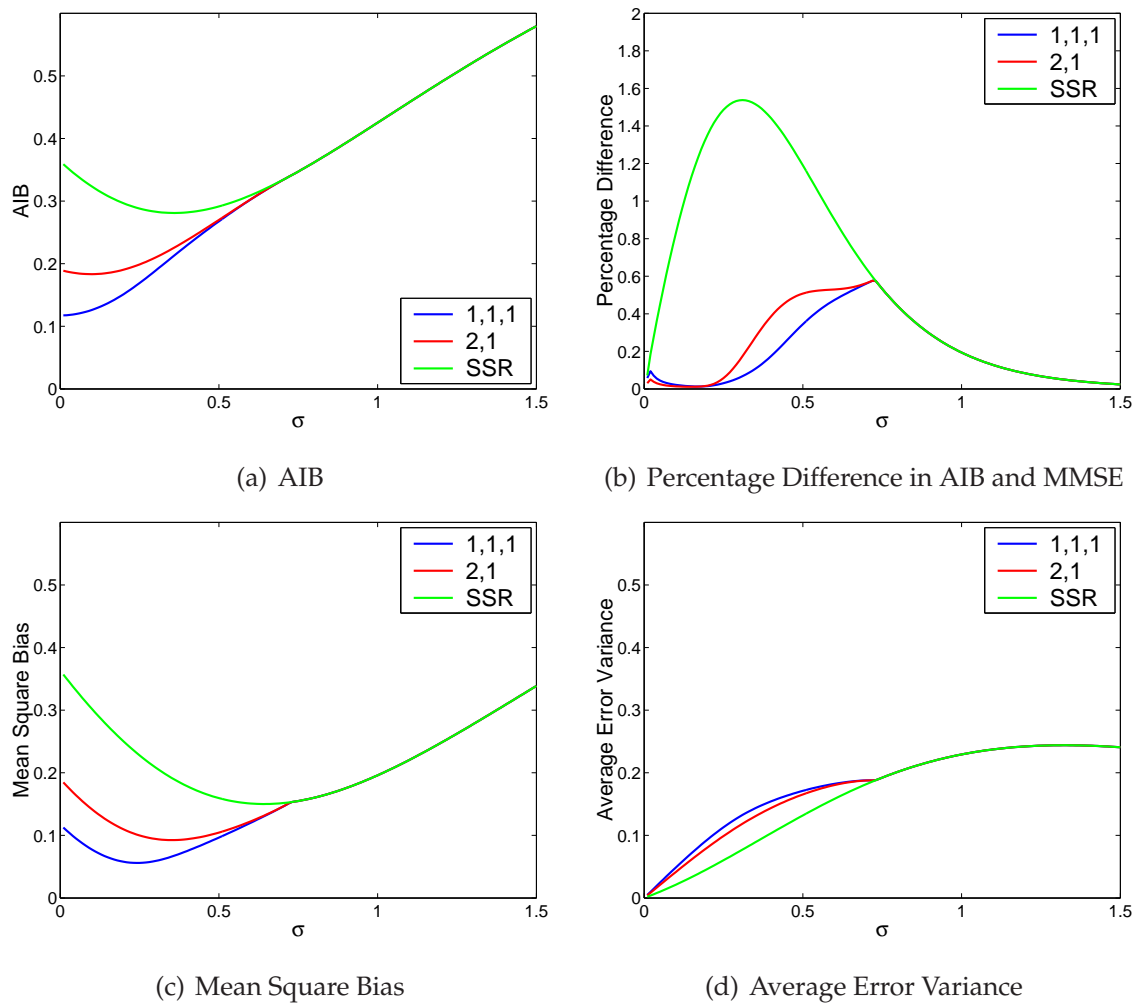


Figure 8.15. Average Information Bound (AIB), Gaussian signal and noise, $N = 3$.

Fig. 8.15(a) shows the AIB of Eqn. (8.54) plotted against increasing σ for the three different local optima. Fig. 8.15(b) shows that the percentage difference between the AIB and the MMSE distortion for each local optimum is very small, indicating that the AIB gives a bound that is very close to the actual MMSE distortion. Figs. 8.15(c) and 8.15(d) show the two components of the AIB, the mean square bias, and the average error variance. Notice how for small σ , the mean square bias for the optimal threshold situation of $\{1,1,1\}$ is far smaller than for SSR, while the average error variance is a small amount larger. This illustrates how optimising the thresholds for small σ optimises the tradeoff between bias and variance. For sufficiently large σ , the optimal tradeoff is provided by the SSR situation.

8.7 Locating the Final Bifurcation

We have seen that for sufficiently large σ , the SSR situation of all thresholds equal to the signal mean becomes optimal. We have also seen that the largest value of σ , which we will call, σ_b , for which SSR is not optimal increases with increasing N . It appears also that in the region of σ just smaller than σ_b that there are two clusters. For even N , these each have size $N/2$, and values $\Theta_1 = -\Theta_2$. If we assume that this will always be the case, it becomes a straightforward task to numerically find the value of σ_b as a function of increasing N . It only requires to set $\Theta_1 = -\Theta_2 = \epsilon$, where ϵ is small, say 0.001, and for a given N to find the value of σ at which the SSR mutual information or MMSE distortion changes from being smaller or larger than the mutual information with two threshold clusters at $\pm\epsilon$.

The result of carrying this out is shown in Fig. 8.16 for mutual information, and Fig. 8.17 for MMSE distortion. The value of N , which we will refer to as N_b , plotted for each value of σ , is the smallest even-valued N for which the mutual information or MMSE distortion with $\Theta_1 = -\Theta_2 = 0.001$ gives better performance than the SSR situation. Thus, for each value of σ , if $N < N_b$, then SSR is optimal.

It is clear that as σ increases, N_b also increases, and increases very rapidly near some critical value of σ . It appears likely that for large N , N_b asymptotically converges towards some fixed value, σ_b^* . This means that for sufficiently large σ , that SSR is always optimal, regardless of the size of N . The value of σ_b^* does however, depend on the measure used, and the signal and noise distribution. Further results on such asymptotic large N behaviour for arbitrary thresholds is left for future work.

8.7 Locating the Final Bifurcation




Figure 8.16. Final bifurcation point, mutual information. This figure shows the smallest even-valued N for which the mutual information with $\Theta_1 = -\Theta_2 = 0.001$ gives better performance than the SSR situation. Thus, the value of σ corresponding to each N is the approximate final bifurcation point. Note that the Laplacian situation is only plotted with N up to 84. This is due to numerical calculations using the Laplacian PDF being less accurate than the Gaussian and logistic cases, due to the Laplacian PDF having a non-differentiable point at $x = 0$.




Figure 8.17. Final bifurcation point, MMSE distortion. This figure shows the smallest even-valued N for which the MMSE distortion with $\Theta_1 = -\Theta_2 = 0.001$ gives better performance than the SSR situation. Thus, the value of σ corresponding to each N is the approximate final bifurcation point.

8.8 Chapter Summary

The introductory Section of this chapter discusses the SSR case of all thresholds being identical, and sets the context for removing this restriction. When this restriction is removed, and all thresholds are free variables, it is of interest to optimally set the threshold values. In the presence of independent threshold noise, such a problem is an optimal stochastic quantisation problem. Section 8.1 therefore also briefly reviews the literature on optimal quantisation and points out that the problem addressed in this Chapter does not appear to have been previously studied.

Section 8.2 mathematically describes generalisation of the SSR model to arbitrary thresholds, outlines a method for recursively calculating the transition probabilities, and mathematically formulates the optimal stochastic quantisation problem we aim to solve. Section 8.3 then describes the solution method we use in solving these problems.

Results are presented in Sections 8.4 and 8.5, and discussed in Section 8.6. The key features of our results are that in general, for sufficiently large σ , the SSR situation is optimal, while for smaller σ , the optimal thresholds tend to cluster to identical values, with the number of clusters decreasing with increasing σ . We also introduce notation to describe these optimal thresholds, and briefly consider the information bound, and an associated lower bound on the MSE distortion. As N increases, the value of σ at which SSR becomes optimal also increases.

Finally, Section 8.7 shows that for sufficiently large σ , it appears that SSR is always optimal, regardless of the magnitude of N .

8.8.1 Original Contributions for Chapter 8

This chapter included the following original contributions:

- A statement of a relationship between the set of $\{P_{1|x,i}\}$ and the transition probabilities, via the moment generating function for the array of threshold devices.
- Derivation of a very general recursive formula, with $O(N^2)$ computational complexity, that makes it straightforward to numerically calculate the transition probabilities, $P(n|x)$, for any given threshold values, and noise distribution.
- Mathematical formulation of the optimal stochastic quantisation problem, for the array of threshold devices, in terms of the vector of optimal thresholds, θ .

- Presentation of numerical solutions to the stochastic optimal quantisation problems of maximising the mutual information, and minimising the MSE distortion. The optimal thresholds found are consistent with the optimal noiseless thresholds, in that for very small noise, the optimal thresholds are very close to the optimal noiseless ones.
- Discovery and discussion of the unexpected bifurcation pattern in the optimal stochastic quantisation results.
- Numerical validation that the SSR situation of all thresholds equal to the signal mean is in fact optimal for sufficiently large noise intensity, for a range of signal and noise distributions, and both the mutual information and MSE distortion measures.
- Derivation of an approximation to the Fisher information for arbitrary thresholds, and its application to a calculation of a lower bound on the MSE distortion. Optimally setting the thresholds is shown to be related to finding the optimal tradeoff between the two components of this lower bound, the mean square bias, and the average error variance.
- Numerical validation that SSR remains optimal for sufficiently large noise intensity, even if N becomes very large.

8.8.2 Further Work

Possible future work and open questions arising from this chapter might include:

- Discussion of other signal and noise distributions than those considered here, including mixed signal and noise distributions, deterministic signals, and distributions with one-sided PDFs such as the Rayleigh distribution.
- Mathematical proofs of the fact that SSR is optimal for large N , and further mathematical analysis of the bifurcational structure.
- Extension of the optimal quantisation problems considered here with simple on-off threshold device, to more realistic neural models, for example, the FitzHugh-Nagumo neuron model.

- Rigorous justification of the fact that Inequality (8.49) can be approximated as an equality under certain conditions. A good starting point for this research question is the material contained in Brunel and Nadal (1998).

This concludes Chapter 8, which studies the extension of the SSR model to a model with arbitrary thresholds. Chapter 9 now examines a further extension of the SSR model to incorporate constraints on energy and information.

Chapter 9

SSR, Neural Coding, and Performance Tradeoffs

ENGINEERED systems usually require finding the right tradeoff between cost and performance. Communications systems are no exception, and much theoretical work has been undertaken to find the limits of achievable performance for the transmission of information. For example, Shannon's celebrated channel capacity formula and coding theorems say that there is an upper limit on the average amount of information that can be transmitted in a channel for error free communication. This limit can be increased if the power of the signal is increased, or the bandwidth in the channel is increased. However, nothing comes for free, and increasing either power or bandwidth can be expensive; hence there is a tradeoff between cost and performance in such a communications system—performance (measured by bit rates) can be increased by increasing the cost (power or bandwidth). This chapter discusses several problems related to the tradeoff between cost and performance in the SSR model. We are interested both in the SSR model as a channel model, from an energy efficient neural coding point of view, as well as the lossy source coding model, where there is a tradeoff between rate and distortion.

9.1 Introduction

Chapter 8 introduces an extension to the Suprathreshold Stochastic Resonance (SSR) model by allowing all thresholds to vary independently, instead of all having the same value. This Chapter further extends the SSR model by introducing an energy constraint into the optimal stochastic quantisation problem. We also examine the tradeoff between rate and distortion, when the SSR model is considered as a stochastic quantiser.

Section 9.2 contains original work on the extension of the SSR model to energy constraints that has been published, in part, in McDonnell *et al.* (2004b). Section 9.3 contains original work on the tradeoff between rate and distortion in the SSR model that has been published, in part, in the open literature (McDonnell *et al.* 2005a, McDonnell *et al.* 2005c, McDonnell *et al.* 2005d).

Recall that the initial work on the SSR model (Stocks 2000a, Stocks and Mannella 2001) was partly motivated by its relevance to neural coding. The constraint we consider in this Chapter is also motivated by this fact. If, as discussed in Chapter 4, the issue is that of information transmission, we can say that ideally the encoding of sensory input by neurons should maximise the mutual information between input and output. However, as with most systems, there is usually some cost associated with maximising a quantity. For neural systems, a cost function that has received recent attention is that of energy efficiency. Hence, we consider the problem of maximising mutual information in the extended SSR model subject to a maximum energy constraint. Such a problem is like the classic information theory problem of finding channel capacity subject to a power constraint on the source, except that we are free to optimise the channel by changing the threshold values.

The second problem we discuss can also be related to the neural coding motivation for the SSR model, but is also of relevance to lossy source coding. Recall how in Chapter 6 we discuss methods for decoding the SSR model's output signal to obtain a new output signal that approximately reconstructs the input signal. We use the Mean Square Error (MSE) distortion as a measure for the performance of such a reconstruction, and find that, for the SSR model, the distortion can be decreased by increasing the number of threshold devices, N , and therefore the number of output states, which is $N + 1$. Suppose we define the *rate* of a quantiser as the log of the number of output states, $\log(N + 1)$. Then this result is a very simple illustration of *rate-distortion* theory (Berger and Gibson 1998). In general, the distortion of a quantiser can be reduced by increasing

the rate, or, conversely, the rate can be reduced by allowing the distortion to increase. Hence, we consider for the SSR model, and its extension to the arbitrary threshold model, the problem of minimising rate subject to a distortion constraint.

Before discussing these problems in more detail, we briefly state relevant results from previous Chapters, and the theory of solving constrained optimisation problems using the method of Lagrange multipliers.

9.1.1 Review of Relevant Material

The model we use is the array of threshold elements shown in Fig. 4.1. As in Chapters 4–8, we assume that the input signal is a sequence of samples drawn from a continuously valued probability distribution with Probability Density Function (PDF), $P(x)$. The output signal, y , is the sum of the individual outputs of each threshold device, and is a discretely valued signal with $N + 1$ states between 0 and N .

As first introduced in Chapter 8, for arbitrary threshold values we let $P_{1|x,i}$ be the probability of threshold device i being ‘on,’ given signal value, x . Then, if the i -th threshold value is θ_i , we have

$$P_{1|x,i} = \int_{\theta_i-x}^{\infty} R(\eta) d\eta = 1 - F_R(\theta_i - x), \quad (9.1)$$

where $R(\cdot)$ is the noise PDF, $F_R(\cdot)$ is the noise Cumulative Distribution Function (CDF), and $i = 1, \dots, N$.

As first discussed in Chapter 4, the mutual information between the input and output signals in the model is given by

$$I(x, y) = - \sum_{n=0}^N P_y(n) \log_2 P_y(n) - \left(- \int_{-\infty}^{\infty} P(x) \sum_{n=0}^N P(n|x) \log_2 P(n|x) dx \right), \quad (9.2)$$

where $P_y(n) = \int_{-\infty}^{\infty} P(n|x)P(x)dx$ is the probability mass function of the output signal, and $P(n|x)$ are the transition probabilities giving the probability that the output is in state $y = n$, given input value x . We saw in Chapter 4 that for SSR, the mutual information is a function of the noise intensity, σ , that is, the ratio of noise standard deviation to signal standard deviation.

For any arbitrary threshold value, θ_i , and noise PDF, $P_{1|x,i}$ can be calculated exactly for any value of x from Eqn. (9.1). Assuming $P_{1|x,i}$ has been calculated for desired values

9.1 Introduction

of x , each $P(n|x)$ can be calculated from the recursive formulation given in Chapter 8 by Eqn. (8.19).

As first discussed in Chapter 6, the optimal Mean Square Error (MSE) distortion decoding is the nonlinear decoding, \hat{x} , with values given by $\hat{x}_n = E[x|n]$, $n = 0, \dots, N$, and an error signal given by

$$\epsilon = x - \hat{x}. \quad (9.3)$$

This decoding results in the minimum possible MSE distortion for given transition probabilities, which is given by

$$\text{MMSE} = E[\epsilon^2] = E[x^2] - E[\hat{x}^2]. \quad (9.4)$$

Like the mutual information, the MMSE depends on the transition probabilities, since

$$\begin{aligned} E[\hat{x}^2] &= \sum_{n=0}^N \hat{x}_n^2 P_y(n) \\ &= \sum_{n=0}^N E[x|n]^2 P_y(n) \\ &= \sum_{n=0}^N \left(\int_x x P(x|n) dx \right)^2 P_y(n) \\ &= \sum_{n=0}^N \frac{\left(\int_x x P(x) P(n|x) dx \right)^2}{P_y(n)}. \end{aligned} \quad (9.5)$$

We will also in this chapter use the expected value of the output signal, y , which is

$$\begin{aligned} E[y] &= \sum_{n=0}^N n P_y(n) \\ &= \sum_{n=0}^N n \int_{-\infty}^{\infty} P(n|x) P(x) dx \\ &= \int_{-\infty}^{\infty} P(x) \left(\sum_{n=0}^N n P(n|x) \right) dx \\ &= \int_{-\infty}^{\infty} P(x) E[y|x] dx. \end{aligned} \quad (9.6)$$

Substituting Eqn. (8.16) from Chapter 8 into Eqn. (9.6) gives

$$\begin{aligned} E[y] &= \int_{-\infty}^{\infty} P(x) \sum_{i=1}^N P_{1|x,i} dx \\ &= \sum_{i=1}^N \int_{-\infty}^{\infty} P(x) P_{1|x,i} dx \\ &= \sum_{i=1}^N \int_{-\infty}^{\infty} P(x) F_R(\theta_i - x) dx, \end{aligned} \quad (9.7)$$

which holds for arbitrary threshold values. As we saw in Chapters 4 and 6, for the SSR case we have $P(n|x)$ given by the binomial formula, as in Eqn. (4.9), and therefore $E[y|x] = NP_{1|x}$. For an even noise PDF, $R(\eta)$, we also have $E[y] = N/2$.

9.1.2 Constrained Optimisation

Suppose we wish to fix the number of threshold devices, N , and for a given value of noise intensity, σ , find the threshold settings that either:

1. maximise the mutual information, subject to a constraint that specifies that $E[y]$ is less than some value, A , or
2. minimise the mutual information, subject to a constraint that the MSE distortion is less than some value, B .

Since $I(x, y)$, $E[y]$, and the MSE distortion are all functions of the transition probabilities, $P(n|x)$, it is possible to formulate such an optimisation as a variational problem, where the aim is to find the optimal set of transition probabilities. This set will consist of $N + 1$ functions of the continuous variable, x , that is, $\{P(n|x)\}$, $n = 0, \dots, N$. However, as we saw in Chapter 8, each $P(n|x)$ depends entirely on the vector of thresholds, θ , and we can therefore solve optimal quantisation problems by finding the optimal N -dimensional vector, θ^* . We are able to take the same approach for constrained optimisation problems, as we now discuss.

Suppose for any given noise intensity, σ , we label the quantity we wish to optimise as $f(\theta)$, and the variable we wish to constrain as $m(\theta)$. Then the problem of maximising the *cost function*, $f(\theta)$, subject to the constraint that $m(\theta) \leq A$, can be expressed as the nonlinear optimisation problem,

$$\begin{aligned} \text{Find:} & \quad \max_{\theta} f(\theta), \\ \text{subject to:} & \quad m(\theta) \leq A, \theta \in \mathbb{R}^N. \end{aligned} \quad (9.8)$$

The method of Lagrange multipliers (Gershenfeld 1999) can be used to solve such constrained optimisation problems. Using the standard approach to this method, we begin by incorporating the constraint, $m(\theta) \leq A$, in a new cost function, $g(\theta)$, as

$$\begin{aligned} \text{Find:} & \quad \max_{\theta} g(\theta) = f(\theta) - \lambda m(\theta), \\ \text{subject to:} & \quad \theta \in \mathbb{R}^N, \lambda > 0. \end{aligned} \quad (9.9)$$

9.2 Information Theory and Neural Coding

It can be shown that solving Problem (9.8), for some constraint value A , is equivalent to solving Problem (9.9), for some corresponding value of λ . The simplest way of ensuring that the constraint is met is to begin with a guess for a value of λ , and then solve Problem (9.9). If for this value of λ the constraint is not met, then the *Lagrange multiplier*, λ , can be varied and Problem (9.9) solved with this new value. This process can be repeated until λ is such that the constraint, $m(\theta) \leq A$, is satisfied.

As in Chapter 8, we will find that there are many local optima for our constrained optimisation problems. However, we can make use of the same techniques as described in Chapter 8 to find the global optimum.

9.1.3 Chapter Structure

The remainder of this Chapter is separated into two main Sections. Section 9.2 considers neurally motivated energy constraint problems. Section 9.3 considers the problem of rate-distortion tradeoff, which is conventionally part of the domain of lossy source coding theory, but which is also applicable in neural coding situations.

9.2 Information Theory and Neural Coding

There is increasing interest in applying the techniques of electronic engineering and signal processing to neuroscience research. This research field is known as *computational neuroscience* (Rieke *et al.* 1997, Eliasmith and Anderson 2003). The motivation for such studies is obvious; the brain uses electrical signals—as well as chemical signals—to propagate, store and process information, and must employ some sort of coding and modulation mechanism as part of this process. The fields of information theory and signal processing have many mature techniques for dealing with signal propagation and coding, and these techniques can be employed to gain new insights into the ways the brain encodes, propagates, stores and processes information.

Of particular relevance to this thesis is the fact that the brain is capable of performing very well when required to obtain information via the senses in very noisy conditions. Often, the signal-to-noise ratio (SNR) of the sensory neuron is orders of magnitude lower than those usually encountered in electronic systems (Bialek *et al.* 1993). As discussed in Chapter 2, many studies have shown that Stochastic Resonance (SR) can occur in neurons, so that it appears possible that certain tasks required in the nervous

system have evolved to be optimally adapted to operating in such noisy conditions, or alternatively, have evolved to generate noise, to enable its neurons to perform optimally (Longtin 1993, Douglass *et al.* 1993).

There is some debate in the literature regarding which measures are appropriate to use to quantify information transmission in neural systems. Although some authors argue against the use of mutual information (Johnson 2002), it has however been the preferred measure in numerous papers (Levin and Miller 1996, Rieke *et al.* 1997, Borst and Theunissen 1999, Stocks and Mannella 2001, Abarbanel and Rabinovich 2001). One reason for this is that mutual information provides a measure that is independent of any decoding. While lossy source coding theory tends to use mean square error to measure distortion, it is debatable whether such a distortion measure is relevant for describing the quality of information perception by the brain. Consider, for example, a stochastically quantised signal; such a signal is a discrete random variable. A known and invertible—that is, not stochastic or lossy—transformation of this signal will cause changes in the MSE distortion, but not the mutual information. If the important feature of a signal is its shape when considered after some known transformation, then it does not make sense to consider MSE distortion, as the same information is available, even if the distortion becomes huge.

However, given that mutual information has been previously used elsewhere in neural coding research, we also use it here to illustrate the main point, that is, to examine how the optimal thresholds in the extended SSR model change when subject to energy constraints. Hence, we now briefly comment on the existing literature on energy constrained neural coding, and define a measure of energy for the extended SSR model.

9.2.1 Energy Constraints

Many recent studies in the field of biophysics have shown that energy is an important constraint in neural operation, and have investigated the role of this constraint in neural information processing (Balasubramanian *et al.* 2001, Laughlin 2001, Wilke and Eurich 2001, Bethge *et al.* 2002, Levy and Baxter 2002, Schreiber *et al.* 2002, Hoch *et al.* 2003b).

Recall that in the SSR model, and its extension to arbitrary thresholds, that when *iid* additive noise is present at the input to each threshold element, the overall output becomes a randomly quantised version of the input signal, which for the right level of

9.2 Information Theory and Neural Coding

noise is highly correlated with the input signal. In this Section, we view the extended SSR model as a population of simple neuron models, and present an investigation of the optimal noisy encoding of a random input signal, subject to constraints on the available energy expenditure. Our model is extremely simplified when compared to real neurons, or realistic models, however it does encapsulate the basic aspects of how stochastic quantisation may occur in a population of neurons.

The only previous work on SSR of relevance to such a goal is work that imposes an energy constraint on the SSR model, and calculates the optimal input signal distribution given that constraint (Hoch *et al.* 2003b).

Previously in this thesis, Chapter 8 deals with optimising the thresholds with no constraints. By contrast, here we look at two energy constrained problems. In the first problem, we fix the input distribution and ‘neural population size,’ N , and aim to find the threshold settings that maximise the mutual information, subject to a maximum average output energy constraint. The second problem we tackle is, the perhaps more biologically relevant problem, of minimising the population size, N —and therefore the energy expenditure—given a minimum mutual information constraint and fixed thresholds.

Before we discuss these problems, however, we firstly define our measures of energy and energy efficiency, and discuss how these measures vary for the SSR model, where all thresholds have the same value.

Average output energy

The simplest approach is to assume that the energy expended is the same constant amount every time a neuron emits a ‘spike’. Therefore, minimising the average output energy consumption requires minimisation of the mean output value, which for our model is $E[y]$, as given by Eqn. (9.6). Note that for arbitrary thresholds, the threshold values control the output probabilities, $P_y(n)$. Therefore by ensuring some thresholds are set high enough, larger values of y can be made less probable in order to reduce the average output energy expended.

Information efficiency

Previously, Schreiber *et al.* (2002) makes use of an efficiency measure defined as the ratio of mutual information to metabolic energy required in a neural coding. We use

a similar metric, the ratio of mutual information to average output energy, which we will call *information efficiency* and denote as

$$\zeta = \frac{I(x, y)}{E[y]}. \quad (9.10)$$

This quantity is a measure of bits per unit energy. Clearly this information efficiency measure is increased if either the mutual information increases, or average output energy decreases. We will see that this measure is only useful if some constraint is placed on either mutual information or energy, since the efficiency can be made near infinite by ensuring that all thresholds are so large that the output of the model is almost always zero. Hence, for practical purposes, maximising the information efficiency requires either maximising the mutual information subject to a maximum output energy constraint, or minimising the average output energy subject to a minimum mutual information constraint.

However, first we will consider the situation where all thresholds are not necessarily zero, but are equal to the same value, θ . We compare how the mutual information and average output energy varies with noise intensity, σ , and θ , as well as how the mutual information varies for specified energy constraints. Examining the behavior of the information efficiency given by Eqn. (9.10) for these situations will provide a useful benchmark for our later constrained optimal quantisation problems.

9.2.2 All Thresholds Equal

If all thresholds are equal to the signal mean of zero, then the model is in the SSR configuration, and if $P(x)$ is an even function the average output energy is $N/2$. If however, the value of all thresholds, θ , is changed to be nonzero, then both the mutual information and average output energy will change. We will consider only the situation of a Gaussian signal and independent Gaussian noise, since we are able to state an exact result for the output energy in this case

Exact output energy for Gaussian signal and noise

Consider the total probability that the i -th device is 'on'. This is a function of that threshold's actual value, θ_i , which we can write as

$$\begin{aligned} f(\theta_i) &= \int_{-\infty}^{\infty} P(x) P_{1|x,i} dx \\ &= 1 - \int_{-\infty}^{\infty} P(x) F_R(\theta_i - x) dx. \end{aligned} \quad (9.11)$$

9.2 Information Theory and Neural Coding

For even noise distributions, Eqn. (9.11) can be written as

$$f(\theta_i) = \int_{-\infty}^{\infty} P(x)F_R(x - \theta_i)dx. \quad (9.12)$$

However, for fully general $R(\cdot)$, taking the derivative of both sides of Eqn. (9.11) with respect to θ_i gives

$$\frac{df(\theta_i)}{d\theta_i} = - \int_{-\infty}^{\infty} P(x)R(\theta_i - x)dx. \quad (9.13)$$

Note that $-\frac{df(\theta_i)}{d\theta_i}$ is the convolution of $P(x)$ and $R(x)$, which are both Gaussians. The convolution of two Gaussians is another Gaussian (Rényi 1970) and we therefore have

$$\frac{df(\theta_i)}{d\theta_i} = -\mathcal{N}(\theta_i, \sigma_x^2 + \sigma_\eta^2), \quad (9.14)$$

where $\mathcal{N}(\mu, s^2)$ is a Gaussian PDF with mean μ and variance s^2 . Thus, we must have

$$\begin{aligned} f(\theta_i) &= \int_{-\infty}^{-\theta_i} \frac{1}{\sqrt{2\pi(\sigma_x^2 + \sigma_\eta^2)}} \exp\left(-\frac{\tau^2}{2(\sigma_x^2 + \sigma_\eta^2)}\right) d\tau \\ &= 0.5 + 0.5\operatorname{erf}\left(\frac{-\theta_i}{\sqrt{2(\sigma_x^2 + \sigma_\eta^2)}}\right) \\ &= 0.5 + 0.5\operatorname{erf}\left(\frac{-\theta_i}{\sigma_x\sqrt{2(1 + \sigma^2)}}\right) \\ &= 0.5 - 0.5\operatorname{erf}\left(\frac{\theta_i}{\sigma_x}\sqrt{\frac{\rho}{2}}\right), \end{aligned} \quad (9.15)$$

where ρ is the correlation coefficient between the input to any two different threshold devices, as discussed in Appendix 6. Thus, $f(\theta_i)$ depends on both σ_η and σ_x , not just their ratio, σ . However, for SSR we have $\theta_i = 0$, and $f(0) = 0.5$, regardless of the actual magnitudes of σ_x and σ_η .

Now, the expected value of y is

$$\begin{aligned} E[y] &= \sum_{i=1}^N f(\theta_i) \\ &= \frac{N}{2} - \frac{1}{2} \sum_{i=1}^N \operatorname{erf}\left(\frac{\theta_i}{\sigma_x\sqrt{2(1 + \sigma^2)}}\right). \end{aligned} \quad (9.16)$$

If all thresholds are equal to the same value, θ , then

$$E[y] = \frac{N}{2} - \frac{N}{2} \operatorname{erf}\left(\frac{\theta}{\sigma_x\sqrt{2(1 + \sigma^2)}}\right), \quad (9.17)$$

and clearly, the average energy is a function of the ratio of θ to σ_x , as well as σ .

We are now able to present results showing how the mutual information and average output energy vary with θ and σ .

Results

Fig. 9.1(a) shows numerical calculations of the mutual information against increasing σ for a Gaussian signal with $\sigma_x = 1$, and Gaussian noise. We have $N = 5$, and consider various threshold values. If the value of σ_x were to be changed, then the thresholds shown would need to be adjusted to keep θ/σ_x constant to give the same results. Fig. 9.1(b) shows the average output energy for the same conditions, calculated from the exact formula of Eqn. (9.17) and validated numerically. The information efficiency is plotted in Fig. 9.1(c). As we expect—given the result in Stocks (2001a) that shows that the optimal threshold value is equal to the signal mean—the mutual information decreases for $\theta \neq 0$. However, as we might also expect, the average output energy is greater than $N/2$ for $\theta < 0$ and less than $N/2$ for $\theta > 0$.

More importantly, Fig. 9.1(c) shows that the information efficiency increases with increasing threshold values, and therefore with decreasing energy. Thus, if we wish to maximise the information efficiency, this illustrates the requirement for placing a minimum mutual information or maximum average output energy constraint for nontrivial results.

Fig. 9.2 shows the mutual information, average output energy, and information efficiency for various values of σ as a function of θ/σ_x . Of the four values of σ shown, at $\theta = 0$ the mutual information is greatest for $\sigma = 0.25$, as we know it should be, since this is the SSR situation discussed in Chapter 4. However, for θ nonzero, the optimal mutual information can occur for other values of σ , as also illustrated in Fig. 9.1(a). At $\theta = 0$, the average output energy is $E[y] = N/2$, as we know it always is for SSR. Fig. 9.2(c) shows that for $\theta \geq 0$, the information efficiency is largest for $\sigma = 0.25$, and for each value of σ increases with increasing θ/σ_x , as also shown in Fig. 9.1(c). For $\theta < 0$, the information efficiency is very small in comparison with $\theta > 0$. This is due to the mutual information decreasing with decreasing σ for $\theta < 0$, while the average energy continues to increase. Fig. 9.2(d) shows the variation in mutual information with $E[y]$. For a given value of σ , the mutual information is maximised at $N/2$, and decreases as the energy decreases from this value.

9.2 Information Theory and Neural Coding

Suppose we require that $E[y] = A$. Then rearranging Eqn. (9.17) gives

$$\frac{\theta}{\sigma_x} = \sqrt{2(1 + \sigma^2)} \operatorname{erf}^{-1} \left(1 - \frac{2A}{N} \right). \quad (9.18)$$

Fig. 9.3 shows the mutual information and information efficiency for various fixed values of $E[y]$. For each value of $E[y]$, the corresponding value of θ was found from Eqn. (9.18).

The main conclusion from these results is that the information efficiency can be made arbitrarily large by decreasing the output energy. However, as the output energy decreases, the mutual information increases. Hence, for any value of σ , there is a tradeoff between mutual information and energy. We now progress to considering two constrained optimal quantisation problems.




Figure 9.1. Information efficiency as a function of σ , all thresholds equal. Fig. 9.1(a) shows the result of numerically calculating the mutual information, $I(x, y)$, as a function of noise intensity, σ , for various threshold values and $\sigma_x = 1$. Note that since both the signal and noise have even PDFs, the mutual information is the same for threshold values with the same magnitude, but opposite sign. Hence, as can be seen in Fig. 9.1(a), the mutual information is identical for $\theta = \pm 0.5$. As the magnitude of the threshold value increases, the mutual information decreases for all σ . However, the value of σ at which the peak value of the mutual information occurs increases with increasing σ . Fig. 9.1(b) shows the average output energy calculated from the exact formula of Eqn. (9.17) and validated numerically. The average energy decreases with increasing σ . The information efficiency is shown in Fig. 9.1(c), which indicates that the efficiency increases with increasing σ , and therefore with increasing output energy.

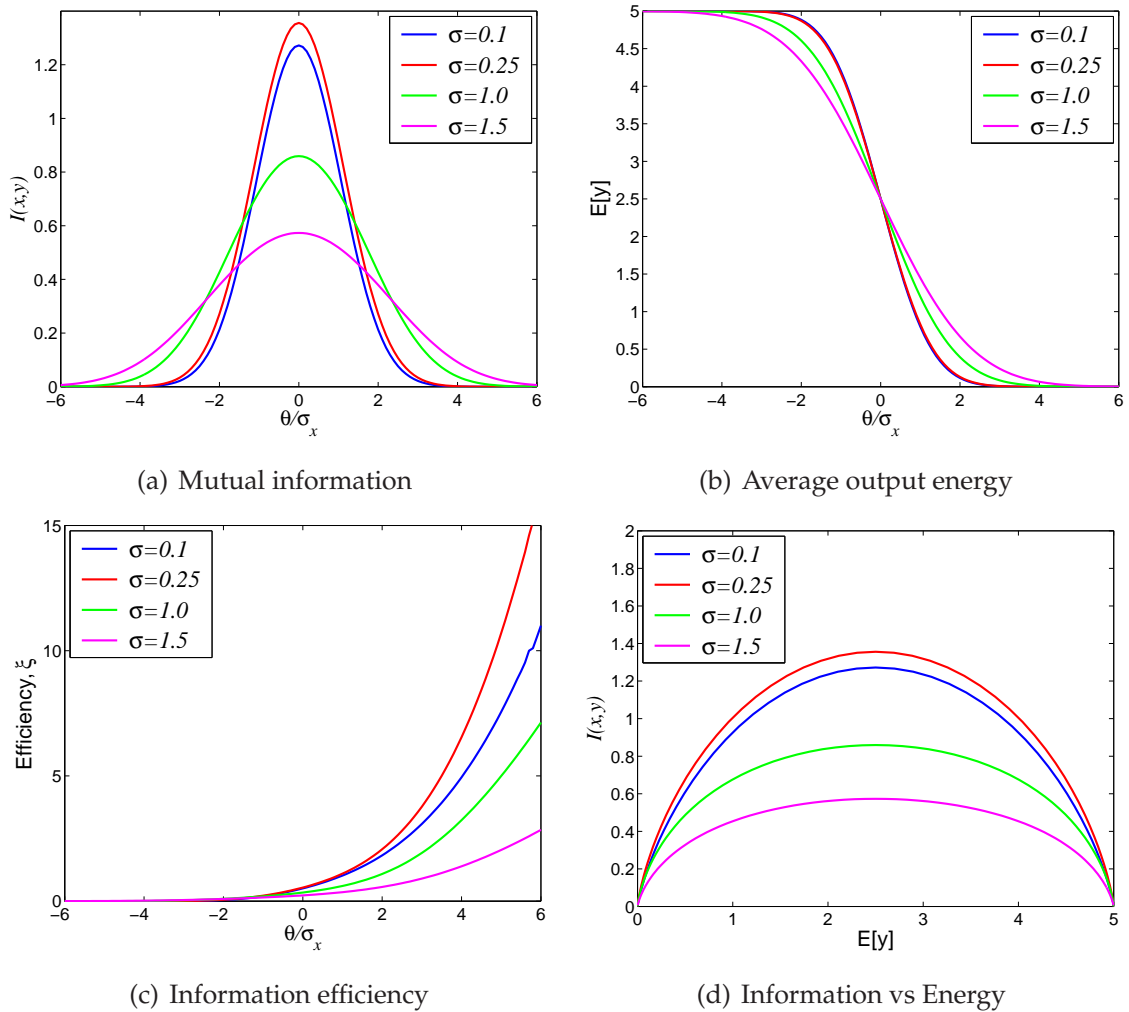


Figure 9.2. Information efficiency as a function of θ/σ_x , all thresholds equal. Fig. 9.2(a) shows the mutual information, $I(x,y)$, calculated numerically for four different values of noise intensity, σ , as a function of θ/σ_x . Of the four values of σ shown, at $\theta = 0$ the mutual information is greatest for $\sigma = 0.25$, as we know it should be, since this is the SSR situation discussed in Chapter 4. However, for θ nonzero, the optimal mutual information can occur for other values of σ , as also illustrated in Fig. 9.1(a). Fig. 9.2(b) shows the average output energy calculated from the exact formula of Eqn. (9.17) and validated numerically. At $\theta = 0$, the average output energy is $E[y] = N/2$, as we know it always does for SSR. Fig. 9.2(c) shows that for $\theta \geq 0$, the information efficiency is largest for $\sigma = 0.25$, and for each value of σ increases with increasing θ/σ_x , as also shown in Fig. 9.1(c). For $\theta < 0$, the information efficiency is very small in comparison with $\theta > 0$. This is due to the mutual information decreasing with decreasing σ for $\theta < 0$, while the average energy continues to increase. Fig. 9.2(d) shows the variation in mutual information with $E[y]$. For a given value of σ , the mutual information is maximised at $N/2$, and decreases as the energy decreases from this value.

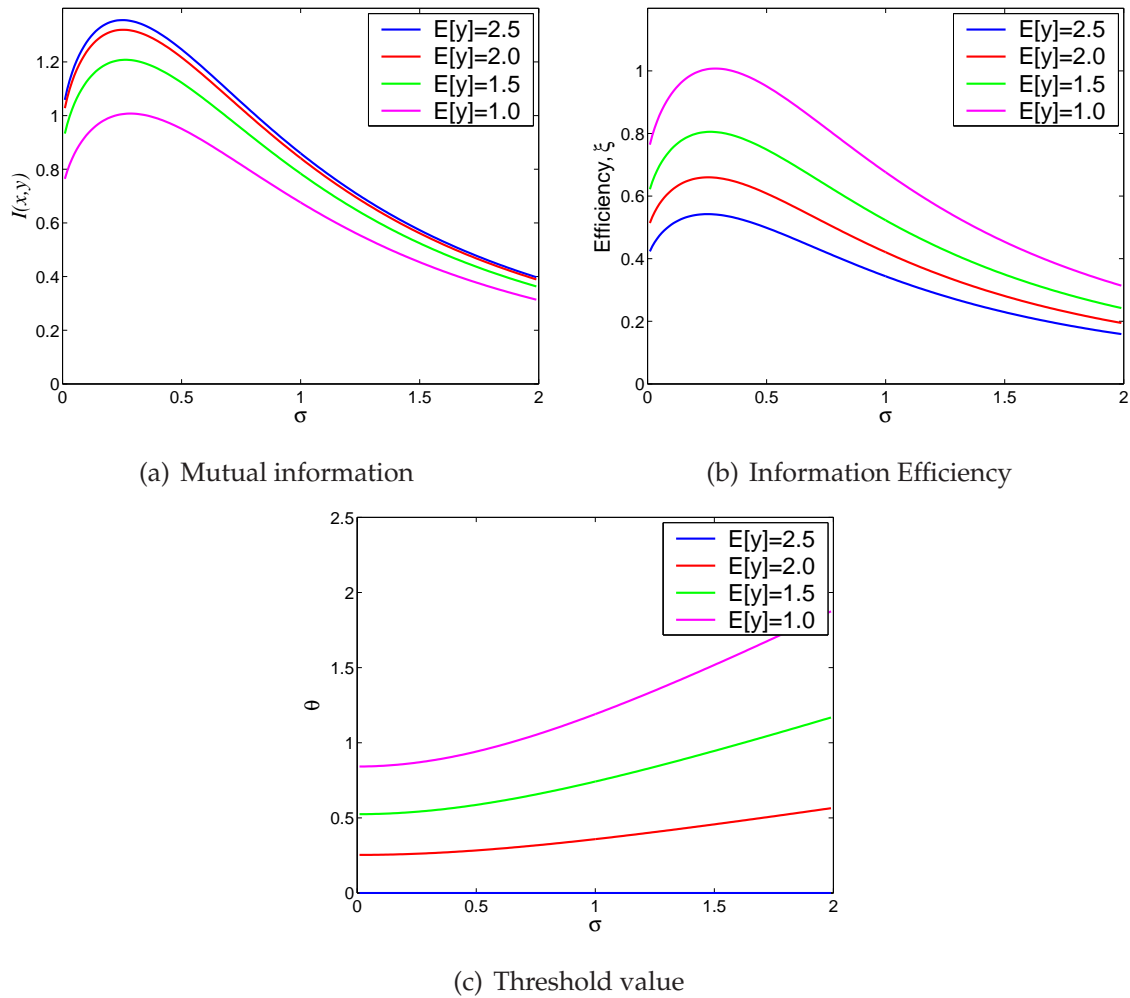


Figure 9.3. Information efficiency for fixed energy. This figure shows numerical calculations of the mutual information, $I(x, y)$, and information efficiency, ξ , for four specified values of output energy, as a function of noise intensity, σ . The threshold value, θ , that achieves the specified value of the average output energy, $E[y]$, is shown in Fig. 9.3(c), and was calculated from the exact formula of Eqn. (9.18).

9.2.3 Energy Constrained Optimal Quantisation

The first constrained problem we aim to solve is the problem of maximising the mutual information subject to the constraint that the average output energy must be no larger than some value, A . This can be expressed as

$$\begin{aligned} \text{Find:} & \quad \max_{\boldsymbol{\theta}} f(\boldsymbol{\theta}) = I(x, y), \\ \text{subject to:} & \quad m(\boldsymbol{\theta}) = \mathbb{E}[y] \leq A, \boldsymbol{\theta} \in \mathbb{R}^N. \end{aligned} \quad (9.19)$$

If we now write this in terms of a Lagrange multiplier, λ , we have

$$\begin{aligned} \text{Find:} & \quad \max_{\boldsymbol{\theta}} g(\boldsymbol{\theta}) = I(x, y) - \lambda \mathbb{E}[y^2], \\ \text{subject to:} & \quad \boldsymbol{\theta} \in \mathbb{R}^N, \lambda \geq 0. \end{aligned} \quad (9.20)$$

As in Chapter 8, we set $\sigma_x = 1$, and expect that the optimal thresholds will scale in inverse proportion to σ_x .

The results of numerical solutions of Problem (9.20) for $N = 5$, Gaussian signal and noise, and various values of A are shown in Figs. 9.4(a), and 9.4(b), which show the mutual information and energy efficiency, for several values of A , and in Fig. 9.5, which shows the optimal thresholds. The case of no constraint is also plotted for comparison. Figs. 9.4(a), and 9.4(b) also show with thin solid lines the mutual information and energy efficiency obtained for each constraint if all thresholds are set to the same value, as considered in Section 9.2.2. Note that the energy constraint is always met with equality, since we find that mutual information always decreases if the energy is made smaller than the constraint value. The value of λ that satisfies each constraint also varies with σ . These values are shown for each energy constraint in Fig. 9.4(c).

Fig. 9.4(a) shows that as the maximum output energy decreases, the mutual information also decreases, but only by a fraction of a bit per sample. From Fig. 9.4(b), the information efficiency increases as the maximum output energy decreases. From this result, it is clear that information efficiency can be substantially increased with only a small decrease in mutual information.

Fig. 9.5 shows that as the maximum average energy constraint decreases, there is an upward shift in the optimal thresholds. This shift increases with increasing σ . Thus, the optimal threshold values are larger than the unconstrained optimal thresholds. This is to be expected, since larger threshold values mean less threshold crossings,

9.2 Information Theory and Neural Coding

and therefore a lower average output energy. However, otherwise the same qualitative behaviour as the corresponding unconstrained problem occurs. For sufficiently large σ , the optimal solution is for all thresholds to have the same value. This value is obtained by substituting the desired value of A into Eqn. (9.18) and solving for θ . For small σ we have the same bifurcational behaviour in the optimal threshold diagram as described in Chapter 8.

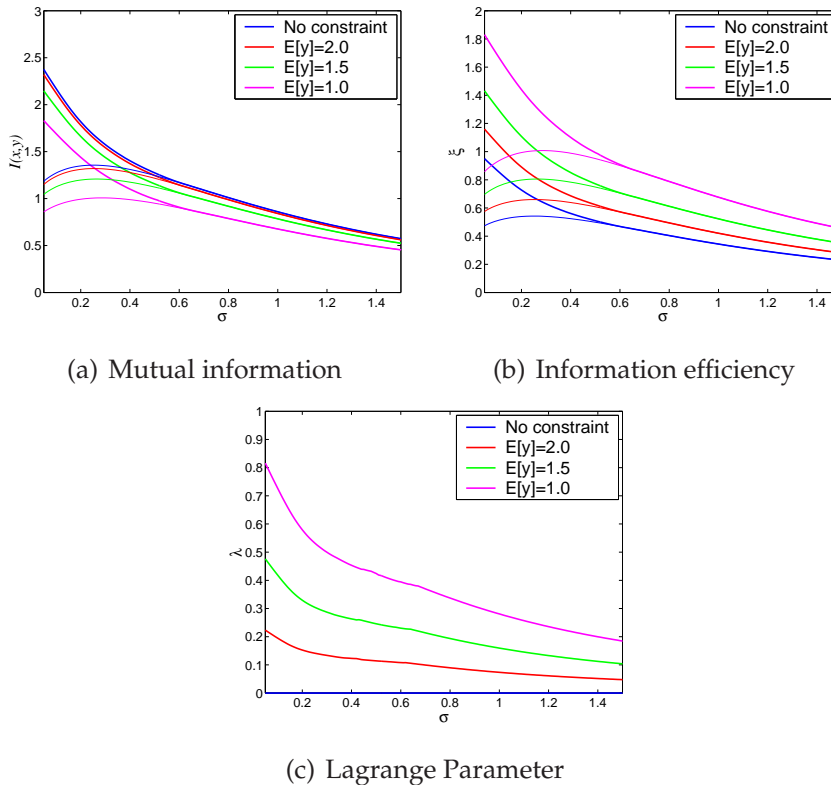


Figure 9.4. Information efficiency for optimised thresholds. Fig. 9.4(a) shows with thick solid lines the mutual information, $I(x,y)$, against increasing noise intensity, σ , obtained by numerically solving Problem (9.20) for three different values of the constraint on the average output energy. It also shows with thin solid lines the mutual information obtained for the same constraint when all thresholds have the same value. Fig. 9.4(b) shows the corresponding energy efficiency, ξ , against σ . Fig. 9.4(c) shows the value of the Lagrange parameter, λ , required to achieve each constraint as a function of σ .




Figure 9.5. Optimal energy constrained thresholds. This figure shows the optimal threshold values obtained when Problem (9.20) is solved for three different values of an energy constraint, for the case of Gaussian signal and noise, and $N = 5$. Also shown are the optimal thresholds for the corresponding unconstrained problem, as solved in Chapter 8. As the maximum average energy constraint decreases, we see an upward shift in the optimal thresholds. Thus, the optimal threshold values are larger than the unconstrained optimal thresholds. This is to be expected, since larger threshold values mean less threshold crossings, and therefore a lower average output energy. However, otherwise the same qualitative behaviour as the unconstrained problem occurs. For sufficiently large σ , the optimal solution is for all thresholds to have the same value.

9.2.4 Fixed Thresholds and a Minimum Information Constraint

In a biological context, increasing the energy efficiency of information encoding by optimising threshold values may not be particularly realistic. To be useful to an organism, a sensory system must convey a certain amount of information, \hat{I} , to the brain. Any less than this and the sensory function is not useful. For example, to avoid becoming prey, a herbivore must be able to distinguish a sound at, say, 100 m whereas 50 m may not be sufficient. To develop such a sensory ability, the animal must use more neurons, and therefore use more energy—in the form of generated spike trains, as well as energy to grow them in the first place—to convey the information to the brain. For maximum evolutionary advantage the animal must do this in the most energy efficient manner—otherwise it may have to eat more than is feasible—so the problem that evolution solves in the animal is not to fix N and maximise $I(x, y)$, but fix $I(x, y)$ and minimise N (Stocks and Morse 2003). If the energy is an increasing function of N , then this problem is equivalent to minimising the energy. This is slightly different to the problem of maximising information subject to an energy constraint, although the most efficient way of transmitting $I(x, y)$ bits of information must be equivalent to finding the value of energy E that gives rise to the optimal amount of information $I(x, y)$.

Thus, for this problem we wish to fix the mutual information and threshold settings, and find the population size, N , that minimises the energy for a range of σ . For the SSR situation where all thresholds are equal to the signal mean, and $P(x)$ and $R(\eta)$ are zero mean even functions, the average output energy is simply $N/2$ and hence increases linearly with increasing N . Therefore the solution to this problem reduces to finding the minimum N that satisfies $I(x, y) \geq B$. Thus, this problem can be expressed as

$$\begin{aligned} \text{Find:} & \quad \min N, \\ \text{subject to:} & \quad I(x, y) \geq B, \boldsymbol{\theta} = \mathbf{0}, N \in \mathbb{Z}^+. \end{aligned} \quad (9.21)$$

The result of solving Problem (9.21) for Gaussian signal and noise and three values of $I(x, y)$ is shown in Fig. 9.6, as a function of σ . It is clear that N is minimised for a nonzero value of σ , which we might expect, given that we have the SSR situation.

9.2.5 Discussion and Further Work

In this Section we have formulated and solved two different problems of energy efficient information transmission in the extended SSR model.




Figure 9.6. Minimising N subject to an information constraint. This plot shows the result of minimising N for three constraints on $I(x, y)$, for Gaussian signal and noise, as a function of σ . It is clear that N is minimised for a nonzero value of σ for each constraint.

For the problem of finding the threshold settings that maximise the information efficiency for fixed N , our results clearly indicate that the information efficiency always increases for decreasing average output energy and always decreases for increasing mutual information. However, a fairly large information efficiency increase can occur for only a small decrease in mutual information. We also show that the maximum information efficiency for optimised thresholds is always strictly decreasing for increasing noise intensity, σ , that is, no SR effect is seen. However, since populations of real neurons are not known for having widely distributed thresholds, the information efficiency results shown in Fig. 9.1(c) are the most biologically relevant. This figure shows that the maximum information efficiency occurs for nonzero noise. This is due to the same mechanism as SSR, where the noise acts to randomly distribute the effective threshold values.

We also briefly discuss the problem of minimising the number of ‘neurons’ required to achieve a certain minimum mutual information. For the SSR model, the minimum N occurs for nonzero σ . In contrast, we expect that if the same problem is solved for arbitrary thresholds, that N will increase with increasing σ . However, we also expect that for small σ , the optimal thresholds will be widely distributed, and for large σ , the optimal thresholds will be the SSR situation of all thresholds equal to the signal mean.

9.3 Rate-Distortion Tradeoff

There is much further work to be completed on these problems, and the results presented are intended to convey ‘proof of principle’ rather than a comprehensive investigation. Future research might, for example, introduce energy constraints and variable thresholds to Problem (9.21). Insight gained from the study of these highly simplified neural population model problems may be of benefit to the understanding of neural coding in the presence of noise in more realistic neuron models and real neurons.

We now, in Section 9.3 consider the rate-distortion tradeoff issue in the SSR model and its extension to arbitrary thresholds.

9.3 Rate-Distortion Tradeoff

From a quantiser design point of view, it is desirable to minimise both the *rate*, and the *distortion*. Although the exact definitions of both rate and distortion can vary, in quantisation, the term ‘rate’ loosely corresponds to the amount of compression obtained by a quantiser, and hence rate-distortion theory falls into the category of *lossy source coding* or *lossy compression* theory. However, since rate and distortion are both dependent on the same variables, they cannot both be simultaneously minimised, and there must be a trade-off between rate and distortion. Thus, the rate-distortion problem is usually formulated as the problem of minimising the rate, subject to a specified constraint that the distortion can be no larger than some fixed value, D . The minimum possible rate that achieves distortion D for a given source distribution is known as the *rate-distortion function*, denoted as $R(D)$. In general, $R(D)$ is a theoretical limit that cannot be achieved in practice.

In such rate-distortion theory, rate is generally defined as the mutual information between the input and output of a quantiser. Therefore, if we consider the SSR model—or its extension to arbitrary thresholds—from such a viewpoint, we must perform the opposite task to that considered in Section 9.2, and minimise the mutual information rather than maximise it. The reason that we can consider both goals without contradiction is that in Section 9.2 we seek to optimise information transmission—that is, we address the question of “what is the largest amount of information, on average, that can be transmitted in a channel, subject to some constraints?” By contrast, this Section

address a compression problem, by aiming to answer the question “What is the smallest amount of information required to represent a signal with, on average, a distortion no larger than some value, D ?”

Of particular relevance to this discussion is a simple formula known as the *information transmission inequality* (Berger and Gibson 1998), which relates rate, distortion, and channel capacity, C . This formula is

$$D \geq R^{-1}(C). \quad (9.22)$$

Note that in this formula, channel capacity is the maximum mutual information for a given channel, and the rate distortion function, $R(D)$, is the minimum mutual information to achieve distortion D . This formula says that if one is trying to transmit data from a source with rate-distortion function $R(D)$, over a channel with capacity C , the average distortion achieved will be greater than the inverse of the rate-distortion function evaluated at C (Berger 1971).

Looked at another way, suppose a communications system designer is required to transmit information through a channel with capacity C , and is aiming for fidelity, or average distortion, D . Suppose also that the information source has a known rate-distortion function, $R(D)$. Then the designer knows immediately whether there is any chance of achieving their requirements, because success will only be possible if $C \geq R(D)$. This illustrates the utility of knowing the theoretical optimal rate-distortion tradeoff, and why there are different reasons for both maximising, and minimising mutual information. We now discuss in more detail the basic theory of the rate-distortion function.

9.3.1 Rate-Distortion Theory

The optimal trade-off between rate and distortion is measured using the rate-distortion function (Berger and Gibson 1998), often expressed as $R(D)$, where R is the rate—defined as the mutual information—and D is some arbitrary distortion measure, often taken to be MSE distortion. The rate distortion function is defined as the solution to the following constrained optimisation problem,

$$\begin{array}{ll} \text{Find:} & \min I(x, y), \\ \text{subject to:} & \text{Distortion} \leq D. \end{array} \quad (9.23)$$

9.3 Rate-Distortion Tradeoff

Thus, $R(D)$ is defined as being the minimum possible rate to achieve distortion D . This implies that in practice, the actual rate achieved in a quantiser will always be $I(x, y) \geq R(D)$.

Exact $R(D)$ for a Gaussian source

Shannon derived an exact analytic result for the rate distortion function for a Gaussian source with variance σ_x^2 , and the MSE distortion measure (Berger and Gibson 1998, Cover and Thomas 1991). This exact result states that

$$R(D) = \begin{cases} 0.5 \log_2 \left(\frac{\sigma_x^2}{D} \right) & 0 \leq D \leq \sigma_x^2, \\ 0 & D \geq \sigma_x^2 \end{cases} \quad (9.24)$$

and says that no quantisation scheme can achieve a distortion less than D with a rate smaller than $R(D)$, for a Gaussian source with variance σ_x^2 . In other words, a quantisation scheme with rate R will provide a mean square distortion no smaller than D . Thus, we have a lower bound for the mutual information for quantisation of a Gaussian source as

$$I(x, y) \geq 0.5 \log_2 \left(\frac{\sigma_x^2}{D} \right). \quad (9.25)$$

We briefly explain why $R(D)$ is zero for $D \geq \sigma_x^2$. Consider the region $D > \sigma_x^2$; here, distortion larger than σ_x^2 can be achieved without transmitting any information. For a zero mean Gaussian signal, the critical distortion, $D = \sigma_x^2$, is achievable by simply always guessing that the input signal is always zero. In this case the MSE distortion is simply the variance of the signal, σ_x^2 . For larger distortions, if we always guess that the input is some nonzero value, say a , then the MSE distortion is $\sigma_x^2 + a^2$. Therefore for distortion $D > \sigma_x^2$, guessing $a = \sqrt{(D - \sigma_x^2)}$ gives distortion D with a rate of zero. However, in the region $D < \sigma_x^2$, to achieve some distortion, $D < \sigma_x^2$, it is necessary to transmit some information, $R(D) > 0$.

The rate-distortion function can be inverted to obtain the *distortion-rate function*, $D(R)$. Carrying this out for a Gaussian source using Eqn. (9.24) gives

$$D(R) = \sigma_x^2 2^{-2R}. \quad (9.26)$$

Eqn. (9.26) says that no quantisation of a Gaussian source can achieve a distortion smaller than D , if the rate is to be no larger than R .

Eqn. (9.26) can be expressed as the signal-to-quantisation-noise ratio (SQNR),

$$\text{SQNR} = \frac{\sigma_x^2}{D(R)} = 2^{2R}. \quad (9.27)$$

Taking the base ten logarithm of both sides of Eqn. (9.26) and multiplying by 10 gives

$$\begin{aligned} 10 \log_{10} (D) &= 10 \log_{10} \sigma_x^2 - 20 \log_{10} (2)R \\ &\simeq 20 \log_{10} \sigma_x - 6.02R. \end{aligned} \quad (9.28)$$

Rearranging Eqn. (9.28) gives the maximum possible SQNR of a Gaussian source with rate R in decibels as

$$10 \log_{10} \left(\frac{\sigma_x^2}{D} \right) = (20 \log_{10} (2)) R \simeq 6.02R \quad \text{dB}. \quad (9.29)$$

This corresponds with a well-known rule of thumb in quantiser design that states that a one bit increase in rate gives about a 6 dB increase in SNR (Gray and Neuhoff 1998). Furthermore, it shows that the maximum possible SQNR for a given R is proportional to R .

$R(D)$ for non-Gaussian source distributions

Exact expressions for $R(D)$ like that given by Eqn. (9.24) are quite rare. For most source distributions, numerical techniques are required to find the rate-distortion function, and fortunately the Arimoto-Blahut algorithm can be used to achieve this (Blahut 1972, Arimoto 1972).¹⁶

However, although analytical formulas for the exact rate-distortion function are difficult to find, much attention has been focused on finding upper or lower bounds to $R(D)$ for non-Gaussian sources. In particular, consider the following reasoning, where y is the output of the encoding operation of a quantiser, \hat{x} is the MMSE distortion decoding of y , and $\epsilon = x - \hat{x}$ is the error signal,

$$\begin{aligned} I(x, y) &= H(x) - H(x|y) \\ &= H(x) - H(x|\hat{x}) \\ &= H(x) - H(x - \hat{x}|\hat{x}) \\ &= H(x) - H(\epsilon|\hat{x}) \\ &\geq H(x) - H(\epsilon). \end{aligned} \quad (9.30)$$

¹⁶The Arimoto-Blahut algorithm can also be used to calculate channel capacity—see Chapter 4.

9.3 Rate-Distortion Tradeoff

The above steps hold since, first, the average conditional entropy of x given y is identical for a decoding of y . Secondly, the conditional entropy of ϵ given \hat{x} is the same as the conditional entropy of x . Thirdly, the conditional entropy of ϵ is always smaller than the unconditioned entropy (Cover and Thomas 1991).

Now, we also know that the entropy of a continuous random variable, v , with variance, σ_v^2 , is maximised by the Gaussian distribution, and is given by

$$H(v) = 0.5 \log_2 (2\pi e \sigma_v^2). \quad (9.31)$$

Therefore, since for optimal decoding the mean of the error is zero, then the variance of the error signal, ϵ , is the MSE distortion, D , and from Inequality (9.30),

$$I(x, y) \geq H(x) - 0.5 \log_2 (2\pi e D). \quad (9.32)$$

Now, let

$$R_L(D) = H(x) - 0.5 \log_2 (2\pi e D). \quad (9.33)$$

It can be shown (Berger and Gibson 1998, Linder and Zamir 1999) that

$$I(x, y) \geq R(D) \geq R_L(D). \quad (9.34)$$

Eqn. (9.33) is therefore known as the *Shannon lower bound* for the rate-distortion function. It applies to any source PDF, $P(x)$, and gives a lower bound on the rate-distortion function.

For the specific case of a Gaussian source, we also have $H(x) = 0.5 \log_2 (2\pi e \sigma_x^2)$, and therefore

$$R_L(D) = 0.5 \log_2 \left(\frac{\sigma_x^2}{D} \right). \quad (9.35)$$

In fact, the reasoning given to arrive at this result forms the first part of the proof of Eqn. (9.24) for a Gaussian source. The second part of the proof is to find a condition for which the bound of Inequality (9.32) is achievable, and is given in Cover and Thomas (1991). This also shows that $R_L(D)$ becomes tight with the actual $R(D)$ as given by Eqn. (9.24), so that $R_L(D) = R(D)$ for a Gaussian source.

For non-Gaussian sources, the Shannon lower bound will always be smaller than that for a Gaussian source with the same variance as the non-Gaussian source. Eqn. (9.33) can also be simplified to be written in a manner similar to Eqn. (9.35) using the concept of *entropy power*. Consider a random variable with PDF $P(x)$ and differential entropy, $H(x)$. The entropy power, Q_σ , of this random variable is defined as being the variance

Table 9.1. Entropy and Shannon lower bound for various distributions.

Distribution	$P(x)$	Variance	$H(x)$	$R_L(D)$
Gaussian	$\frac{1}{\sqrt{2\pi\sigma_x^2}} \exp\left(-\frac{x^2}{2\sigma_x^2}\right)$	σ_x^2	$0.5 \log_2(2\pi e\sigma_x^2)$	$0.5 \log_2\left(\frac{\sigma_x^2}{D}\right)$
Uniform	$\begin{cases} \frac{1}{\sigma_x}, & x \in [-\frac{\sigma_x}{2}, \frac{\sigma_x}{2}], \\ 0, & \text{otherwise} \end{cases}$	$\frac{\sigma_x^2}{12}$	$\log_2(\sigma_x)$	$0.5 \log_2\left(\frac{\sigma_x^2}{2\pi e D}\right)$
Laplacian	$\frac{1}{\sqrt{2}\sigma_x} \exp\left(-\frac{\sqrt{2} x }{\sigma_x}\right)$	σ_x^2	$\log_2(\sqrt{2}e\sigma_x)$	$0.5 \log_2\left(\frac{e\sigma_x^2}{\pi D}\right)$

of a Gaussian random variable that has differential entropy equal to $H(x)$ (Berger and Gibson 1998). Since the differential entropy of a Gaussian distribution is given by Eqn. (9.31), this means that

$$Q_o = \frac{2^{2H(x)}}{2\pi e}, \quad (9.36)$$

and therefore,

$$R_L(D) = 0.5 \log_2\left(\frac{Q_o}{D}\right). \quad (9.37)$$

We also have the SQNR in dB corresponding to $R_L(D)$ as

$$10 \log_{10}\left(\frac{E[x^2]}{D}\right) = 10 \log_{10}\left(\frac{E[x^2]}{Q_o}\right) + (20 \log_{10}(2))R_L(D). \quad (9.38)$$

The entropy, and Shannon lower bound corresponding to three different distributions are shown in Table 9.1. Note that there is no simple analytical expression available for the entropy and Shannon lower bound of the logistic distribution for arbitrary variance, although these quantities can easily be calculated numerically.

$R(D)$ and its relationship to correlation coefficient

Recall the formula in Chapter 6 that expresses the SQNR in terms of the correlation coefficient between the input and decoded output of the SSR model, ρ , as

$$\text{SQNR} = \frac{E[x^2]}{D} = \frac{1}{1 - \rho^2}. \quad (9.39)$$

In terms of decibels this is

$$10 \log_{10}(\text{SQNR}) = -10 \log_{10}(1 - \rho^2). \quad (9.40)$$

This formula holds both for optimal linear (Wiener) decoding, and for MMSE decoding, although the correlation coefficient is smaller for MMSE decoding, for which we achieve the largest possible correlation coefficient, $\rho_{x\hat{x}}$. Substituting Eqn. (9.39) into

9.3 Rate-Distortion Tradeoff

Eqn. (9.35) then gives the rate-distortion function for a Gaussian source—which is also the Shannon lower bound—in terms of the correlation coefficient as

$$R(D) = R_L(D) = -0.5 \log_2 (1 - \rho_{xx}^2). \quad (9.41)$$

Recall also that the correlation coefficient for any linear decoding is independent of the actual decoding, and is the same as the correlation coefficient for no decoding. Hence, for the SSR model we have, for a Gaussian source,

$$R(D) = R_L(D) > -0.5 \log_2 (1 - \rho_{xy}^2). \quad (9.42)$$

Bounds for the mutual information in terms of correlation coefficient, or coherence function, are well known in the literature (Pinsker 1964, Nikitin and Stocks 2004).

We will now consider numerically the relationship between the Shannon lower bound, and the actual mutual information for SSR and various signal and noise distributions.

9.3.2 Rate-Distortion Tradeoff for SSR


We aim to examine the *operational* rate-distortion performance of the SSR model, and compare it to the theoretical Shannon lower bound for the sources given in Table 9.1, as well as the numerically calculated Shannon lower bound for the logistic source.

Operational rate-distortion

Recall that Chapter 4 provides figures showing numerical calculations of the mutual information in the SSR model as a function of noise intensity, σ , for various matched signal and noise distributions. Chapter 6 provides similar figures for the MMSE distortion. Therefore, plotting the value of $I(x, y)$ against the corresponding value of MMSE distortion for each value of σ gives provides a plot of the operational rate-distortion tradeoff for each value of σ in the SSR model. The result of carrying this out for $0 \leq \sigma \leq 5$ is shown in Figs. 9.7 and 9.8 for various values of N , and the cases of matched Gaussian, Laplacian and logistic signal and noise. Fig. 9.7 shows the mutual information plotted against the MMSE distortion while Fig. 9.8 shows the mutual information plotted against the output SQNR. In quantisation theory, it is quite often the case that plots of rate against distortion are shown in terms of SQNR in dB, due to the loglinear nature of the Shannon lower bound. In each figure, points on the curve for

each value of N correspond to a particular value of σ . The thick black lines show the theoretical Shannon lower bound, $R_L(D)$.

For each distribution, the curves for each value of $N > 1$ start at a rate of one bit per sample, then increase with rate and decrease with MMSE distortion, as σ increases. The rate then reaches its maximum before the MMSE distortion reaches its minimum. Then with continuing increasing σ , the curves reach the MMSE distortion minimum, before curling back down towards the $R_L(D)$ curve. Note that this means that—except for large σ —there are two values of σ for which the same distortion can occur, corresponding to two different rates. This is due to the SR behaviour seen in both the mutual information and distortion. If the main goal of a quantiser is to operate with minimum distortion, this observation indicates that the optimal value of input SNR to use is the one which achieves the minimum distortion, rather than the maximum rate. A further observation is the fact that for very large σ —that is, low input SNR—the SSR operational rate-distortion is very close to the Shannon lower bound.



(a) Gaussian signal and noise (b) Laplacian signal and noise (c) Logistic signal and noise

Figure 9.7. Operational rate-distortion tradeoff for SSR. This figure shows plots of the mutual information against MMSE distortion for three matched signal and noise distributions, and a number of values of N . The thick black line shows the Shannon lower bound, $R_L(D)$. Points on the curve for each value of N correspond to different values of noise intensity, σ , where σ starts at zero at the indicated point. Note that there are in general two values of mutual information that achieve the same MMSE distortion and two values of MMSE distortion that achieve the same mutual information. This is due to the SR behavior of the mutual information and MMSE distortion for increasing σ .

Comparison of $I(x, y)$ with $R_L(D)$

Fig. 9.9 shows (with thick lines) the mutual information for each distribution plotted against increasing σ , as in Chapter 4. It also shows (with thin lines) the Shannon lower

9.3 Rate-Distortion Tradeoff




Figure 9.8. Operational rate-distortion tradeoff for SSR, in dB. This figure shows plots of the mutual information against SQNR in dB for three matched signal and noise distributions, and a number of values of N . The thick black line shows the Shannon lower bound, $R_L(D)$, converted to dB. Points on the curve for each value of N correspond to different values of noise intensity, σ , where σ starts at zero at the indicated point.

bound for each value of σ , where the distortion value used for each σ , say $D(\sigma)$, is the operational distortion, obtained with optimal MMSE distortion decoding. The difference between the two curves for each value of N is therefore the difference in rate between the actual rate that achieves $D(\sigma)$ and the Shannon lower bound for the rate that achieves $D(\sigma)$. As is also seen in Figs. 9.7 and 9.8, it is clear that for large σ , the actual mutual information gets closer and closer to the Shannon lower bound.

Exact result for uniform signal and noise and $\sigma \leq 1$

In Chapter 4, we stated an exact result for the SSR mutual information and $\sigma \leq 1$, in terms of N and σ . This result, given in Eqn. (4.58) is first derived in Stocks (2001c). In Chapter 6 we derived a new result for the MMSE distortion under the same conditions, as given by Eqn. (6.100). Therefore, we have enough information to calculate analytical operational rate-distortion points in these circumstances. However, there is no simple way to combine these two expressions to obtain an expression for the mutual information in terms of MMSE distortion. The best we can do is calculate from each equation the mutual information and distortion for a range of values of σ , and then plot the resultant operational rate distortion curve, as shown for other distributions in Fig. 9.7. Since we can also numerically calculate the rate and distortion for $\sigma > 1$, we will not restrict our attention to the exact formulas, but plot in Fig. 9.10 the mutual information against rate for $0 \leq \sigma \leq 5$.




Figure 9.9. Comparison of $I(x,y)$ with $R_L(D)$ for SSR. This figure shows the operational mutual information for SSR (thick lines), as well as the Shannon-lower bound (thin lines) corresponding to the minimum achievable distortion for each value of σ . Clearly, the actual mutual information is far larger than the lower bound for small σ , but gets closer to the bound as σ increases.




Figure 9.10. Rate-distortion tradeoff for uniform signal and noise. This figure shows the operational rate-distortion tradeoff for SSR for uniform signal and noise and various N . Fig. 9.10(a) shows the mutual information plotted against MMSE distortion, while Fig. 9.10(b) shows the mutual information plotted against output SQNR. As with Figs. 9.7 and 9.8, points on the curve for each value of N correspond to different values of noise intensity, σ , where σ starts at zero at the indicated point, and the thick black line shows the Shannon lower bound, $R_L(D)$. Fig. 9.10(c) shows both the operational mutual information, $I(x,y)$ (thick lines) and the Shannon lower bound, $R_L(D)$, (thin lines) plotted against increasing noise intensity, σ .

9.3 Rate-Distortion Tradeoff

For uniform signal and noise and $\sigma = 1$, we saw in Chapter 6 that MMSE decoding has linearly spaced reconstruction points and is therefore a linear decoding scheme. This means that the optimal correlation coefficient is the linear correlation coefficient of Eqn. (6.64) evaluated at $\sigma = 1$, which is

$$\rho_{xy} = \sqrt{\left(\frac{N}{N+2}\right)}. \quad (9.43)$$

Substituting Eqn. (9.43) into Eqn. (9.39) gives

$$\text{SQNR} = \frac{\sigma_x^2}{12D} = \frac{N+2}{2}, \quad (9.44)$$

and therefore

$$\frac{\sigma_x^2}{D} = 6(N+2), \quad (9.45)$$

and we have the Shannon lower bound on rate in terms of the minimum possible distortion at $\sigma = 1$ as

$$R_L(D) = 0.5 \log_2 \left(\frac{3(N+2)}{\pi e} \right). \quad (9.46)$$

We have also from Chapter 5 a large N approximation to the mutual information at $\sigma = 1$ as

$$I(x, y) \simeq 0.5 \log_2 \left(\frac{(N+2)e}{2\pi} \right). \quad (9.47)$$

Therefore, for large N , and uniform signal and noise with $\sigma = 1$ we have

$$I(x, y) - R_L(D) \simeq 0.5 \log_2 \left(\frac{\exp(2)}{6} \right) \simeq 0.15. \quad (9.48)$$

Thus, the mutual information for SSR is always at least approximately 0.15 bits per sample larger than the Shannon lower bound. This difference is clearly visible in Fig. 9.10(c).

9.3.3 Rate-Distortion Tradeoff for Optimised Thresholds

The operational rate-distortion trade-off for the case of optimised thresholds, $N = 5$ and Gaussian signal and noise—as discussed in Chapter 8—is shown in Figure 9.11. Both the cases of maximised mutual information and minimised MMSE distortion are shown, as well as the SSR case. This plot clearly shows that when the MMSE distortion is minimised, the corresponding mutual information is smaller than that obtained by maximising the mutual information. Conversely, maximising the mutual information

results in a larger distortion. When the distortion is minimised, the resultant curve is relatively close to the Shannon lower bound, $R(D)$, but as we might expect, given the SSR results shown previously, does not reach it. This plot also clearly illustrates that the performance of SSR at best reaches about half the optimal mutual information, but only about five times the minimum possible distortion.




Figure 9.11. Operational rate-distortion tradeoff for arbitrary thresholds. This figure shows the operational rate-distortion tradeoff for $N = 5$ and Gaussian signal and noise, for optimised thresholds. Both the cases of maximised mutual information and minimum MMSE distortion are shown, with the SSR case also plotted for comparison. Points on the curve for each situation correspond to different values of noise intensity, σ , where σ starts at zero at the indicated point. The thick solid line shows the Shannon lower bound, $R_L(D)$. It is clear that optimising the thresholds provides a much larger mutual information, and smaller MMSE distortion than is achievable with SSR.

9.4 Chapter Summary

The introductory Section of this Chapter briefly discusses the need for tradeoffs between performance and cost in the design of engineered systems. It also reviews the theory from previous Chapters relevant to this Chapter, and discusses the Lagrange multiplier method for constrained optimisation.

Section 9.2 introduces the problem of energy efficient information transfer to the SSR model and its extension to arbitrary thresholds. We define measures of energy, and energy efficiency, and discuss how these measures vary with the threshold value, θ , and noise intensity, σ , in the SSR model. We then formulate and solve a constrained extension to the optimal quantisation problem first discussed in Chapter 8, and show how

the same qualitative behaviour occurs in the optimal thresholds, but with an increase in the mean threshold value with increasing σ . For sufficiently large σ , the situation of all thresholds equal to the same value is optimal. We also briefly discuss a related problem of minimising the number of threshold devices required to achieve a certain mutual information constraint in the SSR model.

Section 9.3 begins by a discussion of the tradeoff between rate and distortion in a quantiser, and then formalising this discussion by stating relevant theory. We then examine the operational rate-distortion tradeoff in the SSR model, and compare the performance achieved with known lower bounds. We find that for large σ , the operational rate for a given distortion is relatively close to the Shannon lower bound, when compared with small σ .

9.4.1 Original Contributions for Chapter 9

This chapter includes the following original contributions:

- The introduction of energy constraints to the SSR model and the arbitrary threshold extension of the SSR model, and formulation of these constraints into constrained optimisation problems.
- Solutions of an energy constrained stochastic quantisation problem.
- Discussion of the operational rate-distortion tradeoff for the SSR model.

9.4.2 Further Work

Possible future work and open questions arising from this chapter might include:

- Incorporation of energy constraints into optimal stochastic quantisation of arrays of more realistic neuron models, such as the FitzHugh-Nagumo model considered in Stocks and Mannella (2001).
- Consideration of arbitrary thresholds for the problem of minimising N for a given minimum mutual information, or distortion constraint.
- Comparison of the operational rate-distortion tradeoff for the SSR model, and optimally stochastic quantised model, with the operational tradeoff found in conventional scalar quantisers. The conventional quantiser can be considered both

with and without independent threshold noise, and with and without companding.

- Introduction of a constraint on output entropy for the optimal stochastic quantisation model. Such a problem has often been considered in quantisation theory (György and Linder 2000).

This concludes Chapter 9, which considers various tradeoffs between cost and performance in the SSR model, and its extension to arbitrary thresholds. This completes the major chapters of this thesis. The final Chapter makes some general conclusions, recommendations for future work, and summarises the major original contributions of this thesis.

Chapter 10

Conclusions and Future Directions

TO conclude this thesis, we summarise the main results and conclusions, before briefly summarising the most promising areas for future research.

10.1 Thesis Summary

The following subsections summarise the original work contained in this thesis, by listing the most significant contributions. A full list of original contributions is given at the end of each Chapter.

10.1.1 Stochastic Resonance

Chapter 2 presents an original historical review and elucidation of the major epochs in the history of Stochastic Resonance (SR) research, and discussion of the evolution of the term ‘stochastic resonance’.

Chapter 2 also demonstrates qualitatively that SR can actually occur in a single threshold device, where the threshold is set to the signal mean. Although SR cannot occur in the conventional signal-to-noise ratio (SNR) measure in this situation, if ensemble averaging is allowed, then the presence of an optimal noise level can decrease distortion.

Chapter 3 contains an extended discussion and critique of the use of SNR measures to quantify SR, the debate about SNR gains due to SR, and the relationship between SNRs and information theory.

10.1.2 Suprathreshold Stochastic Resonance

Chapter 4 provides an up-to-date literature review of previous work on Suprathreshold Stochastic Resonance (SSR). It also gives numerical results showing SSR occurring for a number of matched and mixed signal and noise distributions not previously considered. A generic change of probability measure in the equations used to determine the mutual information through the SSR model is introduced. This change of probability measure results in a Probability Density Function (PDF) that describes the average transfer function of the SSR model. This PDF is derived for several specific cases, for which it is proved that the mutual information is a function of the noise intensity parameter, σ , rather than a function of both the noise variance, and signal variance independently.

Chapter 5 both improves on previous results and derives several new large N approximations to the mutual information, output entropy, average conditional output entropy, output distribution, and channel capacity in the SSR model. An expression for a

channel capacity achieving input PDF for any given noise PDF is found, which holds under the conditions for which the large N mutual approximation formula holds. This formula gives an upper bound for the achievable channel capacity for SSR.

10.1.3 Stochastic Quantisation with Identical Thresholds

Chapter 6 introduces and applies the concept of decoding to the SSR model, in order to provide a signal that approximates, or reconstructs, the input signal. The noise in such a reconstruction can be measured with the Mean Square Error (MSE) distortion, or equivalently, the signal-to-quantisation-noise ratio. Plots are presented of the MSE distortion and reconstruction points for a number of linear and nonlinear decoding schemes applied to the SSR model, and it is shown that SR occurs in the MSE distortion measure.

Chapter 6 also gives an analytical derivation of an expression for the Minimum Mean Square Error (MMSE) distortion and optimal reconstruction points for the SSR model for the case of uniform signal and noise, with noise intensities, $\sigma \leq 1$.

The *information bound* is applied to find a lower bound, named here as the *average information bound*, on the MSE distortion for decoding of the SSR model. It is shown that a linear decoding scheme for the SSR model is a biased efficient estimator, and that MMSE decoding is biased and does not meet the average information bound with equality. Analysis of the average information bound confirms that the value of noise intensity, σ , which minimises the MSE distortion means finding the best tradeoff between average error variance and mean square bias. This tradeoff is analogous to the tradeoff between output entropy and average conditional output entropy required to maximise the mutual information.

Chapter 7 derives large N expressions for the optimal linear MSE distortion, reconstruction points and linear correlation coefficient. It also demonstrates that the MSE distortion cannot be made asymptotically small for small noise intensities in the SSR model, even for infinite N .

10.1.4 Optimal Stochastic Quantisation

Chapter 8 extends the SSR model to allow all threshold devices to have arbitrary threshold values. In order to calculate the transition probabilities for this extension, we

10.2 Closing Remarks

derive a very general recursive formula, with $O(N^2)$ computational complexity, that makes it straightforward to calculate the transition probabilities for any given noise distribution, and threshold values.

We then mathematically formulate the problem of optimal stochastic quantisation of the array of threshold devices, in terms of the vector of optimal thresholds, θ . The solution to this problem resulted in the discovery of an unexpected bifurcation pattern in the optimal threshold values, as the noise intensity increases. At the same time, this solution numerically validates that the SSR situation of all thresholds equal to the signal mean is in fact optimal for sufficiently large noise intensity, for a range of signal and noise distributions, and both the mutual information and MSE distortion measures. Furthermore, we validate that SSR remains optimal for sufficiently large noise intensity, even if N becomes very large.

Also in Chapter 8, we derive an approximation to the Fisher information for arbitrary thresholds, and apply this approximation to find a lower bound on the MSE distortion. Optimally setting the thresholds was shown to be related to finding the optimal trade-off between the two components of this lower bound, the mean square bias, and the average error variance.

10.1.5 Stochastic Quantiser Performance Tradeoffs

In Chapter 9, we introduce energy constraints to the SSR model and the arbitrary threshold extension of the SSR model, and formulate these constraints into constrained optimisation problems. Solutions of such an energy constrained stochastic quantisation problem show similar bifurcations to the results of Chapter 8.

Chapter 9 also discusses for the first time, the operational rate-distortion tradeoff for the SSR model.

10.2 Closing Remarks

Specific suggestions for further research that may be of interest are given at the end of each Chapter. However, in general there are two directions in which there are many challenges and potential for new results based on this research.

Firstly, for accomplished mathematicians, there is potential for many of the mathematical results presented to be placed in a more rigorous and general mathematical

platform. Furthermore, many of the results given here were obtained by numerical methods alone, and it is the author's belief that the application of advanced mathematical approaches will find analytical explanations or predictions in some cases. In particular, it may be possible to further explain the complex bifurcation patterns plotted in Chapters 8-9, and it may be possible to predict the location of the bifurcations.

Secondly, in the other direction, the challenge for engineers is to design circuits or quantisation schemes based on the theoretical results contained in this thesis. Such work may also require extensions to theory, such as consideration of non-ideal threshold devices (Martorell *et al.* 2004). We already know that DIMUS sonar arrays made use of stochastic quantisation in the 1960s. Furthermore, SSR has been proposed as a means of improving cochlear implant encoding (Stocks *et al.* 2002). Given the current trend towards arrays of small, low-power and low-cost networks of sensors (Pradhan *et al.* 2002), it is feasible that a form of distributed quantisation that utilises ambient noise will find new applications in the near future.

Appendix A

Binary Channel Calculations

A.1 Example 1: Asymmetric Binary Channel

This Section gives derivations required for the example given by Eqn. (3.6) in Section 3.3.2 of Chapter 3. In particular, we derive the fact that $I(s, x) = 1$ bit per sample, and also derive Eqn. (3.8) giving $I(s, y)$.

Joint and conditional probability functions of s and y

The conditional probability function of y given s can be written as

$$p(y|s) = (1 - 2y)s \left(\frac{t-1}{b} - 1 \right) + (1 - y). \quad (\text{A.1})$$

The probability distribution of y can be found by observing that $y = 0$ when $s = 0$ or when $s = 1$ and $n < t - 1$, and that $y = 1$ when $s = 1$ and $n \geq t - 1$. Thus

$$P(y) = 0.5 + 0.5(1 - 2y) \left(\frac{t-1}{b} \right). \quad (\text{A.2})$$

Now the joint distribution of y and s is $p(s, y) = p(s)p(y|s)$. Therefore

$$p(s, y) = 0.5s \left((1 - 2y) \left(\frac{t-1}{b} - 1 \right) \right) + 0.5(1 - y). \quad (\text{A.3})$$

From this joint distribution, it is possible to calculate the probability of an error at the output as $P_e = p(s = 0, y = 1) + p(s = 1, y = 0) = \frac{t-1}{2b}$.

Joint and conditional probability functions of s and x

When $b < 1$, x can never take on values between b and 1. When $b \geq 1$, x can take on any value between 0 and $1 + b$. We restrict our attention to $b < 1$. Note that b has been

A.1 Example 1: Asymmetric Binary Channel

defined such that $t - 1 \leq b < t$, and that $t > 1$. Hence $b > 0$ always, and when $b < 1$ then $t < 2$. Therefore $0 < t - 1 < 1$.

Since $x = s + n$, then $p(x|s) = p(s + n|s) = p_n(x - s)$. Hence

$$p(x|s) = \begin{cases} \frac{1}{b} & s \leq x \leq b + s, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A.4})$$

The joint probability density of s and x is given by $p(s, x) = p(s)p(x|s)$. Hence

$$p(s, x) = \begin{cases} \frac{1}{2b} & s \leq x \leq b + s, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A.5})$$

The probability density function of x is given by

$$p_x(x) = \sum_s p(x|s)p(s) = \begin{cases} \frac{1}{2b} & 0 \leq x \leq b \text{ and } 1 \leq x \leq b + 1, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A.6})$$

Mutual information between s and x

The results in the previous two subsections allow us to write an expression for $I(s, x)$. Recall that $b < 1$. The entropy of x is then

$$\begin{aligned} H(x) &= - \int p_x(x) \log_2 p_x(x) dx \\ &= - \int_0^b \frac{1}{2b} \log_2 \frac{1}{2b} dx - \int_1^{b+1} \frac{1}{2b} \log_2 \frac{1}{2b} dx \\ &= 1 + \log_2 b. \end{aligned} \quad (\text{A.7})$$

This makes sense, as the entropy of a uniform distribution of width c is $\log_2 c$ and the entropy of s is 1. The entropy of y given s is

$$\begin{aligned} H(x|s) &= - \sum_s \int P(s, x) \log_2 P(x|s) dx \\ &= - \sum_s \int_s^{b+s} \frac{1}{2b} \log_2 \frac{1}{b} dx \\ &= - \int_0^b \frac{1}{2b} \log_2 \frac{1}{b} dx + \int_1^{b+1} \frac{1}{2b} \log_2 \frac{1}{b} dx \\ &= \log_2 b. \end{aligned} \quad (\text{A.8})$$

This indicates that $H(x|s)$ is simply the entropy of the noise, which should be expected, since s is known, and both values of s are equally probable. Thus, the mutual information between s and x is $I(s, x) = H(x) - H(x|s) = 1$ bit per sample.

Mutual information between s and y

The entropy of y is

$$\begin{aligned} H(y) &= - \sum_y p_y(y) \log_2 p_y(y) \\ &= - \sum_y \left(0.5 + 0.5(1 - 2y) \left(\frac{t-1}{b} \right) \right) \log_2 \left(0.5 + 0.5(1 - 2y) \left(\frac{t-1}{b} \right) \right). \end{aligned} \quad (\text{A.9})$$

The entropy of y given s is

$$\begin{aligned} H(y|s) &= - \sum_s \sum_y P(s, y) \log_2 P(y|s) \\ &= - \left(\frac{t-1}{2b} \right) \log_2 \left(\frac{t-1}{b} \right) - \left(\frac{b-t+1}{2b} \right) \log_2 \left(\frac{b-t+1}{b} \right). \end{aligned} \quad (\text{A.10})$$

Recall that the probability of error is $P_e = \frac{t-1}{2b}$. Then

$$H(y) = -(0.5 + P_e) \log_2 (0.5 + P_e) - (0.5 - P_e) \log_2 (0.5 - P_e) \quad (\text{A.11})$$

and

$$\begin{aligned} H(y|s) &= - P_e \log_2 2P_e - (0.5 - P_e) \log_2 (1 - 2P_e) \\ &= - 0.5 - P_e \log_2 P_e - (0.5 - P_e) \log_2 (0.5 - P_e). \end{aligned} \quad (\text{A.12})$$

Eqns. (A.11) and (A.12) are in exact agreement with Eqns. (7) and (8) of Chapeau-Blondeau (1997b) for this specific case. Therefore

$$\begin{aligned} I(s, y) &= H(y) - H(y|s) \\ &= 0.5 + P_e \log_2 P_e - (0.5 + P_e) \log_2 (0.5 + P_e) \quad \text{bits per sample.} \end{aligned} \quad (\text{A.13})$$

A.2 Example 2: Chapeau-Blondeau's Erasure Channel

This Section shows how we derive the formula for $I(s, x)$ given by Eqn. (3.14) in Section 3.3.3 of Chapter 3.

The PDF of x depends on the value of σ_n . When $\sigma_n < 2$, s is retrievable without error, using a threshold at 0. In this case, the mutual information between s and x is the entropy of the input signal, s . However, when $\sigma_n \geq 2$, there exists a range of values $s_v - \sigma_n/2 \leq x \leq -s_v + \sigma_n/2$ for which it is impossible to tell whether s is $-s_v$ or s_v .

A.2 Example 2: Chapeau-Blondeau's Erasure Channel

Assume that $p(\pm s_v) = 0.5$. Thus for $\sigma < 2$, the PDF of x is

$$p_x(x) = \sum_s p(x|s)p(s) = \begin{cases} \frac{1}{2\sigma_b} & -1 - \frac{\sigma_n}{2} \leq x \leq -1 + \frac{\sigma_n}{2} \text{ and } 1 - \frac{\sigma_n}{2} \leq x \leq 1 + \frac{\sigma_n}{2}, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A.14})$$

For $\sigma_n \geq 2$,

$$p_x(x) = \begin{cases} \frac{1}{2\sigma_b} & -1 - \frac{\sigma_n}{2} \leq x < 1 - \frac{\sigma_n}{2} \text{ and } -1 + \frac{\sigma_n}{2} < x \leq 1 + \frac{\sigma_n}{2}, \\ \frac{1}{\sigma_b} & 1 - \frac{\sigma_n}{2} \leq x \leq -1 + \frac{\sigma_n}{2}, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A.15})$$

Using Eqns. (A.14) and (A.15) the entropy of x is

$$H(x) = - \int p_x(x) \log_2 P_x(x) dx = \begin{cases} 1 + \log_2 \sigma_n & \sigma_n < 2, \\ \frac{2}{\sigma_n} + \log_2 \sigma_n & \sigma_n \geq 2. \end{cases}$$

Note that as σ_n becomes large, $H(x) \rightarrow \log_2 \sigma_n$, which is the entropy of a uniform distribution. As expected, this indicates that for very large noise x is very unlike the original binary signal, which has an entropy of 1 bit per sample.

As for the binary asymmetric channel, the entropy of x given s is simply the entropy of the uniform noise—see Section A.1. Thus $H(x|s) = \log_2 \sigma_n$. Hence, the mutual information between s and x is

$$I(s, x) = H(x) - H(x|s) = \begin{cases} 1 & \sigma_n < 2, \\ \frac{2}{\sigma_n} & \sigma_n \geq 2. \end{cases} \quad (\text{A.16})$$

Appendix B

Derivations for Suprathreshold Stochastic Resonance Encoding

B.1 Maximum Values and Modes of $P(n|x)$

This Section derives Eqns. (4.11) and (4.12), which give the values of the modes of $P(n|x)$ for SSR, in terms of the Inverse Cumulative Distribution Function (ICDF) of the noise, and are stated in Section 4.3 of Chapter 4.

To begin, the derivative of the transition probabilities, $P(n|x)$, can be expressed in terms of $P_{1|x}$ as

$$\begin{aligned}\frac{dP(n|x)}{dx} &= P(n|x) \left(\frac{n - NP_{1|x}}{P_{1|x}(1 - P_{1|x})} \right) \frac{dP_{1|x}}{dx} \\ &= P(n|x) \left(\frac{n - NP_{1|x}}{P_{1|x}(1 - P_{1|x})} \right) R_\eta(x - \theta) \quad n = 0, \dots, N.\end{aligned}\quad (\text{B.1})$$

Note that for $n = 0$ and $n = N$, $P(0|x) = (1 - P_{1|x})^N$ and $P(N|x) = P_{1|x}^N$.

Setting the derivative to zero in Eqn. (B.1) gives

$$P_{1|x} = \frac{n}{N} \quad n = 1, \dots, N - 1.$$

For $n = 0$, $P(0|x)$ is maximised when $P_{1|x} = 0$, which means at the minimum possible value of x , and for $n = N$, $P(N|x)$ is maximised when $P_{1|x} = 1$, which means at the maximum possible value of x .

Therefore, for all n , the maximum of $P(n|x)$ occurs when

$$1 - F_R(\theta - x) = \frac{n}{N} \quad n = 0, \dots, N,$$

B.2 A Proof of Equation (4.38)

which implies that

$$x = \theta + F_R^{-1} \left(1 - \frac{n}{N} \right) \quad n = 0, \dots, N, \quad (\text{B.2})$$

where $F_R^{-1}(\cdot)$ is the ICDF of the noise. For even noise Probability Density Functions (PDFs),

$$x = \theta + F_R^{-1} \left(\frac{n}{N} \right) \quad n = 0, \dots, N. \quad (\text{B.3})$$

Note that differentiating the natural logarithm of $P(n|x)$ —that is, the log-likelihood function—rather than $P(n|x)$ by itself, gives the same result.

The value of x at which the maximum of $P(n|x)$ occurs is known as the *mode* of $P(n|x)$ for each n . Note that if the noise distribution has a PDF with infinite support, then the mode of $P(0|x)$ and $P(N|x)$ is at $x = \pm\infty$.

B.2 A Proof of Equation (4.38)

This section proves the identity given by Eqn. (4.38) in Section 4.3.2 of Chapter 4,

$$-\sum_{n=0}^N \log_2 \binom{N}{n} = \sum_{n=1}^N (N+1-2n) \log_2 n. \quad (\text{B.4})$$

The proof is

$$\begin{aligned} \text{LHS} &= -\sum_{n=0}^N \log_2 \binom{N}{n} \\ &= \sum_{n=0}^N \log_2 \frac{n!(N-n)!}{N!} \\ &= \sum_{n=1}^N \log_2 n! + \sum_{n=1}^N \log_2 (N-n)! - \sum_{n=1}^N \log_2 N! \\ &= \sum_{n=1}^N \sum_{j=1}^n \log_2 j + \sum_{n=1}^N \sum_{j=1}^{N-n} \log_2 j - \sum_{n=1}^N \sum_{j=1}^N \log_2 j \\ &= \sum_{k=1}^N (N-k+1) \log_2 k + \sum_{k=1}^N (N-k) \log_2 k - N \sum_{k=1}^N \log_2 k \\ &= \sum_{n=1}^N (N-2n+1) \log_2 n \\ &= \text{RHS}. \end{aligned}$$

B.3 Distributions

This Section states the PDF, CDF and ICDF for each of seven different continuously valued probability distributions. It also uses the CDF to state the probability that a single threshold device is 'on', $P_{1|x}$, for threshold value, θ , and uses the ICDF to state the modes of $P(n|x)$. Some of these distributions are used throughout this thesis, but are first used in Section 4.3.3 in Chapter 4.

B.3.1 Gaussian Signal and Noise

If the input signal has a Gaussian distribution with zero mean and variance σ_x^2 , then

$$P(x) = \frac{1}{\sqrt{2\pi\sigma_x^2}} \exp\left(-\frac{x^2}{2\sigma_x^2}\right). \quad (\text{B.5})$$

If the independent noise in each device is Gaussian with zero mean and variance σ_η^2 , then

$$R(\eta) = \frac{1}{\sqrt{2\pi\sigma_\eta^2}} \exp\left(-\frac{\eta^2}{2\sigma_\eta^2}\right). \quad (\text{B.6})$$

The Cumulative Distribution Function (CDF) of the noise evaluated at $\eta = z$ is

$$F_R(z) = 0.5 + 0.5\text{erf}\left(\frac{z}{\sqrt{2}\sigma_\eta}\right), \quad (\text{B.7})$$

where $\text{erf}(\cdot)$ is the error function (Spiegel and Liu 1999). Therefore,

$$P_{1|x} = 0.5 + 0.5\text{erf}\left(\frac{x - \theta}{\sqrt{2}\sigma_\eta}\right). \quad (\text{B.8})$$

The ICDF of the Gaussian noise is

$$F_R^{-1}(w) = \sqrt{2}\sigma_\eta \text{erf}^{-1}(2w - 1), \quad (\text{B.9})$$

where $w \in [0, 1]$, and $\text{erf}^{-1}(\cdot)$ is the inverse error function. From Eqn. (4.12), the values of x at which the maximum of each $P(n|x)$ occurs—that is, the mode of $P(n|x)$, for each n —are

$$x = \theta + \sqrt{2}\sigma_\eta \text{erf}^{-1}\left(\frac{2n}{N} - 1\right) \quad n = 1, \dots, N - 1. \quad (\text{B.10})$$

For $n = 0$ and $n = N$, the mode is at $x = \mp\infty$ respectively.

B.3.2 Uniform Signal and Noise

If the input signal, x is uniformly distributed between $-\sigma_x/2$ and $\sigma_x/2$, with zero mean, then

$$P(x) = \begin{cases} 1/\sigma_x & \text{for } -\sigma_x/2 \leq x \leq \sigma_x/2, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{B.11})$$

If the independent noise η in each device is uniformly distributed between $-\sigma_\eta/2$ and $\sigma_\eta/2$, with zero mean, then

$$R(\eta) = \begin{cases} 1/\sigma_\eta & \text{for } -\sigma_\eta/2 \leq \eta \leq \sigma_\eta/2, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{B.12})$$

The CDF of the noise evaluated at $\eta = z$ is

$$F_R(z) = \begin{cases} 0 & \text{for } z < -\sigma_\eta/2, \\ z/\sigma_\eta + 1/2 & \text{for } -\sigma_\eta/2 \leq z \leq \sigma_\eta/2, \\ 1 & \text{for } z > \sigma_\eta/2. \end{cases} \quad (\text{B.13})$$

Therefore

$$P_{1|x} = \begin{cases} 0 & \text{for } x < \theta - \sigma_\eta/2, \\ x/\sigma_\eta + 1/2 - \theta/\sigma_\eta & \text{for } \theta - \sigma_\eta/2 \leq x \leq \theta + \sigma_\eta/2, \\ 1 & \text{for } x > \theta + \sigma_\eta/2. \end{cases} \quad (\text{B.14})$$

The ICDF of the uniform noise is

$$F_R^{-1}(w) = \sigma_\eta(w - 0.5), \quad (\text{B.15})$$

where $w \in [0, 1]$. From Eqn. (4.12), the mode of each $P(n|x)$ is

$$x = \theta + \sigma_\eta \left(\frac{n}{N} - \frac{1}{2} \right) \quad n = 0, \dots, N. \quad (\text{B.16})$$

B.3.3 Laplacian Signal and Noise

If the input signal x has a Laplacian distribution with zero mean and variance σ_x^2 then

$$P(x) = \frac{1}{\sqrt{2}\sigma_x} \exp\left(\frac{-\sqrt{2}|x|}{\sigma_x}\right). \quad (\text{B.17})$$

If the independent noise η in each device has a Laplacian distribution with zero mean, and variance σ_η^2 then

$$R(\eta) = \frac{1}{\sqrt{2}\sigma_\eta} \exp\left(\frac{-\sqrt{2}|\eta|}{\sigma_\eta}\right). \quad (\text{B.18})$$

The CDF of the noise evaluated at $\eta = z$ is

$$F_R(z) = 0.5 \left(1 + \text{sign}(z) \left(1 - \exp \left(\frac{-\sqrt{2}|z|}{\sigma_\eta} \right) \right) \right), \quad (\text{B.19})$$

where $\text{sign}(\cdot)$ indicates the signum (sign) function. Therefore we have

$$P_{1|x} = 0.5 \left(1 + \text{sign}(x - \theta) \left(1 - \exp \left(\frac{-\sqrt{2}|x - \theta|}{\sigma_\eta} \right) \right) \right). \quad (\text{B.20})$$

It is clearer to write this as

$$P_{1|x} = \begin{cases} 0.5 \exp \left(\frac{-\sqrt{2}\theta}{\sigma_\eta} \right) \exp \left(\frac{\sqrt{2}x}{\sigma_\eta} \right) & \text{for } x \leq \theta, \\ 1 - 0.5 \exp \left(\frac{\sqrt{2}\theta}{\sigma_\eta} \right) \exp \left(\frac{-\sqrt{2}x}{\sigma_\eta} \right) & \text{for } x \geq \theta. \end{cases} \quad (\text{B.21})$$

The ICDF of the Laplacian noise is

$$F_R^{-1}(w) = \begin{cases} \frac{\sigma_\eta}{\sqrt{2}} \ln(2w) & \text{for } w \in [0, 0.5], \\ -\frac{\sigma_\eta}{\sqrt{2}} \ln(2(1-w)) & \text{for } w \in [0.5, 1]. \end{cases} \quad (\text{B.22})$$

From Eqn. (4.12), the mode of each $P(n|x)$ is

$$x = \begin{cases} \theta + \frac{\sigma_\eta}{\sqrt{2}} \ln \left(\frac{2n}{N} \right) & \text{for } 0 < \frac{n}{N} \leq 0.5, \\ \theta - \frac{\sigma_\eta}{\sqrt{2}} \ln \left(2 \left(1 - \frac{n}{N} \right) \right) & \text{for } 0.5 \leq \frac{n}{N} < 1. \end{cases} \quad (\text{B.23})$$

For $n = 0$ and $n = N$, the mode is at $x = \mp\infty$ respectively.

B.3.4 Logistic Signal and Noise

If the input signal x has a logistic distribution with zero mean and variance σ_x^2 then

$$P(x) = \frac{\exp \left(-\frac{x}{b_x} \right)}{b_x \left(1 + \exp \left(-\frac{x}{b_x} \right) \right)^2}, \quad (\text{B.24})$$

where $\sigma_x^2 = \frac{\pi^2 b_x^2}{3}$.

Note that this distribution can also be written in terms of the hyperbolic cosine function as

$$P(x) = \frac{1}{4b_x \cosh^2 \left(\frac{x}{2b} \right)}. \quad (\text{B.25})$$

B.3 Distributions

If the independent noise η in each device has a logistic distribution with zero mean, and variance σ_η^2 then

$$R(\eta) = \frac{\exp\left(-\frac{\eta}{b_\eta}\right)}{b_\eta \left(1 + \exp\left(-\frac{\eta}{b_\eta}\right)\right)^2}. \quad (\text{B.26})$$

where $\sigma_\eta^2 = \frac{\pi^2 b_\eta^2}{3}$.

The CDF of the noise evaluated at $\eta = z$ is

$$F_R(z) = \frac{1}{1 + \exp\left(-\frac{z}{b_\eta}\right)}. \quad (\text{B.27})$$

Note that the logistic CDF is the solution to the logistic equation, that is, the differential equation

$$\frac{dF_R(z)}{dz} = \frac{1}{b_\eta} F_R(z)(1 - F_R(z)). \quad (\text{B.28})$$

Now

$$P_{1|x} = \frac{1}{1 + \exp\left(-\frac{(x-\theta)}{b_\eta}\right)}, \quad (\text{B.29})$$

and hence,

$$\frac{dP_{1|x}}{dx} = \frac{1}{b_\eta} P_{1|x}(1 - P_{1|x}). \quad (\text{B.30})$$

The ICDF of the logistic noise is

$$F_R^{-1}(w) = -b_\eta \ln\left(\frac{1-w}{w}\right), \quad (\text{B.31})$$

where $w \in [0, 1]$. From Eqn. (4.12), the mode of each $P(n|x)$ is

$$x = \theta + \frac{\sqrt{3}\sigma_\eta}{\pi} \ln\left(\frac{n}{N-n}\right) \quad n = 1, \dots, N-1. \quad (\text{B.32})$$

For $n = 0$ and $n = N$, the mode is at $x = \mp\infty$ respectively.

B.3.5 Cauchy Signal and Noise

If the input signal x has a Cauchy (or Lorentzian) distribution with zero mean and parameter $\lambda_x > 0$ then

$$P(x) = \frac{\lambda_x}{\pi} \frac{1}{\lambda_x^2 + x^2}. \quad (\text{B.33})$$

Unlike the other distributions above, the Cauchy distribution has undefined moments, and therefore is not characterised by its variance, but by the parameter λ_x , which is the PDF's *Full Width at Half-Maximum* (FWHM). The FWHM of a function is the distance between points on its curve at which the function reaches half its maximum value.

If the independent noise η in each device has a Cauchy distribution with zero mean, and FWHM λ_η then

$$R(\eta) = \frac{\lambda_\eta}{\pi} \frac{1}{\lambda_\eta^2 + x^2}. \quad (\text{B.34})$$

The CDF of the noise evaluated at $\eta = z$ is

$$F_R(z) = \frac{1}{2} + \frac{1}{\pi} \arctan\left(\frac{z}{\lambda_\eta}\right). \quad (\text{B.35})$$

Therefore we have

$$P_{1|x} = \frac{1}{2} + \frac{1}{\pi} \arctan\left(\frac{x - \theta}{\lambda_\eta}\right). \quad (\text{B.36})$$

The ICDF of the Cauchy noise is

$$F_R^{-1}(w) = \lambda_\eta \tan(\pi(w - 0.5)), \quad (\text{B.37})$$

where $w \in [0, 1]$. From Eqn. (4.12), the mode of each $P(n|x)$ is

$$x = \theta + \lambda_\eta \tan\left(\pi\left(\frac{n}{N} - 0.5\right)\right) \quad n = 1, \dots, N - 1. \quad (\text{B.38})$$

For $n = 0$ and $n = N$, the mode is at $x = \mp\infty$ respectively.

B.3.6 Exponential Signal and Noise

If the input signal $x (\geq 0)$ has an exponential distribution with mean σ_x , then

$$P(x) = \frac{1}{\sigma_x} \exp\left(-\frac{x}{\sigma_x}\right). \quad (\text{B.39})$$

Unlike the previous distributions, the exponential distribution is not an even function, or symmetric about its mean. It is only defined for values greater than or equal to zero.

If the independent noise $\eta (\geq 0)$ in each device has an exponential distribution with mean σ_η , then

$$R(\eta) = \frac{1}{\sigma_\eta} \exp\left(-\frac{\eta}{\sigma_\eta}\right). \quad (\text{B.40})$$

The CDF of the noise evaluated at $\eta = z$ is

$$F_R(z) = 1 - \exp(-z/\sigma_\eta). \quad (\text{B.41})$$

Therefore

$$P_{1|x} = \begin{cases} \exp\left(-\frac{(\theta-x)}{\sigma_\eta}\right) & \text{for } x < \theta, \\ 1 & \text{for } x \geq \theta. \end{cases}$$

The ICDF of the exponential noise is

$$F_R^{-1}(w) = \sigma_\eta \ln\left(\frac{1}{1-w}\right), \quad (\text{B.42})$$

where $w \in [0, 1]$.

B.3.7 Rayleigh Signal and Noise

If the input signal $x (\geq 0)$ has a Rayleigh distribution with mean $\sigma_x \sqrt{\pi/2}$, then

$$P(x) = \frac{x}{\sigma_x^2} \exp\left(-\frac{x^2}{2\sigma_x^2}\right). \quad (\text{B.43})$$

Like the exponential distribution, the Rayleigh distribution is not an even function or symmetric about its mean. It is only defined for values greater than or equal to zero.

If the independent noise $\eta (\geq 0)$ in each device has a Rayleigh distribution with mean $\sigma_\eta \sqrt{\pi/2}$, then

$$R(\eta) = \frac{\eta}{\sigma_\eta^2} \exp\left(-\frac{\eta^2}{2\sigma_\eta^2}\right). \quad (\text{B.44})$$

The CDF of the noise evaluated at $\eta = z$ is

$$F_R(z) = 1 - \exp\left(-\frac{z^2}{2\sigma_\eta^2}\right). \quad (\text{B.45})$$

Therefore

$$P_{1|x} = \begin{cases} \exp\left(-\frac{(\theta-x)^2}{2\sigma_\eta^2}\right) & \text{for } x < \theta, \\ 1 & \text{for } x \geq \theta. \end{cases}$$

The ICDF of the exponential noise is

$$F_R^{-1}(w) = \sigma_\eta \sqrt{2 \ln\left(\frac{1}{1-w}\right)}, \quad (\text{B.46})$$

where $w \in [0, 1]$.

B.4 Proofs for Specific Cases that $Q(\tau)$ is a PDF

This Section provides a proof that $Q(\tau)$ is a PDF for two of the infinite support signal and noise distributions given in Table 4.2 in Section 4.3.3 of Chapter 4. Since each $Q(\tau)$ is nonnegative, for $Q(\tau)$ to be a PDF it is only necessary to show that

$$\int_{\tau=0}^{\tau=1} Q(\tau) d\tau = 1.$$

B.4.1 Gaussian Signal and Noise

For Gaussian signal and noise, and $\sigma \neq 1$, the change of variable from τ to $u = \text{erf}^{-1}(2\tau - 1)$ is useful. The result is

$$\begin{aligned} \int_{\tau=0}^{\tau=1} Q(\tau) d\tau &= \int_{u=-\infty}^{u=\infty} \sigma \exp\left((1 - \sigma^2)u^2\right) \frac{1}{\sqrt{\pi}} \exp(-u^2) du \\ &= \frac{\sigma}{\sqrt{\pi}} \int_{u=-\infty}^{u=\infty} \exp(-\sigma^2 u^2) du \\ &= \frac{\sigma}{\sqrt{\pi}} \sqrt{2\pi} \frac{1}{2\sigma^2} \\ &= 1. \end{aligned}$$

Thus, given this result, $Q(\tau)$ is a PDF for $\sigma \neq 1$.

For the case of $\sigma = 1$, $Q(\tau) = 1 \forall \tau \in [0, 1]$, and is hence a PDF.

B.4.2 Laplacian Signal and Noise

For Laplacian signal and noise and $\sigma \neq 1$,

$$\begin{aligned} \int_{\tau=0}^{\tau=1} Q(\tau) d\tau &= \int_{\tau=0}^{\tau=0.5} \sigma(2\tau)^{(\sigma-1)} d\tau + \int_{\tau=0.5}^{\tau=1} \sigma(2(1-\tau))^{(\sigma-1)} d\tau \\ &= \sigma 2^{(\sigma-1)} \left(\int_{\tau=0}^{\tau=0.5} \tau^{(\sigma-1)} d\tau + \int_{\tau=0.5}^{\tau=1} (1-\tau)^{(\sigma-1)} d\tau \right) \\ &= \sigma 2^{(\sigma-1)} \left(\left[\frac{1}{\sigma} \tau^\sigma \right]_0^{0.5} + \left[\frac{-1}{\sigma} (1-\tau)^\sigma \right]_{0.5}^1 \right) \\ &= \sigma 2^{(\sigma-1)} \left(\frac{1}{\sigma} 2^{-\sigma} + \frac{1}{\sigma} 2^{-\sigma} \right) \\ &= 1. \end{aligned}$$

Thus, given this result, $Q(\tau)$ is a PDF for $\sigma \neq 1$.

For the case of $\sigma = 1$, $Q(\tau) = 1 \forall \tau \in [0, 1]$, and is hence a PDF.

B.5 Numerical Integration of the Mutual Information

This Section gives some details regarding calculating the mutual information between the input and output of the SSR model, via numerical integration. This method is required in Chapters 4, 8 and 9, and analogous methods are required for numerical integration of other quantities elsewhere.

B.5.1 Integrating Over the Input's Support

To obtain $I(x, y)$ numerically, it is necessary to perform numerical integrations. The simplest form of numerical integration is to approximate the signal PDF by a discrete version, with resolution $\Delta x \ll 1/N$, where $N + 1$ is the number of output states in the SSR model. Hence, if in the case of a continuously valued PDF $P(x)$ we have support $x \in [a, b]$, then discretisation with resolution Δx gives discrete values $x = a + i\Delta x$, $i = 0, 1, \dots, (b - a)/\Delta x$. This simple method of numerical integration is easily justified for calculating mutual information—see Cover and Thomas (1991).

Using this discretisation of $P(x)$, the mutual information can be written as

$$I(x, y) = - \sum_{n=0}^N P_y(n) \log_2 P_y(n) - \left(-\Delta x \sum_x P(x) \sum_{n=0}^N P(n|x) \log_2 P(n|x) \right), \quad (\text{B.47})$$

where

$$P_y(n) = \Delta x \sum_x P(n|x)P(x). \quad (\text{B.48})$$

Given these formulas, $P(x)$ and $P(n|x)$ need only be calculated for each specified value of x .

For an input distribution with finite support, such as the uniform distribution, only specification of the resolution is required. However for a distribution that has a PDF with infinite support, such as the Gaussian distribution, it is necessary to restrict the upper and lower bounds of the support of x to finite values. This is achieved in a discretisation of $P(x)$ by setting the maximum and minimum values of x to a multiple, w , of the standard deviation, σ_x . Thus, $x \in [-w\sigma_x, w\sigma_x]$ and $P(x)$ is then discretised to a resolution of Δx .

This method of numerical integration has been found to be sufficient. Its accuracy has been verified by calculating the mutual information by more sophisticated numerical integration schemes, such as Simpson quadrature (Press *et al.* 1992).

B.5.2 Numerical Integration Using $Q(\tau)$

Instead of numerically integrating Eqn. (4.5) and the integral on the RHS of Eqn. (4.17) between $\pm\infty$ (for signal densities with infinite support), the change of variables from x to τ appears to allow a simpler integration of Eqn. (4.24) and the integral on the RHS of Eqn. (4.23), between the limits of zero and unity.

Note that although $Q(\tau)$ is infinite at $\tau = 0$ and $\tau = 1$ for $\sigma < 1$, the integrand in both integrals is zero at these limits, except for $P^*(0)$ when $\tau = 0$, and $P^*(N)$ when $\tau = 1$. Thus, other than these two cases, both integrals are easily obtainable by standard numerical integration techniques. Using the law of total probability avoids the need to deal with the singularity in the integrals of $P^*(0)$ and $P^*(N)$, since by the evenness of the signal and noise densities, $P_y(0) = P_y(N)$. Once $P^*(n)$ has been calculated for all other n , the corresponding $P_y(n)$'s can be derived. Hence $P_y(0) = P_y(N) = 0.5 - 0.5 \sum_{n=1}^{N-1} P_y(n)$.

B.5.3 Comments on Numerical Integration

The main difficulty arising in numerical calculation of the mutual information arises for larger N . Calculations of the output probabilities can require multiplying a small number, $P^*(n)$, by a very large number, $\binom{N}{n}$. For large N , $\binom{N}{n}$ gets very large, and sufficient precision can soon be lost. Even if 32 bit floating point representation is used, the final calculation of $I(x, y)$ can be highly inaccurate. A way to circumvent this problem is to avoid calculating the set of $\binom{N}{n}$ for a given N . This can be carried out by calculating $P(n|x)$ recursively, using the technique described in Section 8.2.3 of Chapter 8.

Appendix C

Derivations for large N Suprathreshold Stochastic Resonance

C.1 Proof of Eqn. (5.9)

This section proves the identity given by Eqn. (5.9) in Section 5.2.1 of Chapter 5.

$$\begin{aligned} - \sum_{n=0}^N P_y(n) \log_2 \binom{N}{n} &= - \frac{1}{N+1} \sum_{n=0}^N \log_2 \binom{N}{n} \\ &= \frac{1}{N+1} \sum_{n=0}^N \log_2 \frac{n!(N-n)!}{N!} \\ &= \frac{1}{N+1} \left(- \sum_{n=0}^N \log_2 N! + \sum_{n=0}^N \log_2 n! + \sum_{n=0}^N \log_2 (N-n)! \right) \\ &= \frac{1}{N+1} \left(-(N+1) \log_2 N! + \sum_{n=0}^N \sum_{j=1}^n \log_2 j + \sum_{n=0}^N \sum_{j=1}^{N-n} \log_2 j \right) \\ &= - \log_2 N! \\ &\quad + \frac{1}{N+1} \left(\sum_{k=1}^N (N-k+1) \log_2 k + \sum_{k=1}^N (N-k+1) \log_2 k \right) \\ &= - \log_2 N! + \frac{2}{N+1} \left(\sum_{n=1}^N (N-n+1) \log_2 n \right) \\ &= - \log_2 N! + 2 \sum_{n=1}^N \log_2 n - \frac{2}{N+1} \sum_{n=1}^N n \log_2 n \\ &= - \log_2 N! + 2 \log_2 N! - \frac{2}{N+1} \sum_{n=1}^N n \log_2 n \\ &= \log_2 N! - \frac{2}{N+1} \sum_{n=1}^N n \log_2 n. \end{aligned} \tag{C.1}$$

C.2 Derivation of Eqn. (5.13)

This section gives the derivation of Eqn. (5.13) in Section 5.2.1 of Chapter 5 that states

$$\frac{2}{N+1} \sum_{n=1}^N n \log_2 n \simeq N \log_2 (N+1) - \frac{N(N+2)}{2 \ln 2(N+1)} + O\left(\frac{\log N}{N}\right). \quad (\text{C.2})$$

The Euler-Maclaurin summation formula (Spiegel and Liu 1999) states that a summation over n terms can be replaced by an integral plus correction terms, as

$$\begin{aligned} \sum_{n=1}^{N-1} F(n) &= \int_{n=0}^N F(n) dn - 0.5F(0) - 0.5F(N) \\ &+ \sum_{p=1}^{\infty} (-1)^{p-1} \frac{B_p}{(2p)!} \left(F^{(2p-1)}(N) - F^{(2p-1)}(0) \right), \end{aligned} \quad (\text{C.3})$$

where $F^{(2p-1)}$ indicates the $(2p-1)$ -th derivative of F , and B_p is the p -th Bernoulli number (Spiegel and Liu 1999).

Noting that $\sum_{n=1}^N n \log_2 n = \sum_{n=1}^{N-1} (n+1) \log_2 (n+1)$, then $F(n) = (n+1) \log_2 (n+1)$, $F(0) = 0$, and $F(N) = (N+1) \log_2 (N+1)$. Denoting the remainder term of Eqn. (C.3) involving derivatives of $F(n)$ as R , then

$$\begin{aligned} \sum_{n=1}^N n \log_2 n &= \sum_{n=1}^{N-1} (n+1) \log_2 (n+1) \\ &= -0.5F(0) - 0.5F(N) + R + \int_{n=0}^{n=N} (n+1) \log_2 (n+1) dn \\ &= 0.5(N+1)^2 \log_2 (N+1) - \frac{N(N+2)}{4 \ln 2} - 0.5(N+1) \log_2 (N+1) + R. \end{aligned} \quad (\text{C.4})$$

Thus, the desired simplification for the second term in Eqn. (5.9) is

$$\frac{2}{N+1} \sum_{n=1}^N n \log_2 n = N \log_2 (N+1) - \frac{N(N+2)}{2 \ln 2(N+1)} + \frac{2R}{N+1}. \quad (\text{C.5})$$

This is an exact expression that holds for any N , and can be obtained by evaluating R . In order to obtain a large N simplification we need to consider the behaviour of R for large N .

Firstly note that for $p=1$, $F'(N) - F'(0) = \log_2 (N+1)$ and for all $p > 1$,

$$F^{(2p-1)}(N) - F^{(2p-1)}(0) = \frac{(2p-3)!}{\ln 2} \left(1 - (N+1)^{-2p+2} \right). \quad (\text{C.6})$$

Substituting Eqn. (C.6) into the remainder term of Eqn. (C.3) and multiplying by $\frac{2}{N+1}$ gives

$$\begin{aligned} \frac{2R}{N+1} &= \frac{1}{6(N+1)} \log_2(N+1) \\ &+ \sum_{p=2}^{\infty} (-1)^{p-1} \frac{B_p}{\ln 2(4p^3 - 7p^2 + 2p)} \left(\frac{1}{N+1} - (N+1)^{-2p+1} \right). \end{aligned} \quad (\text{C.7})$$

Hence, the remainder term scales with $\frac{\log N}{N}$, which approaches zero for large N , and our approximation to the second term in Eqn. (5.9), for large N , is

$$\frac{2}{N+1} \sum_{n=1}^N n \log_2 n \simeq N \log_2(N+1) - \frac{N(N+2)}{2 \ln 2(N+1)} + O\left(\frac{\log N}{N}\right).$$

C.3 Proof that $S(x)$ is a PDF

This section shows that the function, $S(x)$, given by Eqn. (5.72) in Section (5.5.1) is a probability density function (PDF).

As shown in Section D.6 of Appendix D, the Fisher information for the SSR model is given by

$$J(x) = \frac{NR(x)^2}{P_{1|x}(1 - P_{1|x})}. \quad (\text{C.8})$$

Consider the function,

$$S(x) = \frac{\sqrt{J(x)}}{\pi\sqrt{N}} = \frac{R(x)}{\pi\sqrt{P_{1|x}(1 - P_{1|x})}}. \quad (\text{C.9})$$

Since $R(x)$ is a PDF and $P_{1|x} = \int_{-\infty}^x R(\phi) d\phi$ is the CDF of R evaluated at x , we have

$$\frac{dP_{1|x}}{dx} = R(x), \quad (\text{C.10})$$

and $S(x) \geq 0 \forall x$. Letting $f(x) = P_{1|x}$, Eqn. (C.9) can be written as

$$S(x) = \frac{f'(x)}{\pi\sqrt{f(x) - f(x)^2}}. \quad (\text{C.11})$$

C.3 Proof that $S(x)$ is a PDF

Suppose $R(x)$ has support $x \in [-a, a]$. Integrating $S(x)$ over all x gives

$$\begin{aligned}\int_{x=-a}^{x=a} S(x) dx &= \int_{x=-a}^{x=a} \frac{f'(x)}{\pi \sqrt{f(x) - f(x)^2}} dx \\ &= \frac{1}{\pi} \left(2 \arcsin \left(\sqrt{f(x)} \right) \Big|_{x=-a}^{x=a} \right) \\ &= \frac{2}{\pi} \left(\arcsin \left(\sqrt{P_1|_{x=a}} \right) - \arcsin \left(\sqrt{P_1|_{x=-a}} \right) \right) \\ &= \frac{2}{\pi} (\arcsin(1) - \arcsin(0)) \\ &= 1,\end{aligned}\tag{C.12}$$

and hence $S(x)$ is a PDF with the same support as $R(x)$.

Appendix D

Derivations for Suprathreshold Stochastic Resonance Decoding

D.1 Conditional Output Moments

This section derives expressions for the first and second conditional output moments of the SSR model, as well the conditional output variance. These expressions are used throughout Chapter 6.

The m -th moment of the output given x is given by

$$E[y^m|x] = \sum_{n=0}^N n^m P(n|x). \quad (\text{D.1})$$

Instead of using this definition directly, we will make use of the fact that all devices have identical thresholds. Using Eqn. (4.2) from Chapter 4, the first conditional moment is given by

$$\begin{aligned} E[y|x] &= E \left[\frac{1}{2} \sum_{i=1}^N \text{sign}[x + \eta_i] + \frac{N}{2} \middle| x \right] \\ &= \frac{N}{2} E[\text{sign}(x + \eta)|x] + \frac{N}{2}, \end{aligned} \quad (\text{D.2})$$

where since all N additive noise components are *iid*, the subscripts in the η_i s have been removed. This leads to

$$\begin{aligned} E[y|x] &= \frac{N}{2} (-(1 - P_{1|x}) + P_{1|x}) + \frac{N}{2} \\ &= NP_{1|x}. \end{aligned} \quad (\text{D.3})$$

D.2 Output Moments

The conditional variance is given by $\text{var}[y^2|x] = E[y^2|x] - E[y|x]^2$. The second conditional moment is

$$\begin{aligned} E[y^2|x] &= E \left[\left(\frac{1}{2} \sum_{i=1}^N \text{sign}(x + \eta_i) + \frac{N}{2} \right)^2 \mid x \right] \\ &= \frac{1}{4} E \left[\left(\sum_{i=1}^N \text{sign}(x + \eta_i) \right)^2 \mid x \right] + \frac{N^2}{2} E[\text{sign}(x + \eta)|x] + \frac{N^2}{4}. \end{aligned} \quad (\text{D.4})$$

We also have

$$\begin{aligned} E \left[\left(\sum_{i=1}^N \text{sign}(x + \eta_i) \right)^2 \mid x \right] &= NE[(\text{sign}(x + \eta))^2 \mid x] + \\ &\quad N(N-1)E[\text{sign}(x + \eta_j)\text{sign}(x + \eta_{jk}) \mid x] \mid_{\forall j \neq k} \\ &= N + N(N-1)(-2P_{1|x}(1 - P_{1|x}) + (1 - P_{1|x})^2 + P_{1|x}^2) \\ &= N + N(N-1)(2P_{1|x} - 1)^2. \end{aligned} \quad (\text{D.5})$$

Substituting Eqn. (D.5) into Eqn. (D.4) gives

$$\begin{aligned} E[y^2|x] &= \frac{1}{4} \left(N + N(N-1)(2P_{1|x} - 1)^2 \right) + \frac{N^2}{2}(2P_{1|x} - 1) + \frac{N^2}{4} \\ &= NP_{1|x}(1 - P_{1|x}) + N^2P_{1|x}^2. \end{aligned} \quad (\text{D.6})$$

Thus, using Eqns. (D.3) and (D.6) leaves the conditional variance as

$$\text{var}[y|x] = NP_{1|x}(1 - P_{1|x}). \quad (\text{D.7})$$

The correctness of these derivations of the conditional mean and variance can be seen by noting that the probability distribution of the output given the input is the binomial distribution, as given by Eqn. (4.9). It is well known that the expected value of such a binomially distributed variable is $\sum_{n=0}^N nP(n|x) = NP_{1|x}$, and that the variance is $\sum_{n=0}^N n^2P(n|x) - \left(\sum_{n=0}^N nP(n|x) \right)^2 = NP_{1|x}(1 - P_{1|x})$ (Kreyszig 1988).

D.2 Output Moments

This section derives expressions for the mean and mean square value of the output of the SSR model, as well the output variance. These expressions are used throughout Chapter 6.

The m -th moment of the output is given by

$$E[y^m] = E[E[y^m|x]] = \int_x \sum_{n=0}^N n^m P(x, y) dx \quad (\text{D.8})$$

$$= \int_x P(x) \sum_{n=0}^N n^m P(n|x) dx \quad (\text{D.9})$$

$$= \sum_{n=0}^N n^m P_y(n). \quad (\text{D.10})$$

If the output distribution is known, it is generally easier to apply Eqn. (D.10). However, in the following sections it will most often be the case that using Eqn. (D.9) will be more convenient. In particular, we have the conditional mean given by Eqn. (D.3) and the conditional second moment given by Eqn. (D.6). Thus

$$E[y] = NE[P_{1|x}], \quad (\text{D.11})$$

and

$$E[y^2] = NE[P_{1|x}] + N(N-1)E[P_{1|x}^2]. \quad (\text{D.12})$$

The output variance cannot be obtained by integration of Eqn. (D.7), but is easily obtained by noting that $\text{var}[y] = E[y^2] - E[y]^2$ as

$$\text{var}[y] = N(N-1)E[P_{1|x}^2] - NE[P_{1|x}](NE[P_{1|x}] - 1). \quad (\text{D.13})$$

D.2.1 Even Signal and Noise PDFs, All Thresholds Zero

For zero-mean noise PDFs that are even functions, and for all threshold values equal to zero, $P_{1|x}$ is given by Eqn. (4.8), i.e. $P_{1|x} = F_R(x)$, where $F_R(\cdot)$ is the CDF of the noise. This can be written as

$$P_{1|x} = \frac{1}{2} + \int_{\eta=0}^{\eta=x} R(\eta) d\eta. \quad (\text{D.14})$$

Therefore, the expected value of $P_{1|x}$ over the signal distribution is

$$\begin{aligned} E[P_{1|x}] &= \frac{1}{2} + E \left[\int_{\eta=0}^{\eta=x} R(\eta) d\eta \right] \\ &= \frac{1}{2} + \int_{x=-\infty}^{x=\infty} \left(\int_{\eta=0}^{\eta=x} R(\eta) d\eta \right) P(x) dx. \end{aligned} \quad (\text{D.15})$$

Since $R(\eta)$ is even, $\int_0^x R(\eta) d\eta$ is odd with respect to x , and therefore the integral above is zero, as $P(x)$ is even. Thus, $E[P_{1|x}] = 0.5$.

D.3 Correlation and Correlation Coefficient Expressions

Thus, Eqns. (D.11) and (D.12) can be simplified to

$$E[y] = \frac{N}{2}. \quad (\text{D.16})$$

and

$$E[y^2] = N(N-1)E[P_{1|x}^2] + \frac{N}{2}. \quad (\text{D.17})$$

This means that the variance of y is

$$\text{var}[y] = N(N-1)E[P_{1|x}^2] - \frac{N(N-2)}{4}. \quad (\text{D.18})$$

D.3 Correlation and Correlation Coefficient Expressions

This section derives expressions for the correlation coefficient at various stages of the SSR model. These expressions are used throughout Chapter 6.

D.3.1 Input Correlation Coefficient at Any Two Thresholds

Although the noise at the input to any given threshold device in the SSR model is uncorrelated with the noise at the input of any other device, due to the presence of the same input signal on each device there is a correlation between the inputs of any given pair of devices. This can be measured using the correlation coefficient, which for zero-meaned inputs to any two comparators, i and j is given by

$$\begin{aligned} \rho_i &= \frac{\text{cov}[x + \eta_i, x + \eta_j]}{\sqrt{\text{var}[x + \eta_i]\text{var}[x + \eta_j]}} \\ &= \frac{E[(x + \eta_i)(x + \eta_j)]}{\sqrt{(\text{var}[x + \eta])^2}} \\ &= \frac{E[x^2]}{\sqrt{(E[x^2] + E[\eta^2])^2}} \\ &= \frac{\sigma_x^2}{\sigma_x^2 + \sigma_\eta^2} \\ &= \frac{1}{1 + \sigma^2}, \end{aligned} \quad (\text{D.19})$$

where we have denoted the variance of the signal as σ_x^2 and the variance of the noise as σ_η^2 and let $\sigma^2 = \sigma_\eta^2/\sigma_x^2$. It is clear that the correlation coefficient is only unity in the absence of noise, and decreases towards zero as the noise variance increases.

D.3.2 Output Correlation Coefficient at Any Two Thresholds

As in Section D.3.1, which derives a formula for the correlation coefficient between the *inputs* to any two threshold devices, we can derive an expression for the correlation coefficient between the *outputs* of any pair of threshold devices, i and j . This is

$$\begin{aligned}
 \rho_o &= \frac{\text{cov}[y_i, y_j]}{\sqrt{\text{var}[y_i]\text{var}[y_j]}} \\
 &= \frac{\text{E}[y_i y_j] - \text{E}[y_i]\text{E}[y_j]}{\sqrt{(\text{var}[y_i])^2}} \\
 &= \frac{\text{E}[y_i y_j] - \frac{1}{4}}{\frac{1}{4}} \\
 &= \frac{\text{E}[P_{1|x}^2] - \frac{1}{4}}{\frac{1}{4}} \\
 &= 4\text{E}[P_{1|x}^2] - 1.
 \end{aligned} \tag{D.20}$$

D.3.3 Input-Output Correlation

The output of any given threshold device is correlated with its input. This correlation can be expressed as

$$\begin{aligned}
 \text{E}[xy_i] &= \sum_{y_i=0}^1 \int_{-\infty}^{\infty} xy_i P(x, y_i) dx \\
 &= \int_{-\infty}^{\infty} 0x(1 - P_{1|x})P(x) + 1xP_{1|x}P(x) dx \\
 &= \int_{-\infty}^{\infty} xP_{1|x}P(x) dx \\
 &= \text{E}[xP_{1|x}].
 \end{aligned} \tag{D.21}$$

Assuming the input signal has a mean of zero and variance σ_x^2 , the correlation coefficient between x and y_i is

$$\begin{aligned}
 \rho_{xy_i} &= \frac{\text{E}[xy_i]}{\sigma_x \sqrt{\text{var}[y_i^2]}} \\
 &= \frac{\text{E}[xP_{1|x}]}{\sigma_x \sqrt{\text{var}[y_i^2]}}.
 \end{aligned} \tag{D.22}$$

D.4 A Proof of Prudnikov's Integral

The overall system output encoding, y , is also correlated with the input signal, x . The input-output correlation is

$$E[xy] = \int_{-\infty}^{\infty} xE[y|x]P(x)dx \quad (\text{D.23})$$

$$= \sum_{n=0}^N nE[x|n]P_y(n). \quad (\text{D.24})$$

The fact that the above two identities are equal can easily be shown by substituting for the definition of expected value and the use of Bayes' rule, before changing the order of integration and summation. Substituting Eqn. (D.3) into Eqn. (D.23) gives an expression for the input-output correlation in terms of the correlation between the inputs to any two devices as

$$E[xy] = \int_{-\infty}^{\infty} xNP_{1|x}P(x)dx$$

$$= NE[xP_{1|x}] \quad (\text{D.25})$$

$$= NE[xy_i]. \quad (\text{D.26})$$

Hence, the overall input-output correlation is N times larger than the correlation between the input and output of a single device.

Assuming the input signal has a mean of zero and variance σ_x^2 , the correlation coefficient between x and y is

$$\begin{aligned} \rho_{xy} &= \frac{E[xy]}{\sigma_x \sqrt{\text{var}[y]}} \\ &= N \frac{E[xP_{1|x}]}{\sigma_x \sqrt{E[y^2] - E[y]^2}} \\ &= N \frac{E[xP_{1|x}]}{\sigma_x \sqrt{N(N-1)E[P_{1|x}^2] - N^2E[P_{1|x}]^2 + NE[P_{1|x}]}}. \end{aligned} \quad (\text{D.27})$$

To progress further requires knowledge of the noise PDF, $R(\eta)$. Once this is specified, $P_{1|x}$ can be derived, and therefore $E[xP_{1|x}]$ and the output moments.

D.4 A Proof of Prudnikov's Integral

This section gives a proof of an integral used in Section 6.4.2 of Chapter 6 for the calculation of the output variance for the SSR model for Gaussian signal and noise.

The integral we prove here is listed in Prudnikov *et al.* (1986) as

$$f(a) = \int_{x=-\infty}^{x=\infty} \exp(-a^2x^2)\text{erf}^2(x)dx = \frac{2}{a\sqrt{\pi}} \arctan\left(\frac{1}{a\sqrt{a^2+2}}\right). \quad (\text{D.28})$$

Observe that $f(a) = f(-a)$, and that the arctan function is an odd function. Eqn. (D.28) is expressed in terms of the arctan function so that the relationship holds for both positive and negative a . If we are only interested in the case of $a > 0$, the equation can be expressed in terms of the arcsin function as

$$f(a) = \int_{x=-\infty}^{x=\infty} \exp(-a^2x^2)\text{erf}^2(x)dx = \frac{2}{a\sqrt{\pi}} \arcsin\left(\frac{1}{1+a^2}\right). \quad (\text{D.29})$$

The proof of this result follows. Firstly note that the integrand is an even function of x . Hence,

$$f(a) = 2 \int_{x=0}^{x=\infty} \exp(-a^2x^2)\text{erf}^2(x)dx. \quad (\text{D.30})$$

Now,

$$\begin{aligned} \int_{x=0}^{x=\infty} \exp(-a^2x^2)\text{erf}^2(x)dx &= \int_{x=0}^{x=\infty} \exp(-a^2x^2) \left(\frac{2}{\sqrt{\pi}} \int_{u=0}^{u=x} \exp(-u^2)du\right)^2 dx \\ &= \frac{4}{\pi} \int_{x=0}^{x=\infty} \int_{u=0}^{u=x} \int_{v=0}^{v=x} \exp(-(u^2 + v^2 + a^2x^2))dudvdx \end{aligned} \quad (\text{D.31})$$

Letting $w = ax$ and performing a change of variable gives

$$f(a) = \frac{8}{a\pi} \int_{w=0}^{w=\infty} \int_{u=0}^{u=w/a} \int_{v=0}^{v=w/a} \exp(-(u^2 + v^2 + w^2))dudvdw. \quad (\text{D.32})$$

We now convert this triple integral to polar coordinates with $r \in [0, \infty)$, $\theta \in [0, 2\pi)$ and $\phi \in [0, \pi]$. This gives $u = r \cos \theta \sin \phi$, $v = r \sin \theta \sin \phi$ and $w = r \cos \phi$, with the volume element being $dudvdw = r^2 \sin \phi d\phi d\theta dr$. With this conversion we get $u^2 + v^2 + w^2 = r^2$. Thus

$$f(a) = \frac{8}{a\pi} \int_r \int_\theta \int_\phi r^2 \exp(-r^2) \sin \phi d\phi d\theta dr. \quad (\text{D.33})$$

The limits of integration for θ and ϕ can be determined as follows. Firstly, since u and v are integrated between 0 and w/a , this imposes the following inequalities

$$\begin{aligned} 0 &\leq \cos \theta \sin \phi \leq \frac{\cos \phi}{a}, \\ 0 &\leq \sin \theta \sin \phi \leq \frac{\cos \phi}{a}. \end{aligned} \quad (\text{D.34})$$

Since ϕ is defined on $[0, \pi]$, $\sin \phi$ is always nonnegative. Therefore both $\cos \theta \geq 0$ and $\sin \theta \geq 0$, which implies that $\theta \in [0, \pi/2]$. Also, since $a \geq 0$, $\cos \phi \geq 0$ and therefore $\phi \in [0, \pi/2]$.

D.4 A Proof of Prudnikov's Integral

Thus, Inequalities (D.34) can be written as

$$\begin{aligned} 0 &\leq \tan \phi \leq \frac{1}{a \cos \theta}, \\ 0 &\leq \tan \phi \leq \frac{1}{a \sin \theta}. \end{aligned} \quad (\text{D.35})$$

This can be restated as

$$\begin{aligned} \phi &\leq \phi_1 = \arctan \left(\frac{1}{a \cos \theta} \right), \quad \theta \in \left[0, \frac{\pi}{4} \right], \\ \phi &\leq \phi_2 = \arctan \left(\frac{1}{a \sin \theta} \right), \quad \theta \in \left[\frac{\pi}{4}, \frac{\pi}{2} \right]. \end{aligned} \quad (\text{D.36})$$

Furthermore, since in the original integral of Eqn. (D.32), w is integrated between 0 and ∞ , and $\cos \phi \geq 0$, r must be integrated over $[0, \infty)$. Therefore

$$f(a) = \frac{8}{a\pi} \int_{\theta} \int_{\phi} \sin \phi \int_{r=0}^{\infty} r^2 \exp(-r^2) dr d\phi d\theta. \quad (\text{D.37})$$

The inner integral is simply $\int_{r=0}^{\infty} r^2 \exp(-r^2) dr = \sqrt{\pi}/4$. Thus

$$\begin{aligned} f(a) &= \frac{2}{a\sqrt{\pi}} \int_{\theta} \int_{\phi} \sin \phi d\phi d\theta \\ &= \frac{2}{a\sqrt{\pi}} \left(\int_{\theta=0}^{\theta=\pi/4} \int_{\phi=0}^{\phi=\phi_1} \sin \phi d\phi d\theta + \int_{\theta=\pi/4}^{\theta=\pi/2} \int_{\phi=0}^{\phi=\phi_2} \sin \phi d\phi d\theta \right) \\ &= \frac{2}{a\sqrt{\pi}} \left(\int_{\theta=0}^{\theta=\pi/4} [-\cos \phi]_0^{\phi_1} d\theta + \int_{\theta=\pi/4}^{\theta=\pi/2} [-\cos \phi]_0^{\phi_2} d\theta \right) \\ &= \frac{2}{a\sqrt{\pi}} \left(\int_{\theta=0}^{\theta=\pi/4} 1 - \cos \left(\arctan \left(\frac{1}{a \cos \theta} \right) \right) d\theta \right. \\ &\quad \left. + \int_{\theta=\pi/4}^{\theta=\pi/2} 1 - \cos \left(\arctan \left(\frac{1}{a \sin \theta} \right) \right) d\theta \right) \\ &= \frac{2}{a\sqrt{\pi}} \left(\frac{\pi}{2} - 2 \int_{\theta=0}^{\theta=\pi/4} \cos \left(\arctan \left(\frac{1}{a \cos \theta} \right) \right) d\theta \right) \\ &= \frac{2}{a\sqrt{\pi}} \left(\frac{\pi}{2} - 2 \int_{\theta=0}^{\theta=\pi/4} \frac{a \cos \theta}{\sqrt{1 + a^2 \cos^2 \theta}} d\theta \right) \\ &= \frac{2}{a\sqrt{\pi}} \left(\frac{\pi}{2} - 2 \left[\arcsin \left(\frac{a \sin \theta}{\sqrt{1 + a^2}} \right) \right]_0^{\pi/4} \right) \\ &= \frac{2}{a\sqrt{\pi}} \left(\frac{\pi}{2} - 2 \arcsin \left(\frac{a}{\sqrt{2}\sqrt{1 + a^2}} \right) \right) \\ &= \frac{2}{a\sqrt{\pi}} \left(\frac{\pi}{2} - 2 \arcsin \left(\frac{1}{\sqrt{2}} \sqrt{1 - \frac{1}{1 + a^2}} \right) \right). \end{aligned} \quad (\text{D.38})$$

Now, given the identity $A = 2 \arcsin \left(\frac{1}{\sqrt{2}} \sqrt{1 - \cos A} \right)$, it can be seen that Eqn. (D.38) can be written as

$$\begin{aligned} f(a) &= \frac{2}{a\sqrt{\pi}} \left(\frac{\pi}{2} - \arccos \left(\frac{1}{1+a^2} \right) \right) \\ &= \frac{2}{a\sqrt{\pi}} \arcsin \left(\frac{1}{1+a^2} \right). \end{aligned} \quad (\text{D.39})$$

This completes the proof.

D.5 Minimum Mean Square Error Distortion Decoding

This section proves the results stated in Section 6.5.3 of Chapter 6, regarding the minimum possible MSE distortion.

D.5.1 MMSE Reconstruction Points and Distortion

The reconstruction points that provide the Minimum Mean Square Error (MMSE) distortion for a quantisation encoding, y , consisting of possible states, $n = 0, \dots, N$, are

$$\hat{x}_n = E_x[x|n] = \int_x xP(x|n)dx = \frac{1}{P_y(n)} \int_x xP(n|x)P(x)dx. \quad (\text{D.40})$$

A proof of this follows.

If \hat{x} , with possible values given by $\{\hat{x}_0, \dots, \hat{x}_N\}$, is a decoding of a quantised version of a signal, x , the mean square error distortion is

$$\begin{aligned} \text{MSE} &= E[(x - \hat{x})^2] \\ &= \int_{x=-\infty}^{\infty} \sum_{n=0}^N (x - \hat{x}_n)^2 P(x, n) dx \\ &= E[x^2] + \sum_{n=0}^N (\hat{x}_n^2 - 2x\hat{x}_n) P_y(n) P(x|n) dx \\ &= E[x^2] + \sum_{n=0}^N P_y(n) \left(\hat{x}_n^2 \int_{x=-\infty}^{\infty} P(x|n) dx - 2\hat{x}_n \int_{x=-\infty}^{\infty} xP(x|n) dx \right) \\ &= E[x^2] + \sum_{n=0}^N P_y(n) \left(\hat{x}_n^2 - 2\hat{x}_n E[x|n] \right). \end{aligned} \quad (\text{D.41})$$

Notice that since we wish to find the set of \hat{x}_n that minimises the MSE distortion, and since $P_y(n)$ is always positive, we can simply differentiate the term inside the summation in Eqn. (D.41) with respect to \hat{x}_n and set to zero. Also, the second derivative with

D.5 Minimum Mean Square Error Distortion Decoding

respect to \hat{x}_n is equal to 2, which is always greater than zero. This gives $\hat{x}_n = E[x|n]$ as a minimum and we have completed the proof.

Note also that substituting for $\hat{x}_n = E[x|n]$ in Eqn. (D.41) gives the minimum MSE distortion as

$$\text{MMSE} = E[x^2] - E[\hat{x}^2]. \quad (\text{D.42})$$

D.5.2 MMSE Decoded Output is Uncorrelated with the Error

We firstly show that $E[x\hat{x}] = E[\hat{x}^2]$:

$$\begin{aligned} E[x\hat{x}] &= \sum_n \int_x x\hat{x}P(x,n)dx \\ &= \sum_n \int_x x\hat{x}P_y(n)P(x|n)dx \\ &= \sum_n P_y(n)\hat{x} \int_x xP(x|n)dx \\ &= \sum_n P_y(n)\hat{x}\hat{x} \\ &= E[\hat{x}^2]. \end{aligned} \quad (\text{D.43})$$

Thus, if the optimal decoding, \hat{x} , is used, the mean square error is

$$\begin{aligned} E[(\hat{x} - x)^2] &= E_n[\hat{x}^2] - 2E[x\hat{x}] + E_x[x^2] \\ &= E_x[x^2] - E_n[\hat{x}^2], \end{aligned} \quad (\text{D.44})$$

just as in Eqn. (D.42).

We also have,

$$\begin{aligned} E[\hat{x} - x] &= E[\hat{x}] - 0 \\ &= \sum_n P_y(n)E[x|n] \\ &= \sum_n P_y(n) \int_x xP(x|n)dx \\ &= \int_x x \sum_n P_y(n)P(x|n)dx \\ &= \int_x xP(x)dx \\ &= E[x] \\ &= 0. \end{aligned} \quad (\text{D.45})$$

Thus, the mean of the decoded output is zero, the mean error is zero and therefore the MMSE distortion is also the minimum error variance.

Furthermore, with the decoding, \hat{x} , the encoded output, y , is uncorrelated with the error, $\epsilon = x - \hat{x}$. The proof of this is

$$\begin{aligned}
E[\epsilon y] &= E[xy] - E[\hat{x}y] \\
&= \sum_n \int_x xnP(x, n)dx - E[\hat{x}y] \\
&= \sum_n \int_x xnP_y(n)P(x|n)dx - E[\hat{x}y] \\
&= \sum_n nP_y(n) \int_x xP(x|n)dx - E[\hat{x}y] \\
&= \sum_n nP_y(n)E[x|n] - E[\hat{x}y] \\
&= \sum_n n\hat{x}_n P_y(n) - E[\hat{x}y] \\
&= E[\hat{x}y] - E[\hat{x}y] \\
&= 0.
\end{aligned} \tag{D.46}$$

D.5.3 Relationship of MMSE to Backwards Conditional Variance

We are also consider the mean square value of x given output state n . This gives an idea of how variable x is, given that output state. This is

$$E_x[x^2|n] = \int_x x^2 P(x|n)dx = \frac{1}{P_y(n)} \int_x x^2 P(n|x)P(x)dx. \tag{D.47}$$

Consider the variance of x given output state n . We call this the Backwards Conditional Variance (BCV), which we label as

$$\text{BCV}(n) = \text{var}_x[x|n] = E_x[x^2|n] - E_x[x|n]^2 = E_x[x^2|n] - \hat{x}_n^2. \tag{D.48}$$

D.6 Fisher Information

Note that the BCV is a function of n . If we take the expected value of the BCV over all N we get

$$\begin{aligned}
 E_n[\text{BCV}(n)] &= E_n \left[E_x[x^2|n] - \hat{x}_n^2 \right] \\
 &= \sum_n P_y(n) \int_x x^2 P(x|n) dx - E[\hat{x}_n^2] \\
 &= \int_x x^2 P(x) \left(\sum_n P(n|x) \right) dx - E[\hat{x}_n^2] \\
 &= \int_x x^2 P(x) 1 dx - E[\hat{x}_n^2] \\
 &= E[x^2] - E[\hat{x}_n^2]. \tag{D.49}
 \end{aligned}$$

The RHS of Eqn. (D.49) is precisely the MMSE that is obtained using the estimator $\hat{x} = E[x|n]$. So we make note of this by writing explicitly,

$$\text{MMSE} = E_n[\text{BCV}(n)] = E_x[x^2] - E_n[\hat{x}_n^2]. \tag{D.50}$$

D.6 Fisher Information

This section provides two alternative derivations for the Fisher information in the SSR model, an expression which is used in Section 5.4.5 of Chapter 5 and Section 6.7 of Chapter 6.

D.6.1 First Derivation

Consider the individual output signal of each threshold device in the SSR model, $y_i \in \{0, 1\}$, to be an estimator for the input, x . There are two output states, zero and one, with conditional probability functions $P(0|x) = 1 - P_{1|x}$ and $P(1|x) = P_{1|x}$ respectively. Hence, from Eqn. (6.104), the score for each state is

$$V(0) = \frac{1}{(1 - P_{1|x})} \frac{d(1 - P_{1|x})}{dx}, \tag{D.51}$$

and

$$V(1) = \frac{1}{P_{1|x}} \frac{dP_{1|x}}{dx}. \tag{D.52}$$

Therefore, upon substituting into Eqn. (6.114) the Fisher information for each comparator is

$$J_i(x) = (1 - P_{1|x}) \frac{1}{(1 - P_{1|x})^2} \left(\frac{d(1 - P_{1|x})}{dx} \right)^2 + P_{1|x} \frac{1}{P_{1|x}^2} \left(\frac{dP_{1|x}}{dx} \right)^2, \tag{D.53}$$

which simplifies to

$$J_i(x) = \left(\frac{dP_{1|x}}{dx} \right)^2 \frac{1}{P_{1|x}(1 - P_{1|x})}. \quad (\text{D.54})$$

The Fisher information for N iid samples is N times the individual information (Cover and Thomas 1991). Therefore since in the SSR array of N threshold devices, the set of N random variables y_i are all conditionally iid given x , the overall Fisher information is N times the formula given in Eqn. (D.54). Thus

$$J(x) = \left(\frac{dP_{1|x}}{dx} \right)^2 \frac{N}{P_{1|x}(1 - P_{1|x})}. \quad (\text{D.55})$$

This is in agreement with the formula derived for the Fisher information in Hoch *et al.* (2003a) and Hoch *et al.* (2003b).

D.6.2 Second Derivation

This subsection gives a derivation of the Fisher information in the SSR model that does not require using the fact that the overall Fisher information is the sum of the Fisher information in each individual threshold device. Recall that the transition probabilities, $P(n|x)$, are given by the binomial formula in terms of $P_{1|x}$ as

$$P(n|x) = \binom{N}{n} (P_{1|x})^n (1 - P_{1|x})^{N-n} \quad n = 0, \dots, N. \quad (\text{D.56})$$

Differentiation of $P(n|x)$ with respect to x gives

$$\frac{dP(n|x)}{dx} = \frac{(NP_{1|x} - n) \frac{dP_{1|x}}{dx}}{P_{1|x}(1 - P_{1|x})} P(n|x). \quad (\text{D.57})$$

Substituting Eqn. (D.57) into Eqn. (6.116) gives

$$J(x) = \frac{\frac{dP_{1|x}}{dx}^2}{P_{1|x}^2(1 - P_{1|x})^2} \sum_{n=0}^N P(n|x) (NP_{1|x} - n)^2. \quad (\text{D.58})$$

Noting that $\sum_{n=0}^N nP(n|x) = NP_{1|x}$ and $\sum_{n=0}^N n^2P(n|x) = NP_{1|x}(1 - P_{1|x}) + N^2P_{1|x}^2$ Eqn. (D.58) simplifies to

$$J(x) = \left(\frac{dP_{1|x}}{dx} \right)^2 \frac{N}{P_{1|x}(1 - P_{1|x})}. \quad (\text{D.59})$$

This result is in agreement with Eqn. (D.55).

D.7 Proof of the Information and Cramer-Rao Bounds

This section gives proofs of the information bound, and the Cramer-Rao bound. These bound are used in Section 6.7 of Chapter 6.

The Cramer-Rao bound can be easily proved using the Cauchy-Schwarz inequality, which states that

$$(\text{cov}[U, y|x])^2 \leq \text{var}[U|x]\text{var}[y|x]. \quad (\text{D.60})$$

Here, we set U to be the *score function*—the gradient of the log-likelihood function—which is

$$U(x, y) = \frac{d}{dx} \ln P(n|x). \quad (\text{D.61})$$

To prove the Cramer-Rao bound, note firstly that the covariance of U and y can be simplified using the fact that $E[U|x] = 0$ as

$$\begin{aligned} \text{cov}[U, y|x] &= E[(U - E[U|x])(y - E[y|x])|x] \\ &= E[U(y - E[y|x])|x] \\ &= E[Uy|x] - E[y|x]E[U|x] \\ &= E[Uy|x] \\ &= E\left[y \frac{d}{dx} \ln P(n|x)|x\right] \\ &= E\left[y \frac{1}{P(n|x)} \frac{d}{dx} P(n|x)|x\right] \\ &= \sum_{n=0}^N n \frac{d}{dx} P(n|x) \\ &= \frac{d}{dx} \sum_{n=0}^N nP(n|x) \\ &= \frac{d}{dx} E[y|x]. \end{aligned} \quad (\text{D.62})$$

Noting that $J(x) = \text{var}[U|x]$,

$$\left(\frac{d}{dx} E[y|x]\right)^2 \leq J(x)\text{var}[y|x]. \quad (\text{D.63})$$

Therefore

$$\text{var}[y|x] \geq \frac{\left(\frac{d}{dx} E[y|x]\right)^2}{J(x)}. \quad (\text{D.64})$$

This proves the information bound for a biased estimator. Replacing y with a decoding \hat{y} in the above derivation gives

$$\text{var}[\hat{y}|x] \geq \frac{\left(\frac{d}{dx}\text{E}[\hat{y}|x]\right)^2}{J(x)}. \quad (\text{D.65})$$

For an unbiased estimator, $\frac{d}{dx}\text{E}[y|x] = \frac{d}{dx}x = 1$ and

$$\text{var}[y|x] \geq \frac{1}{J(x)}, \quad (\text{D.66})$$

which proves the Cramer-Rao bound for an unbiased estimator.

Bibliography

- ABARBANEL-H. D. I., AND RABINOVICH-M. I. (2001). Neurodynamics: nonlinear dynamics and neurobiology, *Current Opinion in Neurobiology*, **11**, p. 423-430.
- ABBOTT-D. (2001). Overview: Unsolved problems of noise and fluctuations, *Chaos*, **11**, pp. 526–538.
- ABRAMOWITZ-M., AND STEGUN-I. A. (1972). *Handbook of Mathematical Functions*, National Bureau of Standards, Washington DC, USA.
- AKYILDIZ-I. F., SU-W., SANKARASUBRAMANIAM-Y., AND CAYIRCI-E. (2002). A survey on sensor networks, *IEEE Communications Magazine*, **40**, pp. 102–114.
- ALLINGHAM-D., STOCKS-N. G., AND MORSE-R. P. (2003). The use of suprathreshold stochastic resonance in cochlear implant coding, in S. M. Bezrukov., H. Frauenfelder., and F. Moss. (eds.), *Proc. SPIE Fluctuations and Noise in Biological and Biomedical Systems*, Vol. 5110, pp. 92–101.
- ALLINGHAM-D., STOCKS-N. G., MORSE-R. P., AND MEYER-G. F. (2004). Noise enhanced information transmission in a model of multichannel cochlear implantation, in D. Abbott., S. M. Bezrukov., A. Der., and A. Sánchez. (eds.), *Proc. SPIE Fluctuations and Noise in Biological, Biophysical, and Biomedical Systems II*, Vol. 5467, pp. 139–148.
- ALLISON-A., AND ABBOTT-D. (2000). Some benefits of random variables in switched control systems, *Microelectronics Journal*, **31**, pp. 515–522.
- ANDERSON-V. C. (1960). Digital array phasing, *The Journal of the Acoustical Society of America*, **32**, pp. 867–870.
- ANDERSON-V. C. (1980). Nonstationary and nonuniform oceanic background in a high-gain acoustic array, *The Journal of the Acoustical Society of America*, **67**, pp. 1170–1179.
- ANDO-B. (2002). Stochastic resonance and dithering: A matter of classification!, *IEEE Instrumentation and Measurement Magazine*, **5**, pp. 60–63.
- ANDO-B., AND GRAZIANI-S. (2000). *Stochastic Resonance: Theory and Applications*, Kluwer Academic Publishers.
- ANDREWS-G. E. (1976). *The Theory of Partitions*, Addison-Wesley Publishing Company, London, UK.
- ANISCHENKO-V. S., ASTAKOV-V. V., NEIMAN-A. B., VADIVASOVA-T. E., AND SCHIMANSKY-GEIER-L. (2002). *Nonlinear Dynamics of Chaotic and Stochastic Systems*, Springer, New York.
- ANISCHENKO-V. S., NEIMAN-A. B., MOSS-F., AND SCHIMANSKY-GEIER-L. (1999). Stochastic resonance: Noise enhanced order, *Uspekhi Fizicheskikh Nauk*, **169**, pp. 7–38.
- ARIMOTO-S. (1972). An algorithm for computing the capacity of arbitrary discrete memoryless channels, *IEEE Transactions on Information Theory*, **18**, pp. 14–20.
- ASTUMIAN-R. D., ADAIR-R. K., AND WEAVER-J. C. (1997). Stochastic resonance at the single-cell level, *Nature*, **388**, pp. 632–633.

Bibliography

- ASTUMIAN-R. D., AND MOSS-F. (1998). Overview: The constructive role of noise in fluctuation driven transport and stochastic resonance, *Chaos*, **8**, pp. 533–538.
- AYANOĞLU-E. (1990). On optimal quantization of noisy sources, *IEEE Transactions on Information Theory*, **36**, pp. 1450–1452.
- BADZEY-R. L., AND MOHANTY-P. (2005). Coherent signal amplification in bistable nanomechanical oscillators by stochastic resonance, *Nature*, **437**, pp. 995–998.
- BALASUBRAMANIAN-V., KIMBER-D., AND BERRY II-M. J. (2001). Metabolically efficient information processing, *Neural Computation*, **13**, pp. 799–815.
- BARRETT BROWNING-E. (1998). How do I love thee? Let me count the ways, in C. Graham. (ed.), *Elizabeth Barrett Browning*, Everyman Paperbacks, p. 36.
- BARRON-A. R., AND COVER-T. M. (1991). Minimum complexity density estimation, *IEEE Transactions on Information Theory*, **37**, pp. 1034–1054.
- BARTUSSEK-R., HÄNGGI-P., AND JUNG-P. (1994). Stochastic resonance in optical bistable systems, *Physical Review E*, **49**, pp. 3930–3939.
- BASAK-G. K. (2001). Stabilization of dynamical systems by adding a colored noise, *IEEE Transactions on Automatic Control*, **46**, pp. 1107–1111.
- BEASLEY-D., BULL-D. R., AND MARTIN-R. R. (1993). An overview of genetic algorithms: Part I, fundamentals, *University Computing*, **15**, pp. 58–69.
- BEHNAM-S. E., AND ZENG-F.-G. (2003). Noise improves suprathreshold discrimination in cochlear-implant listeners, *Hearing Research*, **186**, pp. 91–93.
- BENDAT-J. S. (1998). *Nonlinear System Techniques and Applications*, John Wiley and Sons, Inc, New York.
- BENNETT-M., WIESENFELD-K., AND JARAMILLO-F. (2004). Stochastic resonance in hair cell mechano-electrical transduction, *Fluctuation and Noise Letters*, **4**, pp. L1–L10.
- BENZI-R. (1980). *NATO International School of Climatology*, Ettore Majorana Center for Scientific Culture, Erice, Italy (unpublished).
- BENZI-R., PARISI-G., SUTERA-A., AND VULPIANI-A. (1982). Stochastic resonance in climatic change, *Tellus*, **34**, pp. 10–16.
- BENZI-R., PARISI-G., SUTERA-A., AND VULPIANI-A. (1983). A theory of stochastic resonance in climatic change, *SIAM Journal on Applied Mathematics*, **43**, pp. 565–578.
- BENZI-R., SUTERA-A., AND VULPIANI-A. (1981). The mechanism of stochastic resonance, *Journal of Physics A: Mathematical and General*, **14**, pp. L453–L457.
- BENZI-R., SUTERA-A., AND VULPIANI-A. (1985). Stochastic resonance in the Landau-Ginzburg equation, *Journal of Physics A: Mathematical and General*, **18**, pp. 2239–2245.
- Berger, A.. (ed.) (1980). *Climatic Variations and Variability: Facts and Theories*, NATO Advanced Study Institutes Series, D. Riedel Publishing Company.

-
- BERGER-T. (1971). *Rate Distortion Theory: A Mathematical Basis for Data Compression*, Prentice-Hall Inc., New Jersey.
- BERGER-T., AND GIBSON-J. D. (1998). Lossy source coding, *IEEE Transactions on Information Theory*, **44**, pp. 2693–2723.
- BERGER-T., ZHANG-Z., AND VISWANATHAN-H. (1996). The CEO problem, *IEEE Transactions on Information Theory*, **42**, pp. 887–902.
- BERNDT-H. (1968). Correlation function estimation by a polarity method using stochastic reference signals, *IEEE Transactions on Information Theory*, **14**, pp. 796–801.
- BERNDT-H., AND JENTSCHHEL-H. J. (2001). Differentially randomized quantization in sigma-delta analog-to-digital converters, *Proc. 8th IEEE International Conference on Electronics, Circuits and Systems*, Vol. 2, pp. 1057–1060.
- BERNDT-H., AND JENTSCHHEL-H. J. (2002). Stochastic quantization transfer functions for high resolution signal estimation, *Proc. 14th International Conference on Digital Signal Processing*, Vol. 2, pp. 881–884.
- BERSHAD-N. J., AND FEINTUCH-P. L. (1974). Sonar array detection of Gaussian signals in Gaussian noise of unknown power, *IEEE Transactions on Aerospace and Electronic Systems*, **10**, pp. 94–99.
- BETHGE-M. (2003). *Codes and Goals of Neuronal Representation*, PhD thesis, Universität Bremen.
- BETHGE-M., ROTERMUND-D., AND PAWELZIK-K. (2002). Optimal short-term population coding: When Fisher information fails, *Neural Computation*, **14**, pp. 2317–2351.
- BEZRUKOV-S. M., AND VOYDANOY-I. (1995). Noise-induced enhancement of signal transduction across voltage-dependent ion channels, *Nature*, **378**, pp. 362–364.
- BEZRUKOV-S. M., AND VOYDANOY-I. (1997). Stochastic resonance in non-dynamical systems without response thresholds, *Nature*, **385**, pp. 319–321.
- BIALEK-W., DEWEESE-M., RIEKE-F., AND WARLAND-D. (1993). Bits and brains: Information flow in the nervous system, *Physica A*, **200**, pp. 581–593.
- BIALEK-W., RIEKE-F., DE RUYTER VAN STEVENINCK-R. R., AND WARLAND-D. (1991). Reading a neural code, *Science*, **252**, pp. 1854–1857.
- BLAHUT-R. E. (1972). Computation of channel capacity and rate-distortion functions, *IEEE Transactions on Information Theory*, **18**, pp. 460–473.
- BLARER-A., AND DOEBELI-M. (1999). Resonance effects and outbreaks in ecological time series, *Ecology Letters*, **2**, pp. 167–177.
- BLUM-R. S. (1995). Quantization in multisensor random signal detection, *IEEE Transactions on Information Theory*, **41**, pp. 204–215.
- BORST-A., AND THEUNISSEN-F. E. (1999). Information theory and neural coding, *Nature Neuroscience*, **2**, pp. 947–957.
- BOWEN-G., AND MANCINI-S. (2004). Noise enhancing the classical information capacity of a quantum channel, *Physics Letters A*, **321**, pp. 1–5.
-

Bibliography

- BRUNEL-N., AND NADAL-J. (1998). Mutual information, Fisher information and population coding, *Neural Computation*, **10**, pp. 1731–1757.
- BUCKLEW-J. A. (1981). Companding and random quantization in several dimensions, *IEEE Transactions on Information Theory*, **IT-27**, pp. 207–211.
- BULSARA-A. (2005). No-nuisance noise, *Nature*, **437**, pp. 962–963.
- BULSARA-A., AND ZADOR-A. (1996). Threshold detection of wideband signals: A noise induced maximum in the mutual information, *Physical Review E*, **54**, pp. R2185–R2188.
- BULSARA-A., HÄNGGI-P., MARCHESONI-F., MOSS-F., AND SHLESINGER-M. (1993). Proceedings of the NATO advanced research workshop—stochastic resonance in physics and biology—preface, *Journal of Statistical Physics*, **70**, pp. 1–2.
- BULSARA-A., JACOBS-E. W., AND ZHOU-T. (1991). Stochastic resonance in a single neuron model—theory and analog simulation, *Journal of Theoretical Biology*, **152**, pp. 531–555.
- BULSARA-A. R., AND GAMMAITONI-L. (1996). Tuning in to noise, *Physics Today*, **49**, pp. 39–45.
- BULSARA-A. R., AND INCHIOSA-M. E. (1996). Noise-controlled resonance behavior in nonlinear dynamical systems with broken symmetry, *Physical Review Letters*, **77**, pp. 2162–2165.
- BULSARA-A. R., AND MOSS-F. E. (1991). Single neuron dynamics: Noise-enhanced signal processing, *IEEE International Joint Conference on Neural Networks*, Vol. 1, pp. 420–425.
- CARBONE-P., AND PETRI-D. (1998). Stochastic-flash analog-to-digital conversion, *IEEE Transactions on Instrumentation and Measurement*, **47**, pp. 65–68.
- CARBONE-P., AND PETRI-D. (2000). Performance of stochastic and deterministic dithered quantizers, *IEEE Transactions on Instrumentation and Measurement*, **49**, pp. 337–340.
- CARTER-G. C. (1993). *Coherence and Time Delay Estimation*, IEEE Press, New Jersey.
- CASADO-PASCUAL-J., DENK-C., GOMEZ-ORDONEZ-J., MARILLO-M., AND HÄNGGI-P. (2003). Gain in stochastic resonance: Precise numerics versus linear response theory beyond the two-mode approximation, *Physical Review E*, **67**, Art. No. 036109.
- CASTANIE-F. (1979). Signal processing by random reference quantizing, *Signal Processing*, **1**, pp. 27–43.
- CASTANIE-F. (1984). Linear mean transfer random quantization, *Signal Processing*, **7**, pp. 99–117.
- CASTANIE-F., HOFFMAN-J. C., AND LACAZE-B. (1974). On the performance of a random reference correlator, *IEEE Transactions on Information Theory*, **IT-20**, pp. 266–269.
- CHALLET-D., AND JOHNSON-N. F. (2002). Optimal combinations of imperfect objects, *Physical Review Letters*, **89**, Art. No. 028701.
- CHANG-C., AND DAVISSON-L. D. (1990). Two iterative algorithms for finding minimax solutions, *IEEE Transactions on Information Theory*, **36**, pp. 126–140.
- CHANG-C. I., AND DAVISSON-L. D. (1988). On calculating the capacity of an infinite-input finite (infinite)-output channel, *IEEE Transactions on Information Theory*, **34**, pp. 1004–1010.

- CHAPEAU-BLONDEAU-F. (1996). Stochastic resonance in the Heaviside nonlinearity with white noise and arbitrary periodic signal, *Physical Review E*, **53**, pp. 5469–5472.
- CHAPEAU-BLONDEAU-F. (1997a). Input-output gains for signal in noise in stochastic resonance, *Physics Letters A*, **232**, pp. 41–48.
- CHAPEAU-BLONDEAU-F. (1997b). Noise-enhanced capacity via stochastic resonance in an asymmetric binary channel, *Physical Review E*, **55**, pp. 2016–2019.
- CHAPEAU-BLONDEAU-F. (1999). Periodic and aperiodic stochastic resonance with output signal-to-noise ratio exceeding that at the input, *International Journal of Bifurcation and Chaos*, **9**, pp. 267–272.
- CHAPEAU-BLONDEAU-F., AND GODIVIER-X. (1996). Stochastic resonance in nonlinear transmission of spike signals: An exact model and an application to the neuron, *International Journal of Bifurcation and Chaos*, **6**, pp. 2069–2076.
- CHAPEAU-BLONDEAU-F., AND GODIVIER-X. (1997). Theory of stochastic resonance in signal transmission by static nonlinear systems, *Physical Review E*, **55**, pp. 1478–1495.
- CHAPEAU-BLONDEAU-F., AND ROJAS-VARELA-J. (2001). Estimation and Fisher information enhancement via noise addition with nonlinear sensors, *Proceedings 2nd International Symposium on Physics in Signal and Image Processing*, Marseille, France, pp. 47–50.
- CHAPEAU-BLONDEAU-F., AND ROUSSEAU-D. (2004). Noise-enhanced performance for an optimal Bayesian estimator, *IEEE Transactions on Signal Processing*, **52**, pp. 1327–1334.
- CHATTERJEE-M., AND OBA-S. I. (2005). Noise improves modulation detection by cochlear implant listeners at moderate carrier levels, *Journal of the Acoustical Society of America*, **118**, pp. 993–1002.
- CHATTERJEE-M., AND ROBERT-M. E. (2001). Noise enhances modulation sensitivity in cochlear implant listeners: Stochastic resonance in a prosthetic sensory system?, *Journal of the Association for Research in Otolaryngology*, **2**, pp. 159–171.
- CHIALVO-D. R., LONGTIN-A., AND MULLER-GERKING-J. (1997). Stochastic resonance in models of neuronal ensembles, *Physical Review E*, **55**, pp. 1798–1808.
- CHONG-C., AND KUMAR-S. P. (2003). Sensor networks: evolution, opportunities and challenges, *Proceedings of the IEEE*, **91**, pp. 1247–1256.
- CHOU-W., AND GRAY-R. M. (1991). Dithering and its effects on sigma-delta and multistage sigma-delta modulation, *IEEE Transactions on Information Theory*, **37**, pp. 500–513.
- CLARKE-B. S., AND BARRON-A. R. (1990). Information-theoretic asymptotics of Bayes methods, *IEEE Transactions on Information Theory*, **36**, pp. 453–471.
- CLARK-G. M. (1986). The University of Melbourne/Cochlear Corporation (Nucleus) program, *Otolaryngologic Clinics of North America*, **19**, pp. 329–354.
- COGDELL-J. R. (1996). *Foundations of Electrical Engineering*, second edn, Prentice Hall International, Inc.
- COLLINS-J. J. (1999). Fishing for function in noise, *Nature*, **402**, pp. 241–242.

Bibliography

- COLLINS-J. J., CHOW-C. C., AND IMHOFF-T. T. (1995). Stochastic resonance without tuning, *Nature*, **376**, pp. 236–238.
- COLLINS-J. J., CHOW-C. C., AND IMHOFF-T. T. (1995a). Aperiodic stochastic resonance in excitable systems, *Physical Review E*, **52**, pp. R3321–R3324.
- COLLINS-J. J., CHOW-C. C., AND IMHOFF-T. T. (1995b). Tuning stochastic resonance, *Nature*, **378**, pp. 341–342.
- COLLINS-J. J., CHOW-C. C., CAPELA-A. C., AND IMHOFF-T. T. (1996a). Aperiodic stochastic resonance, *Physical Review E*, **54**, pp. 5575–5583.
- COLLINS-J. J., IMHOFF-T. T., AND GRIGG-P. (1996b). Noise-enhanced information transmission in rat sa1 cutaneous mechanoreceptors via aperiodic stochastic resonance, *Journal of Neurophysiology*, **76**, pp. 642–645.
- COLLINS-J. J., PRIPLATA-A. A., GRAVELLE-D. C., NIEMI-J., HARRY-J., AND LIPSITZ-L. A. (2003). Noise-enhanced human sensorimotor control, *IEEE Engineering in Medicine and Biology Magazine*, **22**, pp. 76–83.
- COULLET-P. H., ELPHICK-C., AND TIRAPEGUI-E. (1985). Normal form of a Hopf bifurcation with noise, *Physics Letters*, **111A**, pp. 277–282.
- COURANT-R., AND ROBBINS-H. (1996). *What Is Mathematics?: An Elementary Approach to Ideas and Methods*, second edn, Oxford University Press, Oxford, UK.
- COVER-T. M., AND THOMAS-J. A. (1991). *Elements of Information Theory*, John Wiley and Sons, New York.
- DAMGAARD-P. O., AND HUFFEL-H. (1988). *Stochastic Quantization*, World Scientific, Amsterdam.
- DAMPER-R. I. (1995). *Introduction to Discrete-Time Signals and Systems*, Chapman and Hall, London.
- DAVISSON-L. D., AND LEON-GARCIA-A. (1980). A source matching approach to finding minimax codes, *IEEE Transactions on Information Theory*, **IT-26**, pp. 166–174.
- DEBNATH-G., ZHOU-T., AND MOSS-F. (1989). Remarks on stochastic resonance, *Physical Review A*, **39**, pp. 4323–4326.
- DEBYE-P. J. W. (1929). *Polar Molecules*, Chemical Catalog Co., New York.
- DELBRIDGE-A., BERNARD-J. R. L., BLAIR-D., BUTLER-S., PETERS-P., AND YALLOP-C. (1997). *The Macquarie Dictionary*, third edn, The Macquarie Library Pty Ltd.
- DEWEESE-M. (1996). Optimization principles for the neural code, *Network: Computation in Neural Systems*, **7**, pp. 325–331.
- DEWEESE-M., AND BIALEK-W. (1995). Information flow in sensory neurons, *Il Nuovo Cimento*, **17**, pp. 733–741.
- DIMITROV-A. G., AND MILLER-J. P. (2001). Neural coding and decoding: Communication channels and quantization, *Network: Computation in Neural Systems*, **12**, pp. 441–472.

- DIMITROV-A. G., MILLER-J. P., GEDEON-T., ALDWORTH-Z., AND PARKER-A. E. (2003). Analysis of neural coding through quantization with an information-based distortion measure, *Network: Computation in Neural Systems*, **14**, pp. 151–176.
- DOERING-C. R. (1995). Randomly rattled ratchets, *Il Nuovo Cimento*, **17**, pp. 685–697.
- DORMAN-M. F., AND WILSON-B. S. (2004). The design and function of cochlear implants, *American Scientist*, **92**, pp. 436–445.
- DOUGLASS-J. K., WILKENS-L., PANTAZELOU-E., AND MOSS-F. (1993). Noise enhancement of information transfer in crayfish mechanoreceptors by stochastic resonance, *Nature*, **365**, pp. 337–339.
- DRAPER-S. C., AND WORNELL-G. W. (2004). Side information aware coding strategies for sensor networks, *IEEE Journal on Selected Areas in Communications*, **22**, pp. 966–976.
- DROZHDIN-K. (2001). *Stochastic Resonance in Ferroelectric TGS Crystals*, PhD thesis, Mathematisch-Naturwissenschaftlich-Technischen Fakultät der Martin-Luther-Universität Halle-Wittenberg.
- DUNAY-R., KOLLÁR-I., AND WIDROW-B. (1998). Dithering for floating-point number representation, *1st International On-Line Workshop on Dithering in Measurement*, March 1-31, 1998, pp. 9/1–9/12.
- DYKMAN-M. I. (2002). Private communication, Email received 4 January 2002.
- DYKMAN-M. I., AND MCCLINTOCK-P. V. (1998). What can stochastic resonance do?, *Nature*, **391**, p. 344.
- DYKMAN-M. I., LUCHINSKY-D. G., MANELLA-R., MCCLINTOCK-P. V., SHORT-H. E., AND STOCKS-N. G. (1994). Noise-induced linearisation, *Physics Letters A*, **193**, pp. 61–66.
- DYKMAN-M. I., LUCHINSKY-D. G., MANNELLA-R., MCCLINTOCK-P. V., STEIN-N. D., AND STOCKS-N. G. (1995). Stochastic resonance in perspective, *Il Nuovo Cimento*, **17**, pp. 661–683.
- ECKMANN-J.-P., AND THOMAS-L. E. (1982). Remarks on stochastic resonances, *Journal of Physics A: Mathematical and General*, **15**, pp. L261–L266.
- EICHWALD-C., AND WALLECZEK-J. (1997). Aperiodic stochastic resonance with chaotic input signals in excitable systems, *Physical Review E*, **55**, pp. R6315–R6318.
- ELIASMITH-C., AND ANDERSON-C. H. (2003). *Neural Engineering: Computation, Representation, and Dynamics in Neurobiological Systems*, MIT Press, Cambridge, MA.
- FAKIR-R. (1998a). Nonstationary stochastic resonance, *Physical Review E*, **57**, pp. 6996–7001.
- FAKIR-R. (1998b). Nonstationary stochastic resonance in a single neuronlike system, *Physical Review E*, **58**, pp. 5175–5178.
- FAUVE-S., AND HESLOT-F. (1983). Stochastic resonance in a bistable system, *Physics Letters A*, **97A**, pp. 5–7.
- FITELSON-M. (1970). Asymptotic expressions for the mean and variance of the output of clipped DIMUS arrays in a field of uniform uncorrelated Gaussian noise plus signal, *The Journal of the Acoustical Society of America*, **48**, pp. 27–31.
- FRISCH-U., FROESCHLE-C., SCHEIDECKER-J. P., AND SULEM-P. L. (1973). Stochastic resonance in one-dimensional random media, *Physical Review A—General Physics*, **8**, pp. 1416–21.

Bibliography

- FRY-T. (1928). *Probability and its Engineering Uses*, D. Van Nostrand Co., New York.
- GABBIANI-F. (1996). Coding of time-varying signals in spike trains of linear and half-wave rectifying neurons, *Network: Computation in Neural Systems*, **7**, pp. 61–85.
- GABBIANI-F., AND KOCH-C. (1998). Principles of spike train analysis, in C. Koch., and I. Segev. (eds.), *Methods in Neuronal Modeling*, second edn, MIT Press, pp. 312–360.
- GAILEY-P. C., NEIMAN-A., COLLINS-J. J., AND MOSS-F. (1997). Stochastic resonance in ensembles of nondynamical elements: The role of internal noise, *Physical Review Letters*, **79**, pp. 4701–4704.
- GALDI-V., PIERRO-V., AND PINTO-I. M. (1998). Evaluation of stochastic-resonance-based detectors of weak harmonic signals in additive white Gaussian noise, *Physical Review E*, **57**, pp. 6470–6479.
- GAMMAITONI-L. (1995a). Stochastic resonance and the dithering effect in threshold physical systems, *Physical Review E*, **52**, pp. 4691–4698.
- GAMMAITONI-L. (1995b). Stochastic resonance in multi-threshold systems, *Physics Letters A*, **208**, pp. 315–322.
- GAMMAITONI-L., HÄNGGI-P., JUNG-P., AND MARCHESONI-F. (1998). Stochastic resonance, *Reviews of Modern Physics*, **70**, pp. 223–287.
- GAMMAITONI-L., MARCHESONI-F., MENICHELLA-SAETTA-E., AND SANTUCCI-S. (1989a). Stochastic resonance in bistable systems, *Physical Review Letters*, **62**, pp. 349–352.
- GAMMAITONI-L., MENICHELLA-SAETTA-E., MARCHESONI-F., AND PRESILLA-C. (1989b). Periodically time-modulated bistable systems: Stochastic resonance, *Physical Review A*, **40**, pp. 2114–2119.
- GANG-H., DITZINGER-T., NING-C. Z., AND HAKEN-H. (1993). Stochastic resonance without external periodic force, *Physical Review Letters*, **71**, pp. 807–810.
- GAUDET-V. C., AND RAPLEY-A. C. (2003). Iterative decoding using stochastic computation, *Electronics Letters*, **39**, pp. 299–301.
- GERSHENFELD-N. (1999). *The Nature of Mathematical Modeling*, Cambridge University Press, Cambridge, UK.
- GERSHO-A., AND GRAY-R. M. (1992). *Vector Quantization and Signal Compression*, Kluwer Academic Publishers.
- GINGL-Z., KISS-L. B., AND MOSS-F. (1995a). Non-dynamical stochastic resonance: Theory and experiments with white and arbitrarily coloured noise, *Europhysics Letters*, **29**, pp. 191–196.
- GINGL-Z., KISS-L. B., AND MOSS-F. (1995b). Non-dynamical stochastic resonance: Theory and experiments with white and various coloured noises, *Il Nuovo Cimento*, **17**, pp. 795–802.
- GINGL-Z., MAKRA-P., AND VAJTAI-R. (2001). High signal-to-noise ratio gain by stochastic resonance in a double well, *Fluctuation and Noise Letters*, **1**, pp. L181–L188.
- GINGL-Z., VAJTAI-R., AND KISS-L. B. (2000). Signal-to-noise ratio gain by stochastic resonance in a bistable system, *Chaos, Solitons and Fractals*, **11**, pp. 1929–1932.

- GODIVIER-X., AND CHAPEAU-BLONDEAU-F. (1997). Noise-assisted signal transmission by a nonlinear electronic comparator: Experiment and theory, *Signal Processing*, **56**, pp. 293–303.
- GODIVIER-X., AND CHAPEAU-BLONDEAU-F. (1998). Stochastic resonance in the information capacity of a nonlinear dynamic system, *International Journal of Bifurcation and Chaos*, **8**, pp. 581–589.
- GODIVIER-X., ROJAS-VARELA-J., AND CHAPEAU-BLONDEAU-F. (1997). Noise-assisted signal transmission via stochastic resonance in a diode nonlinearity, *Electronics Letters*, **33**, pp. 1666–1668.
- GÓRA-P. F. (2003). Array enhanced stochastic resonance and spatially correlated noises, Preprint arXiv:cond-mat/0308620.
- GOYCHUK-I. (2001). Information transfer with rate-modulated Poisson processes: A simple model for nonstationary stochastic resonance, *Physical Review E*, **64**, Art. No. 021909.
- GOYCHUK-I., AND HÄNGGI-P. (1999). Quantum stochastic resonance in parallel, *New Journal of Physics*, **1**, pp. 14.1–14.14.
- GRAY-R. M., AND NEUHOFF-D. L. (1998). Quantization, *IEEE Transactions on Information Theory*, **44**, pp. 2325–2383.
- GRAY-R. M., AND STOCKHAM-T. G. (1993). Dithered quantizers, *IEEE Transactions on Information Theory*, **39**, pp. 805–812.
- GREENWOOD-P. E., MILLER-U. U., WARD-L. M., AND WEFELMEYER-W. (2003). Statistical analysis of stochastic resonance in a thresholded detector, *Austrian Journal of Statistics*, **32**, pp. 49–70.
- GREENWOOD-P. E., WARD-L. M., AND WEFELMEYER-W. (1999). Statistical analysis of stochastic resonance in a simple setting, *Physical Review E*, **60**, pp. 4687–4695.
- GREENWOOD-P. E., WARD-L. M., RUSSELL-D. F., NEIMAN-A., AND MOSS-F. (2000). Stochastic resonance enhances the electrosensory information available to paddlefish for prey capture, *Physical Review Letters*, **84**, pp. 4773–4776.
- GRUNWALD-P. D., MYUNG-I. J., AND PITT-M. A. (2005). *Advances in Minimum Description Length: Theory and Applications (Neural Information Processing)*, The MIT Press.
- GYÖRGY-A., AND LINDER-T. (2000). Optimal entropy-constrained scalar quantization of a uniform source, *IEEE Transactions on Information Theory*, **46**, pp. 2702–2711.
- HÄNGGI-P. (2002). Stochastic resonance in biology: How noise can enhance detection of weak signals and help improve biological information processing, *Chemphyschem*, **3**, pp. 285–290.
- HÄNGGI-P., INCHIOSA-M. E., FOGLIATTI-D., AND BULSARA-A. R. (2000). Nonlinear stochastic resonance: The saga of anomalous output-input gain, *Physical Review E*, **62**, pp. 6155–6163.
- HARMER-G. P. (2001). *Stochastic Processing for Enhancement of Artificial Insect Vision*, PhD thesis, Department of Electrical and Electronic Engineering, Adelaide University, Australia.
- HARMER-G. P., AND ABBOTT-D. (1999). Losing strategies can win by Parrondo's paradox, *Nature*, **402**, p. 864.
- HARMER-G. P., AND ABBOTT-D. (2001). Motion detection and stochastic resonance in noisy environments, *Microelectronics Journal*, **32**, pp. 959–967.

Bibliography

- HARMER-G. P., DAVIS-B. R., AND ABBOTT-D. (2002). A review of stochastic resonance: Circuits and measurement, *IEEE Transactions on Instrumentation and Measurement*, **51**, pp. 299–309.
- HARRINGTON-J. V. (1955). An analysis of the detection of repeated signals in noise by binary integration, *IRE Transactions on Information Theory*, **1**, pp. 1–9.
- HARRY-J. D., NIEMI-J. B., PRIPLATA-A. A., AND COLLINS-J. J. (2005). Balancing act (noise based sensory enhancement technology), *IEEE Spectrum*, **42**, pp. 36–41.
- HENEGHAN-C., CHOW-C. C., COLLINS-J. J., IMHOFF-T. T., LOWEN-S. B., AND TEICH-M. C. (1996). Information measures quantifying aperiodic stochastic resonance, *Physical Review E*, **54**, pp. R2228–R2231.
- HIBBS-A. D., SINGSAAS-A. L., JACOBS-E. W., BULSARA-A. R., BEKKEDAHL-J. J., AND MOSS-F. (1995). Stochastic resonance in a superconducting loop with a Josephson-junction, *Journal of Applied Physics*, **77**, pp. 2582–2590.
- HOCH-T., WENNING-G., AND OBERMAYER-K. (2003a). Adaptation using local information for maximizing the global cost, *Neurocomputing*, **52-54**, pp. 541–546.
- HOCH-T., WENNING-G., AND OBERMAYER-K. (2003b). Optimal noise-aided signal transmission through populations of neurons, *Physical Review E*, **68**, Art. No. 011911.
- HOHN-N. (2001). *Stochastic Resonance in a Neuron Model with Applications to the Auditory Pathway*, PhD thesis, Department of Otorhynology, The University of Melbourne, Australia.
- HOHN-N., AND BURKITT-A. N. (2001). Shot noise in the leaky integrate-and-fire neuron, *Physical Review E*, **63**, Art. No. 031902.
- HORSTHEMKE-W., AND LEFEVER-R. (1980). Voltage-noise-induced transitions in electrically excitable membranes, *Biophysical Journal*, **35**, pp. 415–432.
- HU-G., GONG-D. C., WEN-X. D., YANG-C. Y., QING-G. R., AND LI-R. (1992). Stochastic resonance in a nonlinear-system driven by an aperiodic force, *Physical Review A*, **46**, pp. 3250–3254.
- IANNELLI-J. M., YARIV-A., CHEN-T. R., AND ZHUANG-Y. H. (1994). Stochastic resonance in a semiconductor distributed feedback laser, *Applied Physics Letters*, **65**, pp. 1983–1985.
- IMKELLER-P., AND PAVLYUKEVICH-I. (2001). Stochastic resonance in two-state markov chains, *Archiv der Mathematik*, **77**, pp. 107–115.
- INCHIOSA-M. E., AND BULSARA-A. R. (1995). Nonlinear dynamic elements with noisy sinusoidal forcing: Enhancing response via nonlinear coupling, *Physical Review E*, **52**, pp. 327–339.
- INCHIOSA-M. E., AND BULSARA-A. R. (1996). Signal detection statistics of stochastic resonators, *Physical Review E*, **53**, pp. R2021–2024.
- INCHIOSA-M. E., AND BULSARA-A. R. (1998). DC signal detection via dynamical asymmetry in a nonlinear device, *Physical Review E*, **58**, pp. 115–127.
- INCHIOSA-M. E., ROBINSON-J. W. C., AND BULSARA-A. R. (2000). Information-theoretic stochastic resonance in noise-floor limited systems: The case for adding noise, *Physical Review Letters*, **85**, pp. 3369–3372.

-
- IYENGAR-S. S., AND BROOKS-R. R. (2004). Special issue introduction—the road map for distributed sensor networks in the context of computing and communication, *Journal of Parallel and Distributed Computing*, **64**, pp. 785–787.
- JARAMILLO-F., AND WIESENFELD-K. (1998). Mechanoelectrical transduction assisted by Brownian motion: a role for noise in the auditory system, *Nature Neuroscience*, **1**, pp. 384–388.
- JAYNES-E. T. (2003). *Probability Theory: The Logic of Science*, Cambridge University Press, Cambridge, UK.
- JOHNSON-D. H. (2002). Four top reasons mutual information does not quantify neural information processing, *Annual Computational Neuroscience Meeting* (Unpublished).
- JUNG-P. (1993). Periodically driven stochastic systems, *Physics Reports - Review Section of Physics Letters*, **234**, pp. 175–295.
- JUNG-P. (1994). Threshold devices: Fractal noise and neural talk, *Physical Review E*, **50**, pp. 2513–2522.
- JUNG-P. (1995). Stochastic resonance and optimal design of threshold detectors, *Physics Letters A*, **207**, pp. 93–104.
- JUNG-P., AND MAYER-KRESS-G. (1995). Stochastic resonance in threshold devices, *Il Nuovo Cimento*, **17**, pp. 827–834.
- KALMYKOV-Y. P., COFFEY-W. T., AND TITOV-S. V. (2004). Bimodal approximation for anomalous diffusion in a potential, *Physical Review E*, **69**, Art. No. 021105.
- KANDEL-E. R., SCHWARTZ-J. H., AND JESSELL-T. M. (1991). *Principles of Neural Science*, third edn, Elsevier, New York.
- KANEFSKY-M. (1966). Detection of weak signals with polarity coincidence arrays, *IEEE Transactions on Information Theory*, **IT-12**, pp. 260–268.
- KAY-S. (2000). Can detectability be improved by adding noise?, *IEEE Signal Processing Letters*, **7**, pp. 8–10.
- KHOVANOV-I. A., AND MCCLINTOCK-P. V. (2003). Comment on “signal-to-noise ratio gain in neuronal systems”, *Physical Review E*, **67**, Art. No. 043901.
- KIKKERT-C. J. (1995). Improving ADC performance by adding dither, *Proc. 13th Australian Microelectronics Conference, Microelectronics: Technology Today for the Future (MICRO '95)*, IREE Soc, Milsons Point, NSW, Australia, pp. 97–102.
- KIRKPATRICK-S., GELATT-C. D., AND VECCHI-M. P. (1983). Optimisation by simulated annealing, *Science*, **220**, pp. 671–680.
- KISH-L. B. (2002). Private communication, Email received October 2002.
- KISH-L. B., HARMER-G. P., AND ABBOTT-D. (2001). Information transfer rate of neurons: Stochastic resonance of Shannon’s information channel capacity, *Fluctuation and Noise Letters*, **1**, pp. L13–L19.

Bibliography

- KISS-L. B. (1996). Possible breakthrough: Significant improvement of signal to noise ratio by stochastic resonance, in R. Katz. (ed.), *Chaotic, Fractal and Nonlinear Signal Processing*, Vol. 375, American Institute of Physics, pp. 382–396.
- KOSKO-B., AND MITAIM-S. (2004). Robust stochastic resonance for simple threshold neurons, *Physical Review E*, **70**, Art. No. 031911.
- KREYSZIG-E. (1988). *Advanced Engineering Mathematics*, 6th edn, John Wiley and Sons, New York.
- LATHI-B. P. (1998). *Modern Digital and Analog Communication Systems*, 3rd edn, Oxford University Press.
- LAUGHLIN-S. B. (2001). Energy as a constraint on the coding and processing of sensory information, *Current Opinion in Neurobiology*, **11**, pp. 475–480.
- LEHMANN-E. L., AND CASELLA-G. (1998). *Theory of Point Estimation*, Springer, New York.
- LEONARD-D. S., AND REICHL-L. E. (1994). Stochastic resonance in a chemical-reaction, *Physical Review E*, **49**, pp. 1734–1737.
- LEVIN-J. E., AND MILLER-J. P. (1996). Broadband neural encoding in the cricket cercal sensory system enhanced by stochastic resonance, *Nature*, **380**, pp. 165–168.
- LEVY-W. B., AND BAXTER-R. A. (2002). Energy-efficient neuronal computation via quantal synaptic failures, *The Journal of Neuroscience*, **22**, pp. 4746–4755.
- LIM-M., AND SALOMA-C. (2001). Noise-enhanced measurement of weak doublet spectra with a Fourier-transform spectrometer and a 1-bit analog-to-digital converter, *Applied Optics*, **40**, pp. 1767–1775.
- LINDER-T., AND ZAMIR-R. (1999). High-resolution source coding for non-difference distortion measures: The rate-distortion function, *IEEE Transactions on Information Theory*, **45**, pp. 533–547.
- LINDER-T., LUGOSI-G., AND ZEGER-K. (1994). Recent trends in lossy source coding, *Journal on Communications (Hungary)*, **45**, pp. 16–22.
- LINDNER-J. F., MEADOWS-B. K., DITTO-W. L., INCHIOSA-M. E., AND BULSARA-A. R. (1995). Array enhanced stochastic resonance and spatiotemporal synchronization, *Physical Review Letters*, **75**, pp. 3–6.
- LIU-F., YU-Y., AND WANG-W. (2001). Signal-to-noise ratio gain in neuronal systems, *Physical Review E*, **63**, Art. No. 051912.
- LLOYD-S. P. (1982). Least squares quantization in PCM, *IEEE Transactions on Information Theory*, **IT-28**, pp. 129–137.
- LOERINCZ-K., GINGL-Z., AND KISS-L. B. (1996). A stochastic resonator is able to greatly improve signal-to-noise ratio, *Physics Letters A*, **224**, pp. 63–67.
- LONGTIN-A. (1993). Stochastic resonance in neuron models, *Journal of Statistical Physics*, **70**, pp. 309–327.

- LONGTIN-A., BULSARA-A., AND MOSS-F. (1991). Time-interval sequences in bistable systems and the noise-induced transmission of information by sensory neurons, *Physical Review Letters*, **67**, pp. 656–659.
- LUCHINSKY-D. G., MANNELLA-R., MCCLINTOCK-P. V., AND STOCKS-N. G. (1999). Stochastic resonance in electrical circuits - I: Conventional stochastic resonance, *IEEE Transactions on Circuits and Systems-II: Analog and Digital Signal Processing*, **46**, pp. 1205–1214.
- LUCHINSKY-D. G., MANNELLA-R., MCCLINTOCK-P. V., AND STOCKS-N. G. (1999). Stochastic resonance in electrical circuits - II: Nonconventional stochastic resonance, *IEEE Transactions on Circuits and Systems-II: Analog and Digital Signal Processing*, **46**, pp. 1215–1224.
- LUCHINSKY-D. G., MCCLINTOCK-P. V., AND DYKMAN-M. I. (1998). Analogue studies of nonlinear systems, *Reports on Progress in Physics*, **61**, pp. 889–997.
- MAKRA-P., GINGL-Z., AND KISH-L. B. (2002). Signal-to-noise ratio gain in non-dynamical and dynamical bistable stochastic resonators, *Fluctuation and Noise Letters*, **2**, pp. L147–L155.
- MANWANI-A., AND KOCH-C. (1998). Synaptic transmission: An information-theoretic perspective, in M. Jordan., M. Kearns., and S. Solla. (eds.), *Advances in Neural Information Processing Systems*, Vol. 10, pp. 201–207.
- MANWANI-A., AND KOCH-C. (1999). Detecting and estimating signals in noisy cable structures, ii: Information theoretical analysis, *Neural Computation*, **11**, pp. 1831–1873.
- MAO-X. M., SUN-K., AND OUYANG-Q. (2002). Stochastic resonance in a financial model, *Chinese Physics*, **11**, pp. 1106–1110.
- MARTINEZ-K., HART-J. K., AND ONG-R. (2004). Sensor network applications, *Computer*, **37**, pp. 50–56.
- MARTORELL-F., MCDONNELL-M. D., ABBOTT-D., AND RUBIO-A. (2004). Generalized noise resonance: Using noise for signal enhancement, in D. Abbott., S. M. Bezrukov., A. Der., and A. Sánchez. (eds.), *Proc. SPIE Fluctuations and Noise in Biological, Biophysical, and Biomedical Systems II*, Vol. 5467, pp. 163–174.
- MARTORELL-F., MCDONNELL-M. D., RUBIO-A., AND ABBOTT-D. (2005). Using noise to break the noise barrier in circuits, in S. F. Al-Sarawi. (ed.), *Proc. SPIE Smart Structures, Devices, and Systems II*, Vol. 5649, pp. 53–66.
- MATSUMOTO-K., AND TSUDA-I. (1985). Information theoretical approach to noisy dynamics, *Journal of Physics A: Mathematical and General*, **18**, pp. 3561–3566.
- MAX-J. (1960). Quantizing for minimum distortion, *IRE Transactions on Information Theory*, **IT-6**, pp. 7–12.
- MCCULLOCH-W. S., AND PITTS-W. (1943). A logical calculus of ideas immanent in nervous activity, *Bulletin of Mathematical Biophysics*, **5**, pp. 115–133.
- MCDONNELL-M. D., ABBOTT-D., AND PEARCE-C. (2003a). Neural mechanisms for analog to digital conversion, in D. V. Nicolau. (ed.), *Proc. SPIE BioMEMS and Nanoengineering*, Vol. 5275, pp. 278–286.

Bibliography

- MCDONNELL-M. D., ABBOTT-D., AND PEARCE-C. E. M. (2002a). An analysis of noise enhanced information transmission in an array of comparators, *Microelectronics Journal*, **33**, pp. 1079–1089.
- MCDONNELL-M. D., ABBOTT-D., AND PEARCE-C. E. M. (2002b). A characterization of suprathreshold stochastic resonance in an array of comparators by correlation coefficient, *Fluctuation and Noise Letters*, **2**, pp. L205–L220.
- MCDONNELL-M. D., AND ABBOTT-D. (2002). Open questions for suprathreshold stochastic resonance in sensory neural models for motion detection using artificial insect vision, in S. M. Bezrukov. (ed.), *UPoN 2002: Third International Conference on Unsolved Problems of Noise and Fluctuations in Physics, Biology, and High Technology*, Vol. 665, American Institute of Physics, pp. 51–58.
- MCDONNELL-M. D., AND ABBOTT-D. (2004a). Optimal quantization in neural coding, *Proc. IEEE International Symposium on Information Theory*, Chicago, USA, p. 496.
- MCDONNELL-M. D., AND ABBOTT-D. (2004b). Signal reconstruction via noise through a system of parallel threshold nonlinearities, *Proc. 2004 IEEE International Conference on Acoustics, Speech, and Signal Processing*, Vol. 2, Montreal, Canada, pp. 809–812.
- MCDONNELL-M. D., PEARCE-C. E. M., AND ABBOTT-D. (2001). Neural information transfer in a noisy environment, in N. W. Bergmann. (ed.), *Proc. SPIE Electronics and Structures for MEMS II*, Vol. 4591, pp. 59–69.
- MCDONNELL-M. D., SETHURAMAN-S., KISH-L. B., AND ABBOTT-D. (2004a). Cross-spectral measurement of neural signal transfer, in Z. Gingl, J. M. Sancho, L. Schimansky-Geier, and J. Kertesz. (eds.), *Proc. SPIE Noise in Complex Systems and Stochastic Dynamics II*, Vol. 5471, pp. 550–559.
- MCDONNELL-M. D., STOCKS-N. G., PEARCE-C. E. M., AND ABBOTT-D. (2002c). Maximising information transfer through nonlinear noisy devices, in D. V. Nicolau. (ed.), *Proc. SPIE Biomedical Applications of Micro and Nanoengineering*, Vol. 4937, pp. 254–263.
- MCDONNELL-M. D., STOCKS-N. G., PEARCE-C. E. M., AND ABBOTT-D. (2003b). The data processing inequality and stochastic resonance, in L. Schimansky-Geier, D. Abbott, A. Neiman, and C. Van den Broeck. (eds.), *Proc. SPIE Noise in Complex Systems and Stochastic Dynamics*, Vol. 5114, pp. 249–260.
- MCDONNELL-M. D., STOCKS-N. G., PEARCE-C. E. M., AND ABBOTT-D. (2003c). Stochastic resonance and the data processing inequality, *Electronics Letters*, **39**, pp. 1287–1288.
- MCDONNELL-M. D., STOCKS-N. G., PEARCE-C. E. M., AND ABBOTT-D. (2004b). Optimal quantization for energy-efficient information transfer in a population of neuron-like devices, in Z. Gingl, J. M. Sancho, L. Schimansky-Geier, and J. Kertesz. (eds.), *Proc. SPIE Noise in Complex Systems and Stochastic Dynamics II*, Vol. 5471, pp. 222–232.
- MCDONNELL-M. D., STOCKS-N. G., PEARCE-C. E. M., AND ABBOTT-D. (2005a). Analog to digital conversion using suprathreshold stochastic resonance, in S. F. Al-Sarawi. (ed.), *Proc. SPIE Smart Structures, Devices, and Systems II*, Vol. 5649, pp. 75–84.
- MCDONNELL-M. D., STOCKS-N. G., PEARCE-C. E. M., AND ABBOTT-D. (2005b). How to use noise to reduce complexity in quantization, in A. Bender. (ed.), *Proc. SPIE Complex Systems*, Vol. 6039, pp. 115–126.

- MCDONNELL-M. D., STOCKS-N. G., PEARCE-C. E. M., AND ABBOTT-D. (2005c). Optimal quantization and suprathreshold stochastic resonance, in N. G. Stocks., D. Abbott., and R. P. Morse. (eds.), *Proc. SPIE Noise in Biological, Biophysical, and Biomedical Systems III*, Vol. 5841, pp. 164–173.
- MCDONNELL-M. D., STOCKS-N. G., PEARCE-C. E. M., AND ABBOTT-D. (2005d). Quantization in the presence of large amplitude threshold noise, *Fluctuation and Noise Letters*, **5**, pp. L457–L468.
- MCDONNELL-M. D., STOCKS-N. G., PEARCE-C. E. M., AND ABBOTT-D. (2006). Optimal information transmission in nonlinear arrays through suprathreshold stochastic resonance, To Appear in *Physics Letters A* (accepted 16 Nov. 2005).
- MCNAMARA-B., AND WIESENFELD-K. (1989). Theory of stochastic resonance, *Physical Review A*, **39**, pp. 4854–4869.
- MCNAMARA-B., WIESENFELD-K., AND ROY-R. (1988). Observation of stochastic resonance in a ring laser, *Physical Review Letters*, **60**, pp. 2626–2629.
- MELNIKOV-V. I. (1993). Schmitt trigger: A solvable model of stochastic resonance, *Physical Review E*, **48**, pp. 2481–2489.
- MICKAN-S., ABBOTT-D., MUNCH-J., ZHANG-X.-C., AND VAN DOORN-T. (2000). Analysis of system trade-offs for terahertz imaging, *Microelectronics Journal*, **31**, pp. 503–514.
- MICKAN-S. P., AND ZHANG-X. C. (2003). T-ray sensing and imaging, *International Journal of High Speed Electronics and Systems*, **13**, pp. 606–676.
- MINGESZ-R., MAKRA-P., AND GINGL-Z. (2005). Cross-spectral analysis of signal improvement by stochastic resonance in bistable systems, in L. B. Kish., K. Lindenberg., and Z. Gingl. (eds.), *Proc. SPIE Noise in Complex Systems and Stochastic Dynamics III*, Vol. 5845, pp. 283–292.
- MITAIM-S., AND KOSKO-B. (1998). Adaptive stochastic resonance, *Proceedings of the IEEE*, **86**, pp. 2152–2183.
- MORSE-R. P., ALLINGHAM-D., AND STOCKS-N. G. (2002). An information-theoretic approach to cochlear implant coding, in S. M. Bezrukov. (ed.), *UPoN 2002: Third International Conference on Unsolved Problems of Noise and Fluctuations in Physics, Biology, and High Technology*, Vol. 665, American Institute of Physics, pp. 125–132.
- MORSE-R. P., AND MEYER-G. F. (2000). The practical use of noise to improve speech coding by analogue cochlear implants, *Chaos, Solitons and Fractals*, **11**, pp. 1885–1894.
- MORSE-R. P., AND ROPER-P. (2000). Enhanced coding in a cochlear-implant model using additive noise: Aperiodic stochastic resonance with tuning, *Physical Review E*, **61**, pp. 5683–5692.
- MORSE-R. P., AND STOCKS-N. G. (2005). Enhanced cochlear implant coding using multiplicative noise, in N. G. Stocks., D. Abbott., and R. P. Morse. (eds.), *Proc. SPIE Fluctuations and Noise in Biological, Biophysical and Biomedical Systems III*, Vol. 5841, pp. 23–30.
- MOSS-F., AND PEI-X. (1995). Neurons in parallel, *Nature*, **376**, pp. 211–212.
- MOSS-F., AND WIESENFELD-K. (1995). The benefits of background noise, *Scientific American*, **273**, pp. 50–53.

Bibliography

- MOSS-F., DOUGLASS-J. K., WILKENS-L., PIERSON-D., AND PANTAZELOU-E. (1993). Stochastic resonance in an electronic FitzHugh-Nagumo model, in J. R. Buchler, and H. E. Kandrup. (eds.), *Stochastic Processes in Astrophysics*, Vol. 706, Annals of the New York Academy of Sciences, pp. 26–41.
- MOSS-F., PIERSON-D., AND OGORMAN-D. (1994). Stochastic resonance - tutorial and update, *International Journal of Bifurcation and Chaos*, **4**, pp. 1383–1397.
- MOSS-F., WARD-L. M., AND SANNITA-W. G. (2004). Stochastic resonance and sensory information processing: a tutorial and review of application, *Clinical Neurophysiology*, **115**, pp. 267–281.
- NADAL-J., AND PARGA-N. (1994). Nonlinear neurons in the low noise limit: A factorial code maximizes information-transfer, *Network: Computation in Neural Systems*, **5**, pp. 565–581.
- NARAYANAN-R. M., AND KUMRU-C. (2005). Implementation of fully polarimetric random noise radar, *IEEE Antennas and Wireless Propagation Letters*, **4**, pp. 125–128.
- NEIMAN-A. (1994). Synchronizationlike phenomena in coupled stochastic bistable systems, *Physical Review E*, **49**, pp. 3484–3487.
- NEIMAN-A., SCHIMANSKY-GEIER-L., AND MOSS-F. (1997). Linear response theory applied to stochastic resonance in models of ensembles of oscillators, *Physical Review E*, **56**, pp. R9–R12.
- NEIMAN-A., SHULGIN-D., ANISHCHENKO-V., EBELING-W., SCHIMANSKY-GEIER-L., AND FREUND-J. (1996). Dynamical entropies applied to stochastic resonance, *Physical Review Letters*, **76**, pp. 4299–4302.
- NICOLIS-C. (1981). Solar variability and stochastic effects on climate, *Solar Physics*, **74**, pp. 473–478.
- NICOLIS-C. (1982). Stochastic aspects of climatic transitions- response to a periodic forcing, *Tellus*, **34**, pp. 1–9.
- NIKITIN-A., AND STOCKS-N. G. (2004). The application of Gaussian channel theory to the estimation of information transmission rates in neural systems, in D. Abbott., S. M. Bezrukov., A. Der., and A. Sánchez. (eds.), *Proc. SPIE Fluctuations and Noise in Biological, Biophysical, and Biomedical Systems II*, Vol. 5467, Maspalomas, Spain, pp. 202–211.
- NOCEDAL-J., AND WRIGHT-S. J. (1999). *Numerical Optimization*, Springer-Verlag, New York.
- NOEST-A. J. (1995). Tuning stochastic resonance, *Nature*, **378**, p. 341.
- ØKSENDAL-B. (1998). *Stochastic Differential Equations: An Introduction With Applications*, fifth edn, Springer, Berlin.
- OLIAEI-O. (2003). Stochastic resonance in sigma-delta modulators, *Electronics Letters*, **39**, pp. 173–174.
- PARADISO-J. A., AND STARNER-T. (2005). Energy scavenging for mobile and wireless electronics, *IEEE Pervasive Computing*, **4**, pp. 18–27.
- PAULSSON-J., AND EHRENBERG-M. (2000). Random signal fluctuations can reduce random fluctuations in regulated components of chemical regulatory networks, *Physical Review Letters*, **84**, pp. 5447–5450.

- PAULSSON-J., BERG-O. G., AND EHRENBERG-M. (2000). Stochastic focusing: Fluctuation-enhanced sensitivity of intracellular regulation, *Proceedings of the National Academy of Sciences of the USA*, **97**, pp. 7148–7153.
- PEI-X., L. WILKENS., AND MOSS-F. (1996). Noise mediated spike timing precision from aperiodic stimuli in an array of Hodgkin-Huxley-type neurons, *Physical Review Letters*, **77**, pp. 4679–4682.
- PETRACCHI-D. (2000). What is the role of stochastic resonance?, *Chaos, Solitons and Fractals*, **11**, pp. 1827–1834.
- PETRACCHI-D., GEBESHUBER-I. C., DEFELICE-L. J., AND HOLDEN-A. V. (2000). Introduction: Stochastic resonance in biological systems, *Chaos, Solitons and Fractals*, **11**, pp. 1819–1822.
- PINSKER-M. S. (1964). *Information and Information Stability of Random Variables and Processes*, Holden-Day, San Francisco, CA.
- PLASKOTA-L. (1996a). How to benefit from noise, *Journal of Complexity*, **12**, pp. 175–184.
- PLASKOTA-L. (1996b). Worst case complexity of problems with random information noise, *Journal of Complexity*, **12**, pp. 416–439.
- POOR-H. V. (1994). *An Introduction to Signal Detection and Estimation*, 2nd edn, Springer-Verlag, New York.
- PRADHAN-S. S., AND RAMCHANDRAN-K. (2005). Generalized coset codes for distributed binning, *IEEE Transactions on Information Theory*, **51**, pp. 3457–3474.
- PRADHAN-S. S., KUSUMA-J., AND RAMCHANDRAN-K. (2002). Distributed compression in a dense microsensor network, *IEEE Signal Processing Magazine*, **19**, pp. 51–60.
- PRESS-W. H., TEUKOLSKY-S. A., VETTERLING-W. T., AND FLANNERY-B. P. (1992). *Numerical Recipes in C: The Art of Scientific Computing*, second edn, Cambridge University Press, Cambridge, UK.
- PRICE-R. (1958). A useful theorem for nonlinear devices having Gaussian inputs, *IRE Transactions on Information Theory*, **IT-4**, pp. 69–72.
- PRIPLATA-A. A., NIEMI-J. B., HARRY-J. D., LIPSITZ-L. A., AND COLLINS-J. J. (2003). Vibrating insoles and balance control in elderly people, *Lancet*, **362**, pp. 1123–1124.
- PRIPLATA-A., NIEMI-J., SALEN-M., HARRY-J., LIPSITZ-L. A., AND COLLINS-J. J. (2004). Noise-enhanced human balance control, *Physical Review Letters*, **89**, Art. No. 238101.
- PROAKIS-J. G., AND MANOLAKIS-D. G. (1996). *Digital Signal Processing: Principles, Algorithms, and Applications*, 3rd edn, Prentice Hall.
- PROAKIS-J. G., AND SALEHI-M. (1994). *Communication Systems Engineering*, Prentice Hall International, Inc.
- PRUDNIKOV-A. P., BRYCHKOV-Y. A., AND MARICHEV-O. I. (1986). *Integrals and Series*, Vol. 2- Special Functions, Gordon and Breach Science Publishers.
- REMLEY-W. R. (1966). Some effects of clipping in array processing, *The Journal of the Acoustical Society of America*, **39**, pp. 702–707.

Bibliography

- RÉNYI-A. (1970). *Probability Theory*, North-Holland Publishing Company, Amsterdam.
- RICE-S. O. (1944). Mathematical analysis of random noise, *The Bell System Technical Journal*, **23**, pp. 282–332.
- RICE-S. O. (1945). Mathematical analysis of random noise, *The Bell System Technical Journal*, **24**, pp. 46–156.
- RICE-S. O. (1948). Statistical properties of a sine wave plus random noise, *The Bell System Technical Journal*, **27**, pp. 109–157.
- RIEKE-F., WARLAND-D., DE RUYTER VAN STEVENINCK-R., AND BIALEK-W. (1997). *Spikes: Exploring the Neural Code*, MIT Press, Cambridge, MA.
- RISSANEN-J. J. (1996). Fisher information and stochastic complexity, *IEEE Transactions on Information Theory*, **42**, pp. 40–47.
- ROBINSON-J. W. C., ASRAF-D. E., BULSARA-A. R., AND INCHIOSA-M. E. (1998). Information-theoretic distance measures and a generalization of stochastic resonance, *Physical Review Letters*, **81**, pp. 2850–2853.
- ROBINSON-J. W. C., RUNG-J., BULSARA-A. R., AND INCHIOSA-M. E. (2001). General measures for signal-noise separation in nonlinear dynamical systems, *Physical Review E*, **63**, Art. No. 011107.
- ROSE-K. (1994). A mapping approach to rate-distortion computation and analysis, *IEEE Transactions on Information Theory*, **40**, pp. 1939–1952.
- ROSE-K. (1998). Deterministic annealing for clustering, compression, classification, regression and related optimization problems, *Proceedings of the IEEE*, **86**, pp. 2210–2239.
- ROSE-K., GUREWITZ-E., AND FOX-G. C. (1990). Statistical mechanics and phase transitions in clustering, *Physical Review Letters*, **65**, pp. 945–948.
- ROSE-K., GUREWITZ-E., AND FOX-G. C. (1992). Vector quantization by deterministic annealing, *IEEE Transactions on Information Theory*, **38**, pp. 1249–1257.
- ROUSSEAU-D., AND CHAPEAU-BLONDEAU-F. (2004). Suprathreshold stochastic resonance and signal-to-noise ratio improvement in arrays of comparators, *Physics Letters A*, **321**, pp. 280–290.
- ROUSSEAU-D., AND CHAPEAU-BLONDEAU-F. (2005). Constructive role of noise in signal detection from parallel arrays of quantizers, *Signal Processing*, **85**, pp. 571–580.
- ROUSSEAU-D., DUAN-F., AND CHAPEAU-BLONDEAU-F. (2003). Suprathreshold stochastic resonance and noise-enhanced Fisher information in arrays of threshold devices, *Physical Review E*, **68**, Art. No. 031107.
- ROZENFELD-R., AND SCHIMANSKY-GEIER-L. (2000). Array-enhanced stochastic resonance in finite systems, *Chaos, Solitons and Fractals*, **11**, pp. 1937–1944.
- RUBINSTEIN-J. T., AND HONG-R. (2003). Signal coding in cochlear implants: exploiting stochastic effects of electrical stimulation, *Annals of Otolaryngology and Laryngology*, **112**, pp. 14–19.
- RUDNICK-P. (1960). Small signal detection in the DIMUS array, *The Journal of the Acoustical Society of America*, **32**, pp. 871–877.

- SATO-A., UEDA-M., AND MUNAKATA-T. (2004). Signal estimation and threshold optimization using an array of bithreshold elements, *Physical Review E*, **70**, Art. No. 021106.
- SCHREIBER-S., MACHENS-C. K., HERZ-A. V., AND LAUGHLIN-S. B. (2002). Energy-efficient coding with discrete stochastic events, *Neural Computation*, **14**, pp. 1323–1346.
- SETHURAMAN-S., AND KISH-L. B. (2003). Cross spectra measure of neural signals and noise, in S. M. Bezrukov., H. Frauenfelder., and F. Moss. (eds.), *Proc. SPIE Fluctuations and Noise in Biological and Biomedical Systems*, Vol. 5110, pp. 244–251.
- SHANNON-C. E. (1948). A mathematical theory of communication, *The Bell System Technical Journal*, **27**, pp. 379–423, 623–656.
- SHATOKHIN-V., WELLENS-T., AND BUCHLEITNER-A. (2004). The noise makes the signal: What a small fry should know about stochastic resonance, *Journal of Modern Optics*, **51**, pp. 851–860.
- SHIRYAEV-A. N. (1996). *Probability*, second edn, Springer, New York.
- SO-H. C., CHAN-Y. T., MA-Q., AND CHING-P. C. (1999). Comparison of various periodograms for sinusoid detection and frequency estimation, *IEEE Transactions on Aerospace and Electronic Systems*, **35**, pp. 945–952.
- SPIEGEL-M. R., AND LIU-J. (1999). *Mathematical Handbook of Formulas and Tables*, McGraw-Hill.
- STOCKS-N. G. (2000a). Suprathreshold stochastic resonance in multilevel threshold systems, *Physical Review Letters*, **84**, pp. 2310–2313.
- STOCKS-N. G. (2000b). Optimising information transmission in model neuronal ensembles: The role of internal noise, in J. A. Freund., and T. Poschel. (eds.), *Future Directions for Intelligent Systems and Information Sciences*, Springer-Verlag, pp. 150–159.
- STOCKS-N. G. (2000c). Suprathreshold stochastic resonance, in D. S. Broomhead., E. A. Luchinskaya., P. V. McClintock., and T. Mullin. (eds.), *Stochastic and Chaotic Dynamics in the Lakes (STOCHAOS)*, Vol. 502, American Institute of Physics, pp. 415–421.
- STOCKS-N. G. (2001a). Information transmission in parallel threshold arrays: Suprathreshold stochastic resonance, *Physical Review E*, **63**, Art. No. 041114.
- STOCKS-N. G. (2001b). Information transmission in parallel threshold networks: Suprathreshold stochastic resonance and coding efficiency, in G. Bosman. (ed.), *Noise in Physical Systems and 1/f Fluctuations, ICNF 2001*, pp. 594–597.
- STOCKS-N. G. (2001c). Suprathreshold stochastic resonance: An exact result for uniformly distributed signal and noise, *Physics Letters A*, **279**, pp. 308–312.
- STOCKS-N. G., ALLINGHAM-D., AND MORSE-R. P. (2002). The application of suprathreshold stochastic resonance to cochlear implant coding, *Fluctuation and Noise Letters*, **2**, pp. L169–L181.
- STOCKS-N. G., AND MANNELLA-R. (2000). Suprathreshold stochastic resonance in a neuronal network model: A possible strategy for sensory coding, in N. Kosobov. (ed.), *Stochastic Processes in Physics, Chemistry and Biology*, Physica-Verlag, pp. 236–247.

Bibliography

- STOCKS-N. G., AND MANNELLA-R. (2001). Generic noise enhanced coding in neuronal arrays, *Physical Review E*, **64**, Art. No. 030902(R).
- STOCKS-N. G., AND MORSE-R. P. (2003). Private communication, Email received 1 October 2003.
- SULCS-S., OPPY-G., AND GILBERT-B. C. (2000). Effective linearization by noise addition in threshold detection and implications for stochastic optics, *Journal of Physics A: Mathematical and General*, **33**, pp. 3997–4007.
- TISHBY-N., PEREIRA-F. C., AND BIALEK-W. (1999). The information bottleneck method, in B. Hajek, and R. S. Sreenivas. (eds.), *Proc. 37th Annual Allerton Conference on Communication, Control and Computing*, University of Illinois, pp. 368–377.
- TORAL-R., MIRASSO-C. R., HERNANDEZ-GARCIA-E., AND PIRO-O. (1999). Synchronization of chaotic systems by common random forcing, in D. Abbott, and L. B. Kish. (eds.), *UPON 1999: Second International Conference on Unsolved Problems of Noise and Fluctuations*, Vol. 511, American Institute of Physics, pp. 255–260.
- TSITSIKLIS-J. N. (1993). Extremal properties of likelihood-ratio quantizers, *IEEE Transactions on Communications*, **41**, pp. 550–558.
- TUTEUR-F. B., AND PRESLEY-J. A. (1981). Spectral estimation of space-time signals with a DIMUS array, *The Journal of the Acoustical Society of America*, **70**, pp. 80–89.
- URICK-R. J. (1967). *Principles of Underwater Sound for Engineers*, McGraw Hill Book Company, New York.
- VAN VLECK-J. H., AND MIDDLETON-D. (1966). The spectrum of clipped noise, *Proceedings of the IEEE*, **54**, pp. 2–19.
- VAUDELLE-F., GAZENGEL-J., RIVOIRE-G., GODIVIER-X., AND CHAPEAU-BLONDEAU-F. (1998). Stochastic resonance and noise-enhanced transmission of spatial signals in optics: The case of scattering, *Journal of the Optical Society of America*, **15**, pp. 2674–2680.
- VILAR-J. M., AND RUBÍ-J. M. (1996). Divergent signal-to-noise ratio and stochastic resonance in monostable systems, *Physical Review Letters*, **77**, pp. 2863–2866.
- VON BAEYER-H. C. (2003). *Information: The New Language of Science*, Orion Book Ltd, London.
- VON NEUMANN-J., AND MORGENSTERN-O. (1944). *Theory of Games and Economic Behaviour*, Princeton University Press.
- WALLACE-R., WALLACE-D., AND ANDREWS-H. (1997). AIDS, tuberculosis, violent crime, and low birthweight in eight US metropolitan areas: Public policy, stochastic resonance, and the regional diffusion of inner-city markers, *Environment and Planning A*, **29**, pp. 525–555.
- WANG-H. S. C. (1972). Quantizer functions and their use in the analyses of digital beamformer performance, *The Journal of the Acoustical Society of America*, **53**, pp. 929–945.
- WANG-H. S. C. (1976). Postfiltering in digital beamformers, *IEEE Transactions on Aerospace and Electronic Systems*, **AES-12**, pp. 718–727.

- WANG-Y., AND WU-L. (2005). Stochastic resonance and noise-enhanced Fisher information, *Fluctuation and Noise Letters*, **5**, pp. L435–L442.
- WANNAMAKER-R. A. (1997). *The Theory of Dithered Quantization*, PhD thesis, The University of Waterloo, Canada.
- WANNAMAKER-R. A., LIPSHITZ-S. P., AND VANDERKOOY-J. (2000a). Stochastic resonance as dithering, *Physical Review E*, **61**, pp. 233–236.
- WANNAMAKER-R. A., LIPSHITZ-S. P., VANDERKOOY-J., AND WRIGHT-J. N. (2000b). A theory of non-subtractive dither, *IEEE Transactions on Signal Processing*, **48**, pp. 499–516.
- WARREN-D., AND WILLETT-P. (1999). Optimum quantization for detector fusion: Some proofs, examples, and pathology, *Journal of the Franklin Institute—Engineering and Applied Mathematics*, **336**, pp. 323–359.
- WELLENS-T., SHATOKHIN-V., AND BUCHLEITNER-A. (2004). Stochastic resonance, *Reports on Progress in Physics*, **67**, pp. 45–105.
- WENNING-G. (2004). *Aspects of Noisy Neural Information Processing*, PhD thesis, Technische Universität Berlin.
- WIDROW-B., KOLLÁR-I., AND LIU-M. (1996). Statistical theory of quantization, *IEEE Transactions on Instrumentation and Measurement*, **45**, pp. 353–361.
- WIESENFELD-K. (1993). An introduction to stochastic resonance, in J. R. Buchler., and H. E. Kandrup. (eds.), *Stochastic Processes in Astrophysics*, Vol. 706 of the Annals of the New York Academy of Sciences, The New York Academy of Sciences, pp. 13–25.
- WIESENFELD-K. (1993). Signals from noise: Stochastic resonance pays off, *Physics World*, **6**, pp. 23–24.
- WIESENFELD-K., AND JARAMILLO-F. (1998). Minireview of stochastic resonance, *Chaos*, **8**, pp. 539–548.
- WIESENFELD-K., AND MOSS-F. (1995). Stochastic resonance and the benefits of noise: From ice ages to crayfish and SQUIDS, *Nature*, **373**, pp. 33–36.
- WIESENFELD-K., PIERSON-D., PANTAZELOU-E., DAMES-C., AND MOSS-F. (1994). Stochastic resonance on a circle, *Physical Review Letters*, **72**, pp. 2125–2129.
- WILKE-S. D., AND EURICH-C. W. (2001). Representational accuracy of stochastic neural populations, *Neural Computation*, **14**, pp. 155–189.
- WOLFF-S. S., THOMAS-J. B., AND WILLIAMS-T. R. (1962). The polarity-coincidence correlator: a non-parametric detection device, *IRE Transactions on Information Theory*, **8**, pp. 5–9.
- XIONG-Z., LIVERIS-A. D., AND CHENG-S. (2004). Distributed source coding for sensor networks, *IEEE Signal Processing Magazine*, **21**, pp. 80–94.
- YATES-R. D., AND GOODMAN-D. J. (2005). *Probability and Stochastic Processes: A Friendly Introduction for Electrical and Computer Engineers*, second edn, John Wiley and Sons, Inc.
- YU-X., AND LEWIS-E. R. (1989). Studies with spike initiators: Linearization by noise allows continuous signal modulation in neural networks, *IEEE Transactions on Biomedical Engineering*, **36**, pp. 36–43.

Bibliography

- ZADOR-P. L. (1982). Asymptotic quantization error of continuous signals and the quantization dimension, *IEEE Transactions on Information Theory*, **IT-28**, pp. 139–149.
- ZAMIR-R., AND FEDER-M. (1992). On universal quantization by randomized uniform/lattice quantizers, *IEEE Transactions on Information Theory*, **38**, pp. 428–436.
- ZAMIR-R., AND FEDER-M. (1995). Rate-distortion performance in coding bandlimited sources by sampling and dithered quantization, *IEEE Transactions on Information Theory*, **41**, pp. 141–154.
- ZOZOR-S., AND AMBLARD-P. (2002). On the use of stochastic resonance in sine detection, *Signal Processing*, **82**, pp. 353–367.

List of Acronyms

Page numbers indicate the first use of the acronym in this thesis.

ADC	Analog-to-Digital Conversion, 35
AESR	Array Enhanced Stochastic Resonance, 21
AIB	Average Information Bound, 242
ASR	Aperiodic Stochastic Resonance, 17
BCPD	Backward Conditional Probability Distribution, 218
BCV	Backwards Conditional Variance, 393
BFGS	Broyden-Fletcher-Goldfarb-Shanno, 285
CB	Chapeau-Blondeau, 58
CDF	Cumulative Distribution Function, 59
CSD	Cross-power Spectral Density, 64
dB	decibels, 191
DFT	Discrete Fourier Transform, 73
DIMUS	DIgital MULtibeam Steering, 31
DPI	Data Processing Inequality, 53
FFT	Fast Fourier Transform, 73
FWHM	Full Width at Half Maximum, 100
ICDF	Inverse Cumulative Distribution Function, 92
iid	independent and identically distributed, 89
LCD	Level Crossing Detector, 26
LHS	Left Hand Side, 95
MAP	Maximum A Posteriori, 218
MASE	Mean Asymptotic Square Error, 266
MDL	Minimum Description Length, 185
MGF	Moment Generating Function, 280
ML	Maximum Likelihood, 220
MMSE	Minimum Mean Square Error, 222
MSE	Mean Square Error, 188
NAS	No Analytical Solution, 103
PDF	Probability Density Function, 27
PSD	Power Spectral Density, 11
RHS	Right Hand Side, 94
rms	root mean square, 27
SNR	Signal-to-Noise Ratio, 2
SQNR	Signal-to-Quantisation-Noise Ratio, 39
SR	Stochastic Resonance, 2
SSR	Suprathreshold Stochastic Resonance, 3

Index

- analog-to-digital conversion (ADC), 9, 10, 13, 26, 27, 35, 37, 40, 87
- Arimoto-Blahut algorithm, 131, 132, 139, 141, 345
- Cauchy-Schwarz inequality, 396
- channel capacity, 52–54, 56, 58, 60, 88, 91, 121, 126, 127, 131, 133, 134, 172, 324, 343, 345
 - binary symmetric channel, 59, 61, 126, 132
 - Chapeau-Blondeau erasure channel, 59, 60
 - Shannon-Hartley formula, 52, 54, 55, 126, 323
 - SSR, 81, 82, 126–132, 135, 139, 140, 143, 144, 146–148, 156, 176–185, 358
 - stochastic resonance, 46, 53, 57, 60
- cochlear implants, 34, 84, 85, 361
- cochlear nerve, xv, 20
- coherence function, 65–67, 69–71, 75, 77–79, 348
- correlation coefficient, 17, 24, 46, 65, 66, 188, 192, 197–202, 204–210, 213, 214, 223, 225, 227–230, 232, 241, 249, 251, 252, 257–260, 268–271, 296, 332, 347, 352, 359, 386, 387
- Cramer-Rao bound, 41, 234–236, 238, 252, 253, 264, 265, 396, 397
- cross-power spectral density, 64–68, 71, 75, 76
- DIMUS sonar arrays, 31, 87, 203, 247, 250, 253, 361
- distortion, 7, 29, 33, 38–40, 42, 185, 190, 323, 325
 - harmonic, 47, 48, 50
 - intermodulation, 47
- dithering, 13, 26–28, 40, 41, 251, 275, 276
- Fisher information, 41, 46, 86, 147, 174, 176, 177, 183, 185, 192, 232, 234–239, 242, 264, 266, 270, 313, 314, 320, 360, 381, 394
- Linear response theory, 15, 26, 50, 51, 64, 69, 251
- Lloyd method I algorithm, 232, 298–302
- Lloyd-Max algorithm, 40, 298
- lossy compression or source coding, xv, 5, 27, 39, 89, 91, 145, 185, 187, 191, 256, 323, 324, 328, 329, 342
- mean square error distortion, 6, 39–41, 46, 187–190, 193–200, 202, 228, 232, 236, 238–242, 244, 248, 250–253, 255, 256, 263–266, 268–271, 273, 274, 289, 296, 303–305, 313, 315, 319, 320, 324, 326, 327, 329, 343, 344, 346, 349, 359, 360, 391, 392
- MMSE distortion, 222, 223, 225, 227, 228, 241–243, 253, 257, 260, 262, 263, 268, 270, 271, 297, 298, 300, 301, 315–318, 345, 348–353, 393
- mutual information, 14, 19, 29, 35, 38, 53–55, 57–60, 62, 63, 65, 89, 126, 127, 131, 135, 138, 144, 348, 364, 376
 - energy constraints, 324, 327, 330, 331, 333, 334, 337, 340, 341, 354
 - neural coding, 329
 - optimal thresholds, 273, 274, 277, 278, 286, 287, 297, 298, 303, 305–307, 309–311, 317, 320, 352, 360
 - rate-distortion, 342–344, 348, 349
 - SSR, 6, 81, 84, 86, 93, 100, 108, 110, 112, 127–129, 133, 135, 136, 139, 140, 143, 144, 146–148, 150, 155, 159, 172, 179, 181–183, 185, 188, 189, 211, 218, 244, 249, 256, 274, 279, 324–326, 348, 350, 352, 358
 - stochastic resonance, 18, 20, 28, 46, 52, 82, 83
- neurons, 19, 24, 34, 64, 65, 69, 78, 152, 329
 - as an excitable system, 16, 17

- coding, 18, 19, 48, 79, 132, 152, 324, 328–330, 342
- energy efficiency, 323, 324, 329, 330
- FitzHugh-Nagumo model, 320, 354
- leaky integrate-and-fire model, 83
- McCulloch-Pitts model, 3
- models, 147, 268
- noise, 16, 34, 329, 342
- populations, 17, 21, 24, 72, 266, 330, 341, 342
- sensory, 3, 16, 19, 324
- spike trains, 72, 76, 78, 340
- SSR in, 84–86
 - FitzHugh-Nagumo model, 84
 - Hodgkin-Huxley model, 85
 - leaky integrate-and-fire model, 85
- stochastic resonance in, 2, 3, 9, 12, 14, 16, 21, 34, 84
 - FitzHugh-Nagumo model, 17, 21
 - Hodgkin-Huxley model, 17
 - models, 11, 12, 16, 34, 72, 73, 75, 77, 79
 - physiological experiments, 11, 12, 16, 17, 34, 72
- periodogram, 73, 74
- point density function, 306, 307
- power spectral density, 11, 47, 48, 64–71, 73–76
- probability distributions
 - binomial, 92, 93, 131, 137, 159–162, 167, 183, 188, 219, 238, 282, 283, 327, 384, 395
 - Cauchy, 99, 100, 102–106, 113, 116, 117, 121, 123, 140, 372
- rate-distortion, 131, 269, 324, 328, 342–346, 348
 - optimal thresholds, 352
 - SSR, 348–350, 354, 360
- signal-to-noise ratio, 2, 4, 11, 14–20, 23, 24, 26, 29, 31, 33, 42, 45–54, 64, 65, 68–73, 75–79, 84, 86, 88, 126, 191, 192, 194, 229, 231, 244, 246–248, 251, 253, 269, 273, 275, 279, 328, 345, 349, 358
- SNR gains, 19, 46, 48–55, 60, 63, 64, 78, 358
- Stirling's formula, 149, 151, 183
- stochastic resonance
 - aperiodic (ASR), 12, 14, 17, 21, 28, 29, 31, 45, 48, 52–54, 64–66, 76, 78, 79, 82, 83, 269
 - array enhanced, 12, 21
- Wiener decoding, 199, 214, 216, 217, 223, 226, 228–230, 232, 233, 253, 257–260, 262, 267, 268, 270, 296, 347
- Wiener-Kolmogorov filtering, 65, 72, 78, 79

Biography



Mark McDonnell received the degrees of Bachelor of Science in the Faculty of Mathematical and Computer Sciences in 1997, Bachelor of Engineering in Electrical and Electronic Engineering in 1998 and a First Class Honours degree in Applied Mathematics in 2001, all from The University of Adelaide in Australia. Since first graduating, he has worked as a research assistant in electromagnetic propagation, and as a computer systems engineer. His main research interests are in the field of nonlinear signal processing, with applications in computational neuroscience, complex systems, and lossy compression. In 2002, Mark was awarded a D. R. Stranks Travelling Fellowship, and in 2003, he was awarded a Santa Fe Institute Complex Systems Summer School Fellowship, as well as the AFUW Doreen MacCarthy Bursary. In 2004 he was the recipient of an Australian Academy of Science Young Researcher's award, and a United States Air Force Research Laboratory travel grant to attend the 2004 International Symposium on Information Theory, in Chicago. During his PhD, he has acted as treasurer and web-officer for EEESAU, the University of Adelaide's student branch of the IEEE, and as the postgraduate representative on the School of Electrical and Electronic Engineering's committee.

Mark is happily married to his wife, Juliet, a high school English teacher, and they have one adored cat named Bella.

NOTE:

This figure/table/image has been removed to comply with copyright regulations. It is included in the print copy of the thesis held by the University of Adelaide Library.

"That's enough. Goodnight!"

```
% phd.m
%
% author: Cecilia
% date: 09/08/05

load THESIS_TOPIC

while (funding==true)
  data = run_experiment(THESIS_TOPIC);
  GOOD_ENOUGH = query(advisor);
  if (data > GOOD_ENOUGH)
    graduate();
    break
  else
    THESIS_TOPIC = new();
    years_in_gradschool += 1;
  end
end
end
```



www.phdcomics.com

Corrigenda

Page 35: Replace all instances of $P(X)$ with $P_X(X)$ and all instances of $P(Y)$ with $P_Y(Y)$. Replace Eqn. (2.3) with

$$D(P_X(X)||P_Y(Y)) = \int_{\eta} P_X(\eta) \log_2 \left(\frac{P_X(\eta)}{P_Y(\eta)} \right) d\eta.$$

Page 36: Replace Eqn. (2.4) with

$$I(X, Y) = \int_X \int_Y P(X, Y) \log_2 \left(\frac{P(X, Y)}{P_X(X)P_Y(Y)} \right) dXdY.$$

Page 49: In Eqn. (3.2), $P_o(\omega_0)$ is the output signal power at frequency ω_0 , $P_i(\omega_0)$ is the input signal power at frequency ω_0 , $S_{N,i}(\omega_0)$ is the power spectral density of the noise at the input at frequency ω_0 , and $S_{N,o}(\omega_0)$ is the power spectral density of the noise at the output at frequency ω_0 .

Page 51: In the first new sentence, the last phrase should read "...and the input SNR has a linearly decreasing curve," rather than "...and the input SNR has a linearly increasing curve."

Page 64: The final sentence of the second paragraph is incorrect, as linear response theory holds for both periodic and aperiodic signals. Therefore, replace this sentence with "This includes, for example, signals that are strong compared to additive noise."

Page 65: Replace Eqn. (3.15) with

$$S_{xx}(f) = \int_{-\infty}^{\infty} R_{xx}(\tau) \exp(-j2\pi\tau f) d\tau,$$

and replace Eqn. (3.16) with

$$S_{xy}(f) = \int_{-\infty}^{\infty} R_{xy}(\tau) \exp(-j2\pi\tau f) d\tau.$$

Page 299: In the caption of Fig. 8.7, delete the second "increasing" on the third line.

Page 399 At the end of the first reference, page numbers should read "pp. 423-430."

Page 400: The final reference on this page should begin "BERGER-A. (ed.) (1980).