



THE UNIVERSITY OF ADELAIDE

The Automated Synchronisation of Independently Moving Cameras

Daniel William Pooley

School of Computer Science
The University of Adelaide

May 2008

TABLE OF CONTENTS

Abstract	v
Declaration	vii
Acknowledgements	ix
List of Figures	xi
List of Tables	xiii
List of Algorithms	xv
List of Examples	xvii
List of Theorems	xix
Notation	xxi
Chapter 1: Introduction	1
1.1 Synchronisation	1
1.2 Outline	5
Chapter 2: Previous Work	7
2.1 Spatial Constraints	7
2.1.1 Essential and Fundamental Matrices	9
2.1.2 Homography Matrices	13
2.1.3 Robust Estimation	16
2.2 Existing Synchronisation Methods	20
2.3 Problem Formulation	27
Chapter 3: Synchronising a Pair of Moving Cameras	29
3.1 The Evolution of a Cost Function	29
3.2 Handling Indices Between Integer Values by Interpolation	33
3.3 Estimating the Frame Offset via an Exhaustive Search	37
3.4 Estimating Both Synchronisation Parameters via an Exhaustive Search .	39
3.5 Synthetic Tests	44

3.5.1	Test Configurations	44
3.5.2	Initial Estimates of Synchronisation Parameters	46
3.5.3	Video Synchronisation Error	47
3.6	Real Video Sequence Pair Tests	51
3.7	Conclusions	54
Chapter 4:	Synchrony Pairs	57
4.1	Finding Synchrony Pairs by Epipolar Line Interpolation	57
4.2	Searching Across a Subset of Frames and Moving Points	59
4.3	Histogram Methods	62
4.3.1	Known Frame Rates	63
4.3.2	Unknown Frame Rates	65
4.3.3	Histogram Method Complexity	67
4.4	Synthetic Tests	69
4.5	Real Video Sequence Pair Tests	70
4.6	Conclusions	74
Chapter 5:	Robust Synchronisation	75
5.1	Efficient Robust Estimation with Random Sampling	76
5.2	Efficient Robust Estimation with Exhaustive Sampling	85
5.3	Robust Synchronisation Details	87
5.3.1	Inlier Classification Function \mathcal{C}	89
5.3.2	Robust Estimation Function \mathcal{R}	92
5.3.3	Estimating Synchronisation From a Single Match	92
5.3.4	Speed Parameter Functions $\mathcal{P}(\psi, \mu)$ and $\mathcal{T}(\psi, \mu)$, and Probability Distributions	96
5.4	Synthetic Tests	103
5.5	Real Video Sequence Pair Tests	109
5.6	Conclusions	113
Chapter 6:	Estimating A Projective Transformation With Moving Scene Points	117
6.1	Problem Formulation	117
6.2	A Constraint For Moving Scene Point Correspondences	118
6.2.1	Plücker Rays	119
6.2.2	The Ray Projection Matrix	121
6.2.3	Ray Transformation	122
6.2.4	The Moving Point Constraint	123
6.3	Estimation of the Ray Transformation Matrix	123
6.3.1	Applying The Three Stationary Scene Point Constraints	128

6.3.2	Applying The Moving Scene Point Constraints	129
6.3.3	Recovering The Projective Transformation \mathbf{G}	134
6.3.4	A Degenerate Camera Pair Motion	135
6.4	Robustly Estimating a Projective Transformation	137
6.4.1	A Robust Algorithm Template	138
6.4.2	Estimation and Evaluation with the Data Subset \mathbf{M}_C	140
6.5	Real Video Sequence Pair Tests	145
6.6	Conclusions	150
Chapter 7:	Conclusion	151
7.1	Summary of Contributions	151
7.2	Future Work	152
Appendix A:	Plücker Ray Theorems	155
Bibliography		165

Abstract

Computer vision is concerned with the recovery of useful scene or camera information from a set of images. One classical problem is the estimation of the 3D scene structure depicted in multiple photographs. Such estimation fundamentally requires determining how the cameras are related in space. For a dynamic event recorded by multiple video cameras, finding the temporal relationship between cameras has a similar importance. Estimating such synchrony is key to a further analysis of the dynamic scene components. Existing approaches to synchronisation involve using visual cues common to both videos, and consider a discrete uniform range of synchronisation hypotheses. These prior methods exploit known constraints which hold in the presence of synchrony, from which both a temporal relationship, and an unchanging spatial relationship between the cameras can be recovered.

This thesis presents methods that synchronise a pair of independently moving cameras. The spatial configuration of cameras is assumed to be known, and a cost function is developed to measure the quality of synchrony even for accuracies within a fraction of a frame. A Histogram method is developed which changes the approach from a consideration of multiple synchronisation hypotheses, to searching for seemingly synchronous frame pairs independently. Such a strategy has increased efficiency in the case of unknown frame rates. Further savings can be achieved by reducing the sampling rate of the search, by only testing for synchrony across a small subset of frames. Two robust algorithms are devised, using Bayesian inference to adaptively seek the sampling rate that minimises total execution time. These algorithms have a general underlying premise, and should be applicable to a wider class of robust estimation problems. A method is also devised to robustly synchronise two moving cameras when their spatial relationship is unknown. It is assumed that the motion of each camera has been estimated independently, so that these motion estimates are unregistered. The algorithm recovers both a synchronisation estimate, and a 3D transformation that spatially registers the two cameras.

Declaration

This work contains no material which 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 for loan and photocopying, subject to the provisions of the Copyright Act 1968.

The author acknowledges that copyright of published works contained within this thesis (as listed below) resides with the copyright holder/s of those works.

Daniel Pooley
October 9, 2008

D.W. Pooley, M.J. Brooks, A.J. van den Hengel and W. Chojnacki. A voting scheme for estimating the synchrony of moving-camera videos. In *International Conference on Image Processing, Barcelona, Spain, September 15-19, 2003*, volume 1, pages 413–416. IEEE Computer Society Press.

D. Pooley, M. Brooks and A. van den Hengel. RATSAC: an adaptive method for accelerated robust estimation, and its application to video synchronisation. In *Digital Image Computing: Techniques and Applications, Adelaide, Australia, December 3-5, 2007*. IEEE Computer Society Press.

Acknowledgements

I would chiefly like to thank my supervisors Professor Michael Brooks and Associate Professor Anton van den Hengel, for their support, guidance and motivation throughout my candidature. In addition, thanks go to Professor Wojciech Chojnacki for his assistance with a variety of mathematical queries, as well as to Dr. Darren Gawley, Dr. Chunhua Shen, Dr. Thorsten Thormählen, Dr. John Bastian, and Mr. Rhys Hill for productive theoretical discussions. I am also appreciative of the support provided by the Cooperative Research Centre for Sensor Signal and Information Processing. Finally, special thanks go to my family and friends, in recognition of their encouragement, moral support, and for providing welcome distractions when most needed.

LIST OF FIGURES

1.1	Time-stamping 3 video sequences	2
1.2	Side-by-side display of synchronised video frames, cropped to a region of interest	4
2.1	The perspective projection of a scene point	8
2.2	Visualisation of the epipolar constraint	11
2.3	Visualisation of the epipolar constraint for projections of a moving scene point in synchronous and asynchronous frame pairs	14
2.4	An application of the homography matrix; original images and the resulting mosaic	15
2.5	An example of pixel-based synchrony with a synthetic scene of a bouncing ball	22
2.6	Key frames from a video sequence of a moving grey square (right), and resulting space-time interest points (left)	25
3.1	Corresponding trajectories between a pair of video sequences, identified prior to synchronisation	30
3.2	Two cases of epipolar errors for synchronous and asynchronous frame pairs	31
3.3	Epipolar Line Interpolation	36
3.4	A Comparison of Interpolated Epipolar Errors	37
3.5	Ranges of frames in each video sequence captured while both cameras recorded	38
3.6	Frame ranges for different values of offset a ($a_{end} > 0$)	40
3.7	Line parameterisation (α, β)	42
3.8	Spatial configuration for synthetic synchronisation tests	45
3.9	Proportions of tests in which each cost function achieved the best synchronisation	50
3.10	VSE results for synthetic test cases synchronised by a search of the parameter space, using scene points with linear motion	52
3.11	VSE results for synthetic test cases synchronised by a search of the parameter space, using scene points with piecewise-linear motion	53
3.12	Approximately synchronous frames from the Lego robot sequences, synchronised using a uniform search for the frame offset	55

4.1	Example synchrony pairs using epipolar line interpolation	59
4.2	Example synchrony pair sets (left) and resulting 2D histograms (right) .	68
4.3	VSE results for synthetic test cases synchronised by a search for synchrony pairs, with a known frame rate ratio and scene points with linear motion	71
4.4	VSE results for synthetic test cases synchronised by a search for synchrony pairs, with an unknown frame rate ratio and scene points with linear motion	72
4.5	Synchrony pair set (top), and resulting 2D (bottom right) and section of 1D (bottom left) histograms for the Lego robot example	73
5.1	Example synchrony pair sets from incorrectly matched trajectories	77
5.2	Effects of a misestimation of σ on a threshold test for Chi-squared distributions of different order	90
5.3	Line parameterisation (α', β')	93
5.4	Relative positions of the moving cameras' optical centres for the bouncing balls example	109
5.5	Two examples of the gradient measure used to assist in ball detection . .	111
5.6	Side-by-side display of synchronised video frames, cropped to a region around a match identified by RATSAC assuming a known frame rate ratio	114
5.7	A reconstruction of 3D ball trajectories, projected to virtual cameras, for matches identified by CATSAC assuming an unknown frame rate ratio .	115
6.1	Two top-down views of degenerate camera pair motions for spatial registration using moving point correspondences	136
6.2	A trajectory pair visible in frames $i_{min} \dots i_{max}$ in video 1, and $j_{min} \dots j_{max}$ in video 2, which can not be a match according to frame offset a	139
6.3	A comparison of optical centre locations for two potential alignments of the second camera in the bouncing balls example	149
6.4	Baselines of approximately synchronous frame pairs for the bouncing balls example	149

LIST OF TABLES

3.1	Synthetic initial estimate performance for a single moving scene point . . .	47
3.2	VSE measures for a single moving scene point	49
3.3	Percentage of successful synchronisations for a single moving scene point	49
3.4	VSE measures for the synchronisation of the Lego robot sequences via a search of the parameter space	54
4.1	Synthetic initial estimate performance for synchrony pair searches using a single moving scene point	69
4.2	VSE measures for the synchronisation of the Lego robot sequences via a search for synchrony pairs	74
5.1	Correct and incorrect matches classified for synthetic test cases, by using RATSAC followed by a minimisation of the robust cost function \mathcal{R} . . .	104
5.2	Correct and incorrect matches classified for synthetic test cases, by using CATSAC followed by a minimisation of the robust cost function \mathcal{R} . . .	105
5.3	Median VSE measures for synthetic test cases, for both CATSAC and RATSAC followed by a minimisation of the cost function \mathcal{R}	107
5.4	Percentages of synthetic tests for which CATSAC and RATSAC, followed by a minimisation of \mathcal{R} , achieve synchronisation with a VSE measure less than 0.5	107
5.5	Initial ψ choices, average times (seconds), and percentages of time saved by using an adaptive speedup, for synthetic test cases with linear scene point motions	108
5.6	VSE measures and matches identified for the bouncing balls example, for both RATSAC and CATSAC, followed by a minimisation of \mathcal{R}	112
5.7	Initial ψ choices, average times (seconds) and percentages of time saved by using an adaptive speedup, for the bouncing balls example	113
6.1	Matches identified for the bouncing balls example, for a robust estimation of both a and \mathbf{G} , followed by a minimisation of \mathcal{R}	147
6.2	Synchronisation quality for the bouncing balls example, for a robust estimation of both a and \mathbf{G} , followed by a minimisation of \mathcal{R}	147
6.3	Spatial alignment quality for the bouncing balls example, for a robust estimation of both a and \mathbf{G} , followed by a minimisation of \mathcal{R}	148

LIST OF ALGORITHMS

3.1	Estimation of a with an Exhaustive Search	41
3.2	Estimation of (α, β) with an Exhaustive Search	45
4.1	An Exhaustive Search for Synchrony Pairs	60
5.1	Random Adaptive Trade-off Sample Consensus	84
5.2	Comprehensive Adaptive Trade-off Sample Consensus	88
5.3	Accelerated Iteration for Estimating Synchronisation	97
6.1	Robust Estimation of a and \mathbf{G}	146

LIST OF EXAMPLES

3.1	Problematic Sum of Squares Example	32
3.2	Problematic Residual Interpolation Example	34
4.1	Non-Synchronous Entries in the Synchrony Pair Set	61
6.1	A Case of Insufficient Stationary Correspondences	119
6.2	Estimating $\hat{\mathbf{H}}$ for Stationary Cameras	124

LIST OF THEOREMS

A.1	The null-space of a Plücker Matrix	155
A.2	Uniqueness up to Scale	155
A.3	Duality Constraint	156
A.4	Ray Intersection	157
A.5	Ray Constraint	158
A.6	Inner Product	159
A.7	Projection to a Line Matrix	160
A.8	Ray Projection	161
A.9	Ray Back-Projection	162
A.10	Ray Transformation	163

NOTATION

k	A scalar (non-bold lower case).
K	A random variable (non-bold upper case).
X	An event (non-bold sans-serif upper case).
\mathbf{S}	A set (bold upper case).
\mathbf{v}	A vector (bold lower case).
\mathbf{A}	A matrix (bold sans-serif upper case).
\mathbf{A}^\top	The transpose of matrix \mathbf{A} .
\mathbf{A}^{-1}	The inverse of matrix \mathbf{A} .
$[\mathbf{v}]_i$	The element i of vector \mathbf{v} .
$[\mathbf{A}]_{i,j}$	The element (i, j) of matrix \mathbf{A} .
$[\mathbf{A}]_{i,*}$	The row i of matrix \mathbf{A} .
$[\mathbf{A}]_{*,j}$	The column j of matrix \mathbf{A} .
$Tr(\mathbf{A})$	The trace of matrix \mathbf{A} .
$Rank(\mathbf{A})$	The rank of matrix \mathbf{A} .
$\ \mathbf{A}\ _F$	The Frobenius norm of matrix \mathbf{A} .
$P(X)$	The probability of event X .
\sim	Equality up to scale of two vectors or matrices.
\approx	Approximate equality up to scale of two vectors or matrices.