Roberto Cipolla, Sebastiano Battiato, and Giovanni Maria Farinella (Eds.)

Computer Vision

Studies in Computational Intelligence, Volume 285

Editor-in-Chief
Prof. Janusz Kacprzyk
Systems Research Institute
Polish Academy of Sciences
ul. Newelska 6
01-447 Warsaw
Poland
E-mail: kacprzyk@ibspan.waw.pl

Further volumes of this series can be found on our homepage: springer.com

Vol. 265. Zbigniew W. Ras and Li-Shiang Tsay (Eds.)
Advances in Intelligent Information Systems, 2009
ISBN 978-3-642-05182-1

Vol. 266. Akitoshi Hanazawa, Tsutom Miki, and Keiichi Horio (Eds.)
Brain-Inspired Information Technology, 2009
ISBN 978-3-642-04024-5

Vol. 267. Ivan Zelinka, Sergej Celikovský, Hendrik Richter, and Guanrong Chen (Eds.)
Evolutionary Algorithms and Chaotic Systems, 2009
ISBN 978-3-642-10706-1

Vol. 268. Johann M.Ph. Schumann and Yan Liu (Eds.)
Applications of Neural Networks in High Assurance Systems, 2009
ISBN 978-3-642-10689-7

Vol. 269. Francisco Fernández de de Vega and Erick Cantú-Paz (Eds.)
Parallel and Distributed Computational Intelligence, 2009
ISBN 978-3-642-10674-3

Vol. 270. Zong Woo Geem
Recent Advances In Harmony Search Algorithm, 2009
ISBN 978-3-642-04316-1

Vol. 271. Janusz Kacprzyk, Frederick E. Petry, and Adnan Yazici (Eds.)
Uncertainty Approaches for Spatial Data Modeling and Processing, 2009
ISBN 978-3-642-10662-0

Vol. 272. Carlos A. Coello Coello, Clarisse Dhaenens, and Laetitia Jourdan (Eds.)
Advances in Multi-Objective Nature Inspired Computing, 2009
ISBN 978-3-642-11217-1

Vol. 273. Fatos Xhafa, Santi Caballé, Ajith Abraham, Thanasis Daradoumis, and Angel Alejandro Juan Perez (Eds.)
Computational Intelligence for Technology Enhanced Learning, 2010
ISBN 978-3-642-11223-2

Vol. 274. Zbigniew W. Raś and Alicja Wieczorkowska (Eds.)
Advances in Music Information Retrieval, 2010
ISBN 978-3-642-11673-5

Vol. 275. Dilip Kumar Pratihar and Lakhmi C. Jain (Eds.)
Intelligent Autonomous Systems, 2010
ISBN 978-3-642-11675-9

Vol. 276. Jacek Mańdziuk
Knowledge-Free and Learning-Based Methods in Intelligent Game Playing, 2010
ISBN 978-3-642-11677-3

Vol. 277. Filippo Spagnolo and Benedetto Di Paola (Eds.)
European and Chinese Cognitive Styles and their Impact on Teaching Mathematics, 2010
ISBN 978-3-642-11679-7

Vol. 278. Radomir S. Stankovic and Jaakko Astola
From Boolean Logic to Switching Circuits and Automata, 2010
ISBN 978-3-642-11681-0

Vol. 279. Manolis Wallace, Ioannis E. Anagnostopoulos, Phivos Mylonas, and Maria Bielikova (Eds.)
Semantics in Adaptive and Personalized Services, 2010
ISBN 978-3-642-11683-4

Vol. 280. Chang Wen Chen, Zhu Li, and Shiguo Lian (Eds.)
Intelligent Multimedia Communication: Techniques and Applications, 2010
ISBN 978-3-642-11685-8

Vol. 281. Robert Babuska and Frans C.A. Groen (Eds.)
Interactive Collaborative Information Systems, 2010
ISBN 978-3-642-11687-2

Vol. 282. Husrev Taha Sencar, Sergio Velastin, Nikolaos Nikolaidis, and Shiguo Lian (Eds.)
Intelligent Multimedia Analysis for Security Applications, 2010
ISBN 978-3-642-11754-1

Vol. 283. Ngoc Thanh Nguyen, Radoslaw Katarzyniak, and Shyi-Ming Chen (Eds.)
Advances in Intelligent Information and Database Systems, 2010
ISBN 978-3-642-12089-3

Vol. 284. Juan R. González, David Alejandro Pelta, Carlos Cruz, Germán Terrazas, and Natalio Krasnogor (Eds.)
Nature Inspired Cooperative Strategies for Optimization (NICSO 2010), 2010
ISBN 978-3-642-12537-9

Vol. 285. Roberto Cipolla, Sebastiano Battiato, and Giovanni Maria Farinella (Eds.)
Computer Vision, 2010
ISBN 978-3-642-12847-9

Roberto Cipolla, Sebastiano Battiato,
and Giovanni Maria Farinella (Eds.)

Computer Vision

Detection, Recognition and Reconstruction

Springer

Prof. Roberto Cipolla
Department of Engineering
University of Cambridge
Cambridge, CB2 1PZ
UK
E-mail: cipolla@eng.cam.ac.uk

Dr. Giovanni Maria Farinella
Dipartimento di Matematica ed Informatica
University of Catania
Viale A. Doria 6, I
95125 - Catania
Italy
E-mail: gfarinella@dmi.unict.it

Prof. Sebastiano Battiato
Dipartimento di Matematica ed Informatica
University of Catania
Viale A. Doria 6, I
95125 – Catania
Italy
E-mail: battiato@dmi.unict.it

ISBN 978-3-642-12847-9 e-ISBN 978-3-642-12848-6

DOI 10.1007/978-3-642-12848-6

Studies in Computational Intelligence ISSN 1860-949X

Library of Congress Control Number: 2010925724

© 2010 Springer-Verlag Berlin Heidelberg

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Typeset & *Cover Design:* Scientific Publishing Services Pvt. Ltd., Chennai, India.

Printed on acid-free paper

9 8 7 6 5 4 3 2 1

springer.com

Editors' Biographies

Roberto Cipolla

Roberto Cipolla obtained the B.A. degree (Engineering) from the University of Cambridge in 1984 and an M.S.E. (Electrical Engineering) from the University of Pennsylvania in 1985. From 1985 to 1988 he studied and worked in Japan at the Osaka University of Foreign Studies (Japanese Language) and Electrotechnical Laboratory. In 1991 he was awarded a D.Phil. (Computer Vision) from the University of Oxford and from 1991-92 was a Toshiba Fellow and engineer at the Toshiba Corporation Research and Development Centre in Kawasaki, Japan. He joined the Department of Engineering, University of Cambridge in 1992 as a Lecturer and a Fellow of Jesus College. He became a Reader in Information Engineering in 1997 and a Professor in 2000. He is Professor of Computer Vision at the Royal Academy of Arts, London (since 2004) and a Director of Toshiba's Cambridge Research Laboratory (since 2007). His research interests are in computer vision and robotics and include the recovery of motion and 3D shape of visible surfaces from image sequences; object detection and recognition; novel man-machine interfaces using hand, face and body gestures; real-time visual tracking for localisation and robot guidance; applications of computer vision in mobile phones, visual inspection and image-retrieval and video search. He has authored 2 books, edited 7 volumes and co-authored more than 300 papers. Professor Cipolla founded (in 2006) and currently directs the International Computer Vision Summer School.

Sebastiano Battiato

Sebastiano Battiato was born in Catania, Italy, in 1972. He received the degree in Computer Science (summa cum laude) in 1995 and his Ph.D. in Computer Science and Applied Mathematics in 1999. From 1999 to 2003 he has leaded the "Imaging" team c/o STMicroelectronics in Catania. Since

2004 he has been a researcher at Department of Mathematics and Computer Science of the University of Catania. His research interests include image enhancement and processing, image coding and camera imaging technology. He published more than 90 papers in international journals, conference proceedings and book chapters. He is a co-inventor of about 15 international patents. He is a reviewer for several international journals and he has been regularly a member of numerous international conference committees. He has participated in many international and national research projects. He is an Associate Editor of the SPIE Journal of Electronic Imaging (Specialty: digital photography and image compression). He is a director (and co-founder) of the International Computer Vision Summer School. He is a Senior Member of the IEEE.

Giovanni Maria Farinella

Giovanni Maria Farinella obtained the degree in Computer Science (egregia cum laude) from University of Catania in 2004. He was awarded a Doctor of Philosophy degree (Computer Vision) in 2008. He joined the Department of Mathematics and Computer Science, University of Catania in 2008 as Contract Researcher. He is Contract Professor of Computer Science, Computer Graphics and Computer Vision at the Academy of Arts, Catania (since 2004) and Associate Member of the Computer Vision and Robotics Research Group at University of Cambridge (since 2006). His research interests lie in the fields of Computer Vision, Computer Graphics, Pattern Recognition and Machine Learning. He has co-authored a number of papers in international journals and international conferences proceedings. He also serves as reviewer international journals and conferences. He founded (in 2006) and currently directs the International Computer Vision Summer School.

Preface

Computer vision is the science and technology of making machines that see. It is concerned with the theory, design and implementation of algorithms that can automatically process visual data to recognize objects, track and recover their shape and spatial layout.

The International Computer Vision Summer School - ICVSS was established in 2007 to provide both an objective and clear overview and an in-depth analysis of the state-of-the-art research in Computer Vision. The courses are delivered by world renowned experts in the field, from both academia and industry, and cover both theoretical and practical aspects of real Computer Vision problems. The school is organized every year by University of Cambridge (Computer Vision and Robotics Group) and University of Catania (Image Processing Lab). Different topics are covered each year. A summary of the past Computer Vision Summer Schools can be found at: http://www.dmi.unict.it/icvss

This edited volume contains a selection of articles covering some of the talks and tutorials held during the first two editions of the school on topics such as Recognition, Registration and Reconstruction. The chapters provide an in-depth overview of these challenging areas with key references to the existing literature.

The book starts with two chapters devoted to introducing the reader to the exciting field of Vision. In Chapter 1 a discussion about the fundamentals of the discipline is presented. Human vision is analyzed and a number of basic principles of biological vision that might be of interest to the machine vision community are identified. Chapter 2 introduces a methodology to evaluate the effectiveness of local features when employed for recognition tasks. A novel mathematical characterisation of the co-variance properties of the features which accounts for deviation from the usual idealised image affine (de)formation model together with a novel metrics to evaluate the features are described.

In Chapter 3 and Chapter 4 computational techniques based on Dynamic Graph Cuts and Discriminative Graphical Models are presented and

employed in problems such as image and video segmentation, pose estimation and context based classification. An overview of the Mutual SubSpace Method and its applications in face and character recognition is presented in Chapter 5. The book continues with three chapters that cover recent approaches for detection, classification and recognition of objects, scenes and activities from images. Specifically, Chapter 6 concentrates on the task of activity recognition by using graphical models which combine information from both object recognition and scene classification. Chapter 7 examines Semantic Texton Forests and evaluates their use for image categorization and semantic segmentation, whereas Chapter 8 focuses on finding a suitable representation that can efficiently capture the intrinsic three-dimensional and multi-view nature of object categories to help the recognition and categorization task.

In Chapter 9 a vision-based system for touch-free interaction with a display at a distance is presented after a deep revision of the state of the art techniques on hand tracking. The problem of tracking multiple objects taking into account multiple views is introduced in Chapter 10.

Finally, two Chapters discussing the problem and existing solutions for 3D reconstruction through multiview and photometric stereo conclude the book.

It is our hope that graduate students, young and senior researchers, and academic/industrial professionals will find the book useful for reviewing current approaches and for teaching Computer Vision, thereby continuing the mission of the International Computer Vision Summer School.

Sicily, December 2009

Roberto Cipolla
Sebastiano Battiato
Giovanni Maria Farinella

Acknowledgements

We would like to take this opportunity to thank all contributors of this book, and all people involved in the organization of ICVSS. A special thanks to Professor Giovanni Gallo for his invaluable encouragement, guidance and support.

Contents

1 Is Human Vision Any Good? 1
Jan J. Koenderink
1.1 Introduction 1
1.1.1 What Is "Human Vision"? 1
1.1.2 How Is Human Vision Implemented? 2
1.2 Frameworks 5
1.2.1 The Spatial Framework 6
1.2.2 The Photometric Framework 9
1.3 A Case Study: "Shape from Shading" 13
1.3.1 The So Called "Shape from Shading Problem" 13
1.3.2 Setting up the Problem 15
1.3.3 The Local Shape from Shading Problem 18
1.4 Final Remarks 24
References 25

2 Knowing a Good Feature When You See It: Ground Truth and Methodology to Evaluate Local Features for Recognition 27
Andrea Vedaldi, Haibin Ling, Stefano Soatto
2.1 Introduction 27
2.2 Empirical Studies of Local Features 29
2.2.1 Some Nomenclature 30
2.3 Constructing a Rigorous Ground Truth 30
2.3.1 Modeling the Detector 30
2.3.2 Modeling Correspondences 33
2.3.3 Ground-Truth Correspondences 34
2.3.4 Comparing Ground-Truth and Real-World Correspondences 36
2.3.5 The Data 37
2.4 Learning to Compare Invariant Features 40

2.4.1 Wide Baseline Motion Statistics 41
2.4.2 Learning to Rank Matches 41
2.4.3 Learning to Accept Matches 43
2.5 Discussion 44
Appendix 1: Calculations 44
Appendix 2: Algorithms 45
References 48

3 Dynamic Graph Cuts and Their Applications in Computer Vision 51
Pushmeet Kohli, Philip H.S. Torr
3.1 Introduction 51
3.1.1 Markov and Conditional Random Fields 52
3.2 Graph Cuts for Energy Minimization 54
3.2.1 The st-Mincut Problem 54
3.2.2 Formulating the Max-Flow Problem 55
3.2.3 Augmenting Paths, Residual Graphs 56
3.2.4 Minimizing Functions Using Graph Cuts 56
3.3 Minimizing Dynamic Energy Functions Using Dynamic Graph Cuts 58
3.3.1 Dynamic Computation 58
3.3.2 Energy and Graph Reparameterization 60
3.4 Recycling Computation 61
3.4.1 Updating Residual Graphs 62
3.4.2 Computational Complexity of Update Operations 64
3.5 Improving Performance by Recycling Search Trees 65
3.5.1 Reusing Search Trees 65
3.5.2 Tree Recycling for Dynamic Graph Cuts 66
3.6 Dynamic Image Segmentation 67
3.6.1 CRFs for Image Segmentation 68
3.6.2 Image Segmentation in Videos 70
3.6.3 Experimental Results 70
3.6.4 Reusing Flow vs. Reusing Search Trees 73
3.7 Simultaneous Segmentation and Pose Estimation of Humans 74
3.7.1 Pose Specific CRF for Image Segmentation 78
3.7.2 Formulating the Pose Inference Problem 82
3.7.3 Experiments 85
3.7.4 Shape Priors for Reconstruction 87
3.7.5 Discussion 91
3.7.6 Summary and Future Work 92
3.8 Measuring Uncertainty in Graph Cut Solutions 93
3.8.1 Preliminaries 94
3.8.2 Computing Min-Marginals Using Graph Cuts 97

3.8.3 Min-Marginals and Flow Potentials . . . 99
3.8.4 Extension to Multiple Labels . . . 100
3.8.5 Minimizing Energy Function Projections Using Dynamic Graph Cuts . . . 101
3.8.6 Computational Complexity and Experimental Evaluation . . . 101
3.8.7 Applications of Min-Marginals . . . 102
References . . . 104

4 Discriminative Graphical Models for Context-Based Classification . . . 109
Sanjiv Kumar
4.1 Contextual Dependencies in Images . . . 109
4.1.1 The Nature of Contextual Interactions . . . 110
4.2 Markov Random Field (MRF) . . . 112
4.3 Conditional Random Field (CRF) . . . 113
4.3.1 Association Potential . . . 115
4.3.2 Interaction Potential . . . 117
4.4 Parameter Learning and Inference . . . 119
4.4.1 Maximum Likelihood Parameter Learning . . . 119
4.4.2 Inference . . . 121
4.5 Extensions . . . 121
4.5.1 Multiclass CRF . . . 122
4.5.2 Hierarchical CRF . . . 123
4.6 Applications . . . 127
4.6.1 Man-Made Structure Detection . . . 127
4.6.2 Image Classification and Contextual Object Detection . . . 129
4.7 Related Work and Further Readings . . . 133
References . . . 133

5 From the Subspace Methods to the Mutual Subspace Method . . . 135
Ken-ichi Maeda
5.1 Introduction . . . 135
5.2 The Subspace Methods . . . 139
5.2.1 A Brief History of the Subspace Methods . . . 139
5.2.2 Basic Idea . . . 139
5.2.3 Subspace Construction . . . 140
5.3 The Mutual Subspace Method . . . 142
5.3.1 Basic Idea . . . 142
5.3.2 Application to Chinese Character Recognition . . . 144
5.3.3 Application to Face Recognition . . . 147
5.3.4 Application to 3-D Face Recognition . . . 148
5.4 Conclusion . . . 154
References . . . 155

6 What, Where and Who? Telling the Story of an Image by Activity Classification, Scene Recognition and Object Categorization 157
Li Fei-Fei, Li-Jia Li
6.1 Introduction and Motivation 157
6.2 Overall Approach 159
6.3 Literature Review 160
6.4 The Integrative Model 160
6.4.1 Labeling an Unknown Image 163
6.5 Learning the Model 164
6.6 System Implementation 165
6.7 Experiments and Results 166
6.7.1 Dataset 166
6.7.2 Experimental Setup 166
6.7.3 Results 169
6.8 Conclusion 169
References 170

7 Semantic Texton Forests 173
Matthew Johnson, Jamie Shotton
7.1 Introduction 173
7.2 Related Work 174
7.3 Randomized Decision Forests 175
7.3.1 Training the Forest 176
7.3.2 Experiments 180
7.4 Image Categorization 183
7.4.1 Tree Histograms and Pyramid Matching 184
7.4.2 Categorization Results 186
7.5 Semantic Segmentation 188
7.6 Soft Classification of Pixels 189
7.7 Image-Level Semantic Constraints 191
7.7.1 Categorization Results 192
7.7.2 The Image Level Prior 192
7.8 Compositional Constraints 194
7.9 Experiments 194
7.9.1 MSRC21 Dataset 196
7.9.2 VOC 2007 Segmentation Dataset 200
7.10 Discussion 200
References 201

8 Multi-view Object Categorization and Pose Estimation 205
Silvio Savarese, Li Fei-Fei
8.1 Introduction 206
8.2 Literature Review 209
8.3 The Part-Based Multi-view Model 210

8.3.1 Overview . . . 210
8.3.2 Canonical Parts and Linkage Structure . . . 211
8.4 Building the Model . . . 212
8.4.1 Extract Features . . . 213
8.4.2 Form Parts . . . 213
8.4.3 Find Canonical Parts Candidates . . . 214
8.4.4 Create the Model . . . 216
8.5 View Synthesis . . . 217
8.5.1 Representing an Unseen View . . . 217
8.6 Recognizing Object Class in Unseen Views . . . 219
8.6.1 Extract Features and Get Part Candidates . . . 220
8.6.2 Recognition Procedure: First Step . . . 221
8.6.3 Recognition Procedure: Second Step . . . 222
8.7 Experiments and Results . . . 222
8.7.1 Experiment I: Comparison with Thomas et al. [63] . . . 222
8.7.2 Experiment II: Detection and Pose Estimation Results on the Dataset in [54] . . . 223
8.7.3 Experiment III: Detection and Pose Estimation Results on the Dataset in [55] . . . 227
8.8 Conclusion . . . 227
References . . . 228

9 A Vision-Based Remote Control . . . 233
Björn Stenger, Thomas Woodley, Roberto Cipolla
9.1 Introduction . . . 233
9.2 Prior Work . . . 235
9.2.1 Hand Tracking for Human Computer Interfaces . . . 235
9.2.2 Commercial Gesture Interface Systems . . . 238
9.2.3 Visual Tracking of a Single Object . . . 239
9.3 Tracking with Multiple Observers . . . 240
9.3.1 Observation Models . . . 240
9.3.2 Evaluating Single Observers . . . 243
9.3.3 Evaluating Multiple Observers . . . 245
9.3.4 Parallel Evaluation . . . 246
9.3.5 Cascaded Evaluation . . . 247
9.3.6 Dynamic Model Discussion . . . 247
9.3.7 Experimental Results . . . 248
9.3.8 Individual Observers . . . 248
9.3.9 Observer Combinations . . . 249
9.3.10 Tracker Evaluation on Selected Combinations . . . 252
9.4 Gesture Interface System . . . 253
9.4.1 Visual Attention Mechanism . . . 253
9.4.2 Tracking Mechanism . . . 254
9.4.3 Selection Mechanisms . . . 255

9.5 Summary and Conclusion 257
References 258

10 Multi-view Multi-object Detection and Tracking 263
Murtaza Taj, Andrea Cavallaro
10.1 Introduction 263
10.2 Problem Formulation 265
10.3 Calibration and Fusion 266
10.3.1 Single-Level Homography 266
10.3.2 Multi-Level Homography 268
10.4 Track-First Approaches 269
10.4.1 Independent Tracking 271
10.4.2 Collaborative Tracking 271
10.5 Fuse-First Approaches 272
10.5.1 Detection-Based Tracking 273
10.5.2 Track-Before-Detect 276
10.6 Conclusions 278
References 278

11 Shape from Photographs: A Multi-view Stereo Pipeline 281
Carlos Hernández, George Vogiatzis
11.1 Introduction 281
11.2 Multi-view Stereo Pipeline: From Photographs to 3D Models 284
11.3 Computing Photo-Consistency from a Set of Calibrated Photographs 285
11.3.1 Normalized Cross Correlation for Depth-Map Computation 287
11.3.2 Depth Map Estimation 288
11.3.3 Photo-Consistency 3D Map from a Set of Depth-Maps 292
11.4 Extracting a 3D Surface from a 3D Map of Photo-Consistency 292
11.4.1 Multi-view Stereo Using Multi-resolution Graph-Cuts 292
11.4.2 Discontinuity Cost from a Set of Depth-Maps 294
11.4.3 Graph Structure 295
11.4.4 Labeling Cost from a Set of Depth-Maps 295
11.4.5 Probabilistic Fusion of Depth Sensors 296
11.4.6 Deformable Models 299
11.5 Experiments 302
11.5.1 Depth Map Evaluation 302
11.5.2 Multi-view Stereo Evaluation 304

11.5.3 Digitizing Works of Art 304
11.6 Discussion 307
References 309

12 Practical 3D Reconstruction Based on Photometric Stereo 313
George Vogiatzis, Carlos Hernández
12.1 Introduction 313
12.2 Photometric Stereo with Coloured Light 314
12.2.1 Classic Three-Source Photometric Stereo 315
12.2.2 Multi-spectral Sources and Sensors 316
12.2.3 Calibration 317
12.2.4 Comparison with Photometric Stereo 317
12.2.5 The Problem of Shadows 318
12.2.6 Facial Capture Experiments 325
12.2.7 Related Work 328
12.3 Multi-view Photometric Stereo 330
12.3.1 Related Work 331
12.3.2 Algorithm 333
12.3.3 Experiments 338
References 342

Index 347

List of Contributors

Andrea Cavallaro
School of Electronic Engineering and Computer Science
Queen Mary University of London
Mile End Road, London E1 4NS, UK
andrea.cavallaro@elec.qmul.ac.uk

Roberto Cipolla
Department of Engineering
University of Cambridge
CB2 1PZ, Cambridge UK
cipolla@eng.cam.ac.uk

Li Fei-Fei
Department of Computer Science,
Stanford University,
353 Serra Mall, Gates Building,
CA 94305-9020, Stanford, USA
feifeili@cs.stanford.edu

Carlos Hernández
Computer Vision Group
Toshiba Research Cambridge Ltd,
CB4 0GZ, Cambridge, UK
carlos.hernandez@crl.toshiba.co.uk

Matthew Johnson
Nokia Spear Tower, 1 Market Plaza
CA 94105, San Francisco, USA
matthew.3.johnson@nokia.com

Jan J. Koenderink
EEMCS, Delft University of Technology
P.O. Box 5031, 2600GA Delft,
The Netherlands
jan.koenderink@telfort.nl

Pushmeet Kohli
Microsoft Research,
7 J.J. Thomson Ave,
CB3 0FB, Cambridge, UK
pkohli@microsoft.com

Sanjiv Kumar
Google Research
76 Ninth Ave, New York, NY, USA
sanjivk@google.com

Li-Jia Li
Department of Computer Science
Stanford University
353 Serra Mall, Gates Building
CA 94305-9020, Stanford, USA
lijiali@cs.stanford.edu

Haibin Ling
Center for Information Science & Technology
Department of Computer & Information Sciences

324 Wachman Hall,
Temple University
1805 N. Broad St.,
Philadelphia, PA 19122
hbling@temple.edu

Ken-ichi Maeda
Corporate Research &
Development Center
Toshiba Corporation
1 Komukai-Toshiba-Cho, Kawasaki,
212-8582, Japan
ken.maeda@toshiba.co.jp

Silvio Savarese
Electrical and Computer Engineering
University of Michigan
1301 Beal Avenue, Ann Arbor,
MI 48109-2122, USA
silvio@umich.edu

Jamie Shotton
Microsoft Research,
7 J.J. Thomson Ave,
CB3 0FB, Cambridge, UK
jamie@shotton.org

Stefano Soatto
Department of Computer Science
University of California
at Los Angeles
Boelter Hall 3531D, 405 Hilgard Ave
Los Angeles, CA 90095-1596, USA
soatto@ucla.edu

Björn Stenger
Computer Vision Group
Toshiba Research Cambridge Ltd,
CB4 0GZ, Cambridge, UK
bjorn@cantab.net

Murtaza Taj
School of Electronic Engineering and
Computer Science
Queen Mary University of London
Mile End Road, London E1 4NS, UK
murtaza@elec.qmul.ac.uk

Philip H.S. Torr
Department of Computing
Oxford Brookes University
Wheatley, Oxford, OX33 1HX, UK
philiptorr@brookes.ac.uk

Andrea Vedaldi
Department of Computer Science
University of California
at Los Angeles
Boelter Hall 3531D, 405 Hilgard Ave
Los Angeles, CA 90095-1596, USA
vedaldi@cs.ucla.edu

George Vogiatzis
Aston University
Birmingham, UK
g.vogiatzis@aston.ac.uk

Thomas Woodley
Department of Engineering
University of Cambridge
CB2 1PZ, Cambridge UK
tew32@cam.ac.uk