

LNCS 3952

Aleš Leonardis
Horst Bischof
Axel Pinz (Eds.)

Computer Vision – ECCV 2006

9th European Conference on Computer Vision
Graz, Austria, May 2006
Proceedings, Part II

2
Part II



 Springer

Aleš Leonardis Horst Bischof
Axel Pinz (Eds.)

Computer Vision – ECCV 2006

9th European Conference on Computer Vision
Graz, Austria, May 7-13, 2006
Proceedings, Part II

Volume Editors

Aleš Leonardis
University of Ljubljana
Faculty of Computer and Information Science
Visual Cognitive Systems Laboratory
Trzaska 25, 1001 Ljubljana, Slovenia
E-mail: alesl@fri.uni-lj.si

Horst Bischof
Graz University of Technology
Institute for Computer Graphics and Vision
Inffeldgasse 16, 8010 Graz, Austria
E-mail: bischof@icg.tu-graz.ac.at

Axel Pinz
Graz University of Technology
Institute of Electrical Measurement and Measurement Signal Processing
Schießstattgasse 14b, 8010 Graz, Austria
E-mail: Axel.Pinz@tugraz.at

Library of Congress Control Number: 2006924180

CR Subject Classification (1998): I.4, I.3.5, I.5, I.2.9-10

LNCS Sublibrary: SL 6 – Image Processing, Computer Vision, Pattern Recognition, and Graphics

ISSN	0302-9743
ISBN-10	3-540-33834-9 Springer Berlin Heidelberg New York
ISBN-13	978-3-540-33834-5 Springer Berlin Heidelberg New York

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, re-use of illustrations, recitation, broadcasting, reproduction on microfilms 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.

Springer is a part of Springer Science+Business Media
springer.com

© Springer-Verlag Berlin Heidelberg 2006
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 11744047 06/3142 5 4 3 2 1 0

Preface

These are the proceedings of the 9th European Conference on Computer Vision (ECCV 2006), the premium European conference on computer vision, held in Graz, Austria, in May 2006.

In response to our conference call, we received 811 papers, the largest number of submissions so far. Finally, 41 papers were selected for podium presentation and 151 for presentation in poster sessions (a 23.67% acceptance rate).

The double-blind reviewing process started by assigning each paper to one of the 22 area chairs, who then selected 3 reviewers for each paper. After the reviews were received, the authors were offered the possibility to provide feedback on the reviews. On the basis of the reviews and the rebuttal of the authors, the area chairs wrote the initial consolidation report for each paper. Finally, all the area chairs attended a two-day meeting in Graz, where all decisions on acceptance/rejection were made. At that meeting, the area chairs responsible for similar sub-fields thoroughly evaluated the assigned papers and discussed them in great depth. Again, all decisions were reached without the knowledge of the authors' identity. We are fully aware of the fact that reviewing is always also subjective, and that some good papers might have been overlooked; however, we tried our best to apply a fair selection process.

The conference preparation went smoothly thanks to several people. We first wish to thank the ECCV Steering Committee for entrusting us with the organization of the conference. We are grateful to the area chairs, who did a tremendous job in selecting the papers, and to more than 340 Program Committee members and 220 additional reviewers for all their professional efforts. To the organizers of the previous ECCV 2004 in Prague, Vaclav Hlaváč, Jirí Matas and Tomáš Pajdla for providing many insights, additional information, and the superb conference software. Finally, we would also like to thank the authors for contributing a large number of excellent papers to support the high standards of the ECCV conference.

Many people showed dedication and enthusiasm in the preparation of the conference. We would like to express our deepest gratitude to all the members of the involved institutes, that is, the Institute of Electrical Measurement and Measurement Signal Processing and the Institute for Computer Graphics and Vision, both at Graz University of Technology, and the Visual Cognitive Systems Laboratory at the University of Ljubljana. In particular, we would like to express our warmest thanks to Friedrich Fraundorfer for all his help (and patience) with the conference software and many other issues concerning the event, as well as Johanna Pfeifer for her great help with the organizational matters.

February 2006

Aleš Leonardis,
Horst Bischof,
Axel Pinz

Organization

Conference Chair

Axel Pinz

Graz University of Technology, Austria

Program Chairs

Horst Bischof

Graz University of Technology, Austria

Aleš Leonardis

University of Ljubljana, Slovenia

Organization Committee

Markus Brandner

Local Arrangements

Graz Univ. of Technology,
Austria

Friedrich Fraundorfer

Local Arrangements

Graz Univ. of Technology,
Austria

Matjaž Jogan

Tutorials Chair

Univ. of Ljubljana, Slovenia

Andreas Opelt

Local Arrangements

Graz Univ. of Technology,
Austria

Johanna Pfeifer

Conference Secretariat

Graz Univ. of Technology,
Austria

Matthias R  ther

Local Arrangements

Graz Univ. of Technology,
Austria

Danijel Sko  aj

Workshops Chair

Univ. of Ljubljana, Slovenia

Conference Board

Hans Burkhardt

University of Freiburg, Germany

Bernard Buxton

University College London, UK

Roberto Cipolla

University of Cambridge, UK

Jan-Olof Eklundh

Royal Institute of Technology, Sweden

Olivier Faugeras

INRIA, Sophia Antipolis, France

Anders Heyden

Lund University, Sweden

Bernd Neumann

University of Hamburg, Germany

Mads Nielsen

IT University of Copenhagen, Denmark

Tom    Pajdla

CTU Prague, Czech Republic

Giulio Sandini

University of Genoa, Italy

David Vernon

Trinity College, Ireland

Area Chairs

Michael Black

Brown University, USA

Joachim M. Buhmann

ETH Z  rich, Switzerland

Rachid Deriche	INRIA Sophia Antipolis, France
Pascal Fua	EPFL Lausanne, Switzerland
Luc Van Gool	KU Leuven, Belgium & ETH Zürich, Switzerland
Edwin Hancock	University of York, UK
Richard Hartley	Australian National University, Australia
Sing Bing Kang	Microsoft Research, USA
Stan Li	Chinese Academy of Sciences, Beijing, China
David Lowe	University of British Columbia, Canada
Jirí Matas	CTU Prague, Czech Republic
Nikos Paragios	Ecole Centrale de Paris, France
Marc Pollefeys	University of North Carolina at Chapel Hill, USA
Long Quan	HKUST, Hong Kong, China
Bernt Schiele	Darmstadt University of Technology, Germany
Amnon Shashua	Hebrew University of Jerusalem, Israel
Peter Sturm	INRIA Rhône-Alpes, France
Chris Taylor	University of Manchester, UK
Bill Triggs	INRIA Rhône-Alpes, France
Joachim Weickert	Saarland University, Germany
Daphna Weinshall	Hebrew University of Jerusalem, Israel
Andrew Zisserman	University of Oxford, UK

Program Committee

Motilal Agrawal	Stan Birchfield	Octavia Camps
Jörgen Ahlberg	Laure Blanc-Feraud	David Capel
Miguel Alemán-Flores	Nicolas P. de la Blanca	Barbara Caputo
Yiannis Aloimonos	Volker Blanz	Stefan Carlsson
Amir Amini	Rein van den Boomgaard	Vicent Caselles
Arnon Amir	Patrick Bouthemy	Tat-Jen Cham
Elli Angelopoulou	Richard Bowden	Mike Chantler
Adnan Ansar	Edmond Boyer	Francois Chaumette
Helder Araujo	Yuri Boykov	Rama Chellappa
Tal Arbel	Francois Bremond	Tsuhan Chen
Antonis Argyros	Thomas Breuel	Dmitry Chetverikov
Karl Astrom	Lisa Brown	Ondrej Chum
Shai Avidan	Michael Brown	James Clark
Vemuri Baba	Thomas Brox	Bob Collins
Subhashis Banerjee	Alfred Bruckstein	Dorin Comaniciu
Aharon Bar-Hillel	Andres Bruhn	Tim Cootes
Kobus Barnard	Roberto Brunelli	Joao Costeira
Joao Pedro Barreto	Antoni Buades	Daniel Cremers
Chiraz Ben Abdelkader	Michael Burl	Antonio Criminisi
Marie-Odile Berger	Brian Burns	James Crowley
Marcelo Bertalmio	Darius Burschka	Kristin Dana
Ross Beveridge	Aurelio Campilho	Kostas Daniilidis

Trevor Darrell
 James W. Davis
 Fernando DelaTorre
 Herve Delingette
 Frank Dellaert
 Frederic Devernay
 Michel Dhome
 Sven Dickinson
 Zachary Dodds
 Ondrej Drbohlav
 Mark S. Drew
 Zoran Duric
 Pinar Duygulu
 Charles Dyer
 Alexei Efros
 Jan-Olof Eklundh
 James Elder
 Ahmed Elgammal
 Mark Everingham
 Aly Farag
 Paolo Favaro
 Ronald Fedkiw
 Michael Felsberg
 Rob Fergus
 Cornelia Fermüller
 Vittorio Ferrari
 Frank P. Ferrie
 James Ferryman
 Mario Figueiredo
 Graham Finlayson
 Bob Fisher
 Patrick Flynn
 Wolfgang Förstner
 Hassan Foroosh
 David Forsyth
 Friedrich Fraundorfer
 Daniel Freedman
 Andrea Fusiello
 Xiang Gao
 Nikolas Gebert
 Yakup Genc
 Guido Gerig
 Jan-Mark Geusebroek
 Christopher Geyer
 Georgy Gimel'farb

Joshua Gluckman
 Jacob Goldberger
 Dmitry Goldgof
 Venu Govindaraju
 Etienne Grossmann
 Frederic Guichard
 Yanlin Guo
 Allan Hanbury
 Horst Haussecker
 Eric Hayman
 Tamir Hazan
 Martial Hebert
 Bernd Heisele
 Anders Heyden
 R. Andrew Hicks
 Adrian Hilton
 Jeffrey Ho
 Tin Kam Ho
 David Hogg
 Ki-Sang Hong
 Anthony Hoogs
 Joachim Hornegger
 Kun Huang
 Slobodan Ilic
 Atsushi Imiya
 Sergey Ioffe
 Michael Isard
 Yuri Ivanov
 Allan D. Jepson
 Hailin Jin
 Peter Johansen
 Nebojsa Jojic
 Mike Jones
 Fredrik Kahl
 J.K. Kamarainen
 Chandra Kambhamettu
 Yoshinari Kameda
 Kenichi Kanatani
 Qifa Ke
 Daniel Keren
 Renaud Keriven
 Benjamin Kimia
 Ron Kimmel
 Nahum Kiryati
 Josef Kittler

Georges Koepfler
 Vladimir Kolmogorov
 Pierre Kornprobst
 Jana Kosecka
 Danica Kragic
 Kiriakos Kutulakos
 InSo Kweon
 Shang-Hong Lai
 Ivan Laptev
 Erik Learned-Miller
 Sang Wook Lee
 Bastian Leibe
 Christophe Lenglet
 Vincent Lepetit
 Thomas Leung
 Stephen Lin
 Michael Lindenbaum
 Jim Little
 Yanxi Liu
 Alex Loui
 Brian Lovell
 Claus Madsen
 Marcus Magnor
 Shyjan Mahamud
 Atsuto Maki
 Tom Malzbender
 R. Manmatha
 Petros Maragos
 Sebastien Marcel
 Eric Marchand
 Jorge Marques
 Jose Luis Marroquin
 David Martin
 Aleix M. Martinez
 Bogdan Matei
 Yasuyuki Matsushita
 Iain Matthews
 Stephen Maybank
 Helmut Mayer
 Leonard McMillan
 Gerard Medioni
 Etienne Memin
 Rudolf Mester
 Dimitris Metaxas
 Krystian Mikolajczyk

Majid Mirmehdi	Paolo Remagnino	Rahul Sukthankar
Anurag Mittal	Xiaofeng Ren	Josephine Sullivan
J.M.M. Montiel	Tammy Riklin-Raviv	Changming Sun
Theo Moons	Ehud Rivlin	David Suter
Philippos Mordohai	Antonio Robles-Kelly	Tomáš Svoboda
Greg Mori	Karl Rohr	Richard Szeliski
Pavel Mrázek	Sami Romdhani	Tamas Sziranyi
Jane Mulligan	Bodo Rosenhahn	Hugues Talbot
Joe Mundy	Arun Ross	Tieniu Tan
Vittorio Murino	Carsten Rother	Chi-keung Tang
Hans-Hellmut Nagel	Nicolas Rougon	Xiaouu Tang
Vic Nalwa	Mikael Rousson	Hai Tao
Srinivasa Narasimhan	Sebastien Roy	Sibel Tari
P.J. Narayanan	Javier Sanchez	Gabriel Taubin
Oscar Nestares	Jose Santos-Victor	Camillo Jose Taylor
Heiko Neumann	Guillermo Sapiro	Demetri Terzopoulos
Jan Neumann	Radim Sara	Ying-li Tian
Ram Nevatia	Jun Sato	Carlo Tomasi
Ko Nishino	Yoichi Sato	Antonio Torralba
David Nister	Eric Saund	Andrea Torsello
Thomas O'Donnell	Hanno Scharr	Panos Trahanias
Masatoshi Okutomi	Daniel Scharstein	Mohan Trivedi
Ole Fogh Olsen	Yoav Y. Schechner	Emanuele Trucco
Tomáš Pajdla	Otmar Scherzer	David Tschumperle
Chris Pal	Christoph Schnörr	Yanghai Tsin
Theodore Papadopoulos	Stan Sclaroff	Matthew Turk
Nikos Paragios	Yongduek Seo	Tinne Tuytelaars
Ioannis Pavlidis	Mubarak Shah	Nuno Vasconcelos
Vladimir Pavlovic	Gregory Shakhnarovich	Olga Veksler
Shmuel Peleg	Ying Shan	Svetha Venkatesh
Marcello Pelillo	Eitan Sharon	David Vernon
Francisco Perales	Jianbo Shi	Alessandro Verri
Sylvain Petitjean	Ilan Shimshoni	Luminita Aura Vese
Matti Pietikainen	Ali Shokoufandeh	Rene Vidal
Filiberto Pla	Kaleem Siddiqi	Markus Vincze
Robert Pless	Greg Slabaugh	Jordi Vitria
Jean Ponce	Cristian Sminchisescu	Julia Vogel
Rich Radke	Stefano Soatto	Toshikazu Wada
Ravi Ramamoorthi	Nir Sochen	Tomáš Werner
Deva Ramanan	Jon Sporring	Carl-Fredrik Westin
Visvanathan Ramesh	Anuj Srivastava	Yonatan Wexler
Ramesh Raskar	Chris Stauffer	Ross Whitaker
Christopher Rasmussen	Drew Steedly	Richard Wildes
Carlo Regazzoni	Charles Stewart	Chris Williams
James Rehg	Tomáš Suk	James Williams

Lance Williams	Jie Yang	Cha Zhang
Richard Wilson	Ming-Hsuan Yang	Song-Chun Zhu
Lior Wolf	Ruigang Yang	Todd Zickler
Kwan-Yee K. Wong	Jingyi Yu	Michael Zillich
Ming Xie	Ramin Zabih	Larry Zitnick
Yasushi Yagi	Changshui Zhang	Lilla Zöllei
Hulya Yalcin	Zhengyou Zhang	Steven Zucker

Additional Reviewers

Vitaly Ablavsky	Chi-Wei Chu	Leo Grady
Jeff Abrahamson	Andrea Colombari	Kristen Grauman
Daniel Abretske	Jason Corso	Ralph Gross
Amit Adam	Bruce Culbertson	Nicolas Guilbert
Gaurav Aggarwal	Goksel Dedeoglu	Abdenour Hadid
Amit Agrawal	David Demirdjian	Onur Hamsici
Timo Ahonen	Konstantinos Derpanis	Scott Helmer
Amir Akbarzadeh	Zvi Devir	Yacov Hel-Or
H. Can Aras	Stephan Didas	Derek Hoiem
Tamar Avraham	Miodrag Dimitrijevic	Byung-Woo Hong
Harlyn Baker	Ryan Eckbo	Steve Hordley
Patrick Baker	Christopher Engels	Changbo Hu
Hynek Bakstein	Aykut Erdem	Rui Huang
Olof Barr	Erkut Erdem	Xinyu Huang
Adrien Bartoli	Anders Ericsson	Camille Izard
Paul Beardsley	Kenny Erleben	Vidit Jain
Isabelle Bégin	Steven Eschrich	Vishal Jain
Ohad Ben-Shahar	Francisco Estrada	Christopher Jaynes
Møarten Björkman	Ricardo Fabbri	Kideog Jeong
Mark Borg	Xiaodong Fan	Björn Johansson
Jake Bouvrie	Craig Fancourt	Marie-Pierre Jolly
Bernhard Burgeth	Michela Farenzena	Erik Jonsson
Frédéric Cao	Han Feng	Klas Josephson
Gustavo Carneiro	Doug Fidaleo	Michael Kaess
Nicholas Carter	Robert Fischer	Rahul Khare
Umberto Castellani	Andrew Fitzhugh	Dae-Woong Kim
Bruno Cernuschi-Frias	Francois Fleuret	Jong-Sung Kim
Ming-Ching Chang	Per-Erik Forssén	Kristian Kirk
Roland Chapuis	Ben Fransen	Dan Kushnir
Thierry Chateau	Clement Fredembach	Ville Kyrki
Hong Chen	Mario Fritz	Pascal Lagger
Xilin Chen	Gareth Funka-Lea	Prasun Lala
Sen-ching Cheung	Darren Gawely	Michael Langer
Tat-Jun Chin	Atiyeh Ghoreyshi	Catherine Laporte
Mario Christhoudias	Alvina Goh	Jean-Marc Lavest

Albert Law	Julien Pilet	Leonid Taycher
Jean-Pierre Lecadre	David Pisinger	Ashwin Thangali
Maxime Lhuillier	Jean-Philippe Pons	David Thirde
Gang Li	Yuan Quan	Mani Thomas
Qi Li	Ariadna Quattoni	Tai-Peng Tian
Zhiguo Li	Kevin Quennesson	David Tolliver
Hwasup Lim	Ali Rahimi	Nhon Trinh
Sernam Lim	Ashish Raj	Ambrish Tyagi
Zicheng Liu	Ananath Ranganathan	Raquel Urtasun
Wei-Lwun Lu	Avinash Ravichandran	Joost Van-de-Weijer
Roberto Lubliner	Randall Rojas	Andrea Vedaldi
Simon Lucey	Mikael Rousson	Dejun Wang
Gian Luca Mariottini	Adit Sahasrabudhe	Hanzi Wang
Scott McCloskey	Roman Sandler	Jingbin Wang
Changki Min	Imari Sato	Liang Wang
Thomas Moeslund	Peter Savadjiev	Martin Welk
Kooksang Moon	Grant Schindler	Adam Williams
Louis Morency	Konrad Schindler	Bob Woodham
Davide Moschini	Robert Schwanke	Stefan Wörz
Matthias Mühlich	Edgar Seemann	Christopher Wren
Artiom Myaskovskiy	Husrev Taha Sencar	Junwen Wu
Kai Ni	Ali Shahrokni	Wen Wu
Michael Nielsen	Hong Shen	Rong Yan
Carol Novak	Fan Shufei	Changjiang Yang
Fredrik Nyberg	Johan Skoglund	Qing-Xiong Yang
Sang-Min Oh	Natalia Slesareva	Alper Yilmaz
Takahiro Okabe	Jan Sochman	Jerry Yokono
Kenki Okuma	Jan Erik Solem	David Young
Carl Olsson	Jonathan Starck	Quan Yuan
Margarita Osadchy	Jesse Stewart	Alan Yuille
Magnus Oskarsson	Henrik Stewenius	Micheal Yurick
Niels Overgaard	Moritz Stoerring	Dimitrios Zarpalas
Ozge Ozcanli	Svetlana Stolpner	Guoying Zhao
Mustafa Ozuysal	Mingxuan Sun	Tao Zhao
Vasu Parameswaran	Ying Sun	Song-Feng Zheng
Prakash Patel	Amir Tamrakar	Jie Zhu
Massimiliano Pavan	Robby Tan	Loe Zhu
Patrick Perez	Tele Tan	Manli Zhu
Michael Phelps	Donald Tanguay	

Sponsoring Institutions

Advanced Computer Vision, Austria
 Graz University of Technology, Austria
 University of Ljubljana, Slovenia

Table of Contents – Part II

Energy Minimization

Comparison of Energy Minimization Algorithms for Highly Connected Graphs <i>Vladimir Kolmogorov, Carsten Rother</i>	1
A Comparative Study of Energy Minimization Methods for Markov Random Fields <i>Richard Szeliski, Ramin Zabih, Daniel Scharstein, Olga Veksler, Vladimir Kolmogorov, Aseem Agarwala, Marshall Tappen, Carsten Rother</i>	16
Measuring Uncertainty in Graph Cut Solutions - Efficiently Computing Min-marginal Energies Using Dynamic Graph Cuts <i>Pushmeet Kohli, Philip H.S. Torr</i>	30

Tracking and Motion

Tracking Dynamic Near-Regular Texture Under Occlusion and Rapid Movements <i>Wen-Chieh Lin, Yanxi Liu</i>	44
Simultaneous Object Pose and Velocity Computation Using a Single View from a Rolling Shutter Camera <i>Omar Ait-Aider, Nicolas Andreff, Jean Marc Lavest, Philippe Martinet</i>	56
A Theory of Multiple Orientation Estimation <i>Matthias Mühlich, Til Aach</i>	69

Poster Session II

Tracking and Motion

Resolution-Aware Fitting of Active Appearance Models to Low Resolution Images <i>Göksel Dedeoğlu, Simon Baker, Takeo Kanade</i>	83
--	----

High Accuracy Optical Flow Serves 3-D Pose Tracking: Exploiting
Contour and Flow Based Constraints
*Thomas Brox, Bodo Rosenhahn, Daniel Cremers,
Hans-Peter Seidel*..... 98

Enhancing the Point Feature Tracker by Adaptive Modelling of the
Feature Support
Siniša Šegvić, Anthony Remazeilles, François Chaumette 112

Tracking Objects Across Cameras by Incrementally Learning
Inter-camera Colour Calibration and Patterns of Activity
Andrew Gilbert, Richard Bowden..... 125

Monocular Tracking of 3D Human Motion with a Coordinated Mixture
of Factor Analyzers
Rui Li, Ming-Hsuan Yang, Stan Sclaroff, Tai-Peng Tian 137

Multiview Geometry and 3D Reconstruction

Balanced Exploration and Exploitation Model Search for Efficient
Epipolar Geometry Estimation
Liran Goshen, Ilan Shimshoni 151

Shape-from-Silhouette with Two Mirrors and an Uncalibrated Camera
Keith Forbes, Fred Nicolls, Gerhard de Jager, Anthon Voigt 165

Robust and Efficient Photo-Consistency Estimation for Volumetric 3D
Reconstruction
Alexander Hornung, Leif Kobbelt..... 179

An Affine Invariant of Parallelograms and Its Application to Camera
Calibration and 3D Reconstruction
F.C. Wu, F.Q. Duan, Z.Y. Hu 191

Nonrigid Shape and Motion from Multiple Perspective Views
René Vidal, Daniel Abretske..... 205

3D Surface Reconstruction Using Graph Cuts with Surface Constraints
Son Tran, Larry Davis 219

Statistical Models and Visual Learning

Trace Quotient Problems Revisited
Shuicheng Yan, Xiaoou Tang 232

Learning Nonlinear Manifolds from Time Series <i>Ruei-Sung Lin, Che-Bin Liu, Ming-Hsuan Yang, Narendra Ahuja, Stephen Levinson</i>	245
Accelerated Convergence Using Dynamic Mean Shift <i>Kai Zhang, James T. Kwok, Ming Tang</i>	257
Efficient Belief Propagation with Learned Higher-Order Markov Random Fields <i>Xiangyang Lan, Stefan Roth, Daniel Huttenlocher, Michael J. Black</i>	269
Non Linear Temporal Textures Synthesis: A Monte Carlo Approach <i>Andrea Masiero, Alessandro Chiuso</i>	283
Low-Level Vision, Image Features	
Curvature-Preserving Regularization of Multi-valued Images Using PDE's <i>David Tschumperlé</i>	295
Higher Order Image Pyramids <i>Joshua Gluckman</i>	308
Image Specific Feature Similarities <i>Ido Omer, Michael Werman</i>	321
Coloring Local Feature Extraction <i>Joost van de Weijer, Cordelia Schmid</i>	334
Defocus Inpainting <i>Paolo Favaro, Enrico Grisan</i>	349
Viewpoint Induced Deformation Statistics and the Design of Viewpoint Invariant Features: Singularities and Occlusions <i>Andrea Vedaldi, Stefano Soatto</i>	360
Face/Gesture/Action Detection and Recognition	
Spatio-temporal Embedding for Statistical Face Recognition from Video <i>Wei Liu, Zhifeng Li, Xiaoou Tang</i>	374
Super-Resolution of 3D Face <i>Gang Pan, Shi Han, Zhaohui Wu, Yueming Wang</i>	389

Estimating Gaze Direction from Low-Resolution Faces in Video
Neil Robertson, Ian Reid 402

Learning Effective Intrinsic Features to Boost 3D-Based Face
Recognition
Chenghua Xu, Tieniu Tan, Stan Li, Yunhong Wang, Cheng Zhong . . . 416

Human Detection Using Oriented Histograms of Flow and Appearance
Navneet Dalal, Bill Triggs, Cordelia Schmid 428

Cyclostationary Processes on Shape Spaces for Gait-Based Recognition
David Kaziska, Anuj Srivastava 442

Segmentation and Grouping

Multiclass Image Labeling with Semidefinite Programming
Jens Keuchel 454

Automatic Image Segmentation by Positioning a Seed
Branislav Mičušík, Allan Hanbury 468

Patch-Based Texture Edges and Segmentation
Lior Wolf, Xiaolei Huang, Ian Martin, Dimitris Metaxas 481

Unsupervised Texture Segmentation with Nonparametric Neighborhood
Statistics
Suyash P. Awate, Tolga Tasdizen, Ross T. Whitaker 494

Detecting Symmetry and Symmetric Constellations of Features
Gareth Loy, Jan-Olof Eklundh 508

Discovering Texture Regularity as a Higher-Order Correspondence
Problem
James Hays, Marius Leordeanu, Alexei A. Efros, Yanxi Liu 522

Object Recognition, Retrieval and Indexing

Exploiting Model Similarity for Indexing and Matching to a Large
Model Database
Yi Tan, Bogdan C. Matei, Harpreet Sawhney 536

Shift-Invariant Dynamic Texture Recognition
Franco Woolfe, Andrew Fitzgibbon 549

Modeling 3D Objects from Stereo Views and Recognizing Them in Photographs <i>Akash Kushal, Jean Ponce</i>	563
A Boundary-Fragment-Model for Object Detection <i>Andreas Opelt, Axel Pinz, Andrew Zisserman</i>	575
Region Covariance: A Fast Descriptor for Detection and Classification <i>Oncel Tuzel, Fatih Porikli, Peter Meer</i>	589
Segmentation	
Affine-Invariant Multi-reference Shape Priors for Active Contours <i>Alban Foulonneau, Pierre Charbonnier, Fabrice Heitz</i>	601
Figure/Ground Assignment in Natural Images <i>Xiaofeng Ren, Charless C. Fowlkes, Jitendra Malik</i>	614
Background Cut <i>Jian Sun, Weiwei Zhang, Xiaoou Tang, Heung-Yeung Shum</i>	628
POSECUT: Simultaneous Segmentation and 3D Pose Estimation of Humans Using Dynamic Graph-Cuts <i>Matthieu Bray, Pushmeet Kohli, Philip H.S. Torr</i>	642
Author Index	657

Comparison of Energy Minimization Algorithms for Highly Connected Graphs

Vladimir Kolmogorov¹ and Carsten Rother²

¹ University College London
vnk@adastral.ucl.ac.uk

² Microsoft Research Ltd., Cambridge, UK
carrot@microsoft.com

Abstract. Algorithms for discrete energy minimization play a fundamental role for low-level vision. Known techniques include graph cuts, belief propagation (BP) and recently introduced tree-reweighted message passing (TRW). So far, the standard benchmark for their comparison has been a 4-connected grid-graph arising in pixel-labelling stereo. This minimization problem, however, has been largely solved: recent work shows that for many scenes TRW finds the global optimum. Furthermore, it is known that a 4-connected grid-graph is a poor stereo model since it does not take occlusions into account.

We propose the problem of stereo with occlusions as a new test bed for minimization algorithms. This is a more challenging graph since it has much larger connectivity, and it also serves as a better stereo model. An attractive feature of this problem is that increased connectivity does not result in increased complexity of message passing algorithms. Indeed, one contribution of this paper is to show that sophisticated implementations of BP and TRW have the same time and memory complexity as that of 4-connected grid-graph stereo.

The main conclusion of our experimental study is that for our problem graph cut outperforms both TRW and BP considerably. TRW achieves consistently a lower energy than BP. However, as connectivity increases the speed of convergence of TRW becomes slower. Unlike 4-connected grids, the difference between the energy of the best optimization method and the lower bound of TRW appears significant. This shows the hardness of the problem and motivates future research.

1 Introduction

Many early vision problems can be naturally formulated in terms of energy minimization where the energy function has the following form:

$$E(\mathbf{x}) = \sum_{p \in \mathcal{V}} D_p(x_p) + \sum_{(p,q) \in \mathcal{E}} V_{pq}(x_p, x_q). \quad (1)$$

Set \mathcal{V} usually corresponds to pixels; x_p denotes the label of pixel p which must belong to some finite set. For motion or stereo, the labels are disparities, while for image restoration they represent intensities. This energy is often derived in the context of Markov Random Fields [1]: unary terms D_p represent data likelihoods, and pairwise terms V_{pq} encode a prior over labellings. Energy minimization framework has been applied with great success to many vision applications such as stereo [2, 3, 4, 5, 6, 7],

image restoration [2], image segmentation [8], texture synthesis [9]. Algorithms for minimizing energy E are therefore of fundamental importance in vision. In this paper we consider three different algorithms: Graph Cut, belief propagation (BP) and tree-reweighted message passing (TRW). For the problem of stereo matching these methods are among the best performing optimization techniques [10]. A comparison of their advantages and disadvantages is at the end of this section.

So far, comparison studies of these optimization methods have been rather limited in the sense that they only consider energy functions with a particular graph structure [11, 12, 13, 14]. The algorithms have been tested on the energy function arising in stereo matching problem [2]. This energy is defined on a graph with a 4-neighborhood system, where nodes correspond to pixels in the left image. Occlusion are not modeled since this gives a more complex and highly connected graph structure. The comparison studies consistently concluded that the lowest energy is obtained by TRW, graph cuts come second and BP comes third [11, 12, 13, 14]. Very recently, it has been shown [13] that TRW even achieves the global optimum for standard benchmark stereo pairs [10]. Consequently, this problem, which was considered to be very challenging a decade ago, has now largely been solved. The comparison studies also showed that the proposed energy gives large error statistics compared with state-of-the-art methods, and consequently progress in this field can only be achieved by improving the energy formulation itself, as stated in [11, 13].

The main goal of this paper is to test how different optimization methods perform on graphs with larger connectivity. Our study has two motivations. First, such energy functions are becoming increasingly important in vision [3, 4, 5, 6, 7, 15]. They typically arise when we need to match two images while imposing regularization on the deformation field. Pixels (or features) in one image can potentially match to many pixels (features) in the other image, which yields a highly connected graph structure.

Our second motivation is to understand better intrinsic properties of different algorithms. One way to achieve this is to consider a very difficult problem: Algorithm's weaknesses then become more apparent, which may suggest ways of improving the method. It is known that the presence of short cycles in the graph makes the problem harder for message passing techniques. From this point of view, the problem that we are considering is much more challenging than 4-connected grid graphs. Another indicator of the difficulty of our problem will be shown by our experiments.

We choose the energy function arising in the problem of stereo with occlusions [4]. In this case there are nodes corresponding to pixels in the left and right image, and each node has $K + 4$ neighbors where K is the number of disparities. We propose this problem as a new challenging test bed for minimization algorithms. Our experiments also show that modeling occlusions gives a significantly better stereo model, since the energy of the ground truth is close to the energy of the best optimization method, and the value of the energy correlates with the error statistics derived from ground truth.

When applying BP or TRW to this energy, we immediately run into efficiency problems. There are K labels and $O(NK)$ edges, so a straightforward implementation would take $O(NK^2)$ memory and time for one iteration, even with the distance transform technique in [16]. By exploiting a special structure of the energy we show that