

2013 IEEE Conference on Computer Vision and Pattern Recognition

(CVPR 2013)

**Portland, Oregon, USA
23 - 28 June 2013**

Pages 1-746



**IEEE Catalog Number: CFP13003-POD
ISBN: 978-1-4673-6410-2**

2013 IEEE Conference on Computer Vision and Pattern Recognition

CVPR 2013

Table of Contents

Message from the General and Program

Chairs.....	xxxvii
Organizing Committee.....	xxxix
Reviewers.....	xlii

Orals 1A: 3D Imaging and Reasoning

3D-Based Reasoning with Blocks, Support, and Stability	1
<i>Zhaoyin Jia, Andrew Gallagher, Ashutosh Saxena, and Tsuhan Chen</i>	
Physically Plausible 3D Scene Tracking: The Single Actor Hypothesis	9
<i>Nikolaos Kyriazis and Antonis Argyros</i>	
Intrinsic Scene Properties from a Single RGB-D Image	17
<i>Jonathan T. Barron and Jitendra Malik</i>	
Depth Acquisition from Density Modulated Binary Patterns	25
<i>Zhe Yang, Zhiwei Xiong, Yueyi Zhang, Jiao Wang, and Feng Wu</i>	
Understanding Indoor Scenes Using 3D Geometric Phrases	33
<i>Wongun Choi, Yu-Wei Chao, Caroline Pantofaru, and Silvio Savarese</i>	

Orals 1B: Statistics and Learning

Rolling Riemannian Manifolds to Solve the Multi-class Classification Problem	41
<i>Rui Caseiro, Pedro Martins, João F. Henriques, Fátima Silva Leite, and Jorge Batista</i>	
Exploring Compositional High Order Pattern Potentials for Structured Output Learning	49
<i>Yujia Li, Daniel Tarlow, and Richard Zemel</i>	
Discrete MRF Inference of Marginal Densities for Non-uniformly Discretized Variable Space	57
<i>Masaki Saito, Takayuki Okatani, and Koichiro Deguchi</i>	

GeoF: Geodesic Forests for Learning Coupled Predictors	65
<i>Peter Kotschieder, Pushmeet Kohli, Jamie Shotton, and Antonio Criminisi</i>	
Kernel Methods on the Riemannian Manifold of Symmetric Positive Definite Matrices	73
<i>Sadeep Jayasumana, Richard Hartley, Mathieu Salzmann, Hongdong Li, and Mehrtash Harandi</i>	
Posters 1A: 3D and Stereo	
Manhattan Scene Understanding via XSlit Imaging	81
<i>Jinwei Ye, Yu Ji, and Jingyi Yu</i>	
Discovering the Structure of a Planar Mirror System from Multiple Observations of a Single Point	89
<i>Ilya Reshetouski, Alkhazur Manakov, Ayush Bandhari, Ramesh Raskar, Hans-Peter Seidel, and Ivo Ihrke</i>	
Joint 3D Scene Reconstruction and Class Segmentation	97
<i>Christian Häne, Christopher Zach, Andrea Cohen, Roland Angst, and Marc Pollefeys</i>	
Tensor-Based Human Body Modeling	105
<i>Yinpeng Chen, Zicheng Liu, and Zhengyou Zhang</i>	
City-Scale Change Detection in Cadastral 3D Models Using Images	113
<i>Aparna Taneja, Luca Ballan, and Marc Pollefeys</i>	
Improving the Visual Comprehension of Point Sets	121
<i>Sagi Katz and Ayellet Tal</i>	
Mirror Surface Reconstruction from a Single Image	129
<i>Miaomiao Liu, Richard Hartley, and Mathieu Salzmann</i>	
Detecting Changes in 3D Structure of a Scene from Multi-view Images Captured by a Vehicle-Mounted Camera	137
<i>Ken Sakurada, Takayuki Okatani, and Koichiro Deguchi</i>	
Templateless Quasi-rigid Shape Modeling with Implicit Loop-Closure	145
<i>Ming Zeng, Jiaxiang Zheng, Xuan Cheng, and Xinguo Liu</i>	
Understanding Bayesian Rooms Using Composite 3D Object Models	153
<i>Luca Del Pero, Joshua Bowdish, Bonnie Kermgard, Emily Hartley, and Kobus Barnard</i>	
Shape from Silhouette Probability Maps: Reconstruction of Thin Objects in the Presence of Silhouette Extraction and Calibration Error	161
<i>Amy Tabb</i>	
Joint Geodesic Upsampling of Depth Images	169
<i>Ming-Yu Liu, Oncel Tuzel, and Yuichi Taguchi</i>	

Relative Volume Constraints for Single View 3D Reconstruction	177
<i>Eno Töppe, Claudia Nieuwenhuis, and Daniel Cremers</i>	
Is There a Procedural Logic to Architecture?	185
<i>Julien Weissenberg, Hayko Riemenschneider, Mukta Prasad, and Luc Van Gool</i>	
Category Modeling from Just a Single Labeling: Use Depth Information to Guide the Learning of 2D Models	193
<i>Quanshi Zhang, Xuan Song, Xiaowei Shao, Ryosuke Shibasaki, and Huijing Zhao</i>	
Bayesian Grammar Learning for Inverse Procedural Modeling	201
<i>Andelo Martinovic and Luc Van Gool</i>	
Fusing Depth from Defocus and Stereo with Coded Apertures	209
<i>Yuichi Takeda, Shinsaku Hiura, and Kosuke Sato</i>	
Bayesian Depth-from-Defocus with Shading Constraints	217
<i>Chen Li, Shuochen Su, Yasuyuki Matsushita, Kun Zhou, and Stephen Lin</i>	
Multi-scale Curve Detection on Surfaces	225
<i>Michael Kolomenkin, Ilan Shimshoni, and Ayellet Tal</i>	
Intrinsic Characterization of Dynamic Surfaces	233
<i>Tony Tung and Takashi Matsuyama</i>	
Pattern-Driven Colorization of 3D Surfaces	241
<i>George Leifman and Ayellet Tal</i>	
Three-Dimensional Bilateral Symmetry Plane Estimation in the Phase Domain	249
<i>Ramakrishna Kakarala, Prabhu Kaliamoorthi, and Vittal Premachandran</i>	
Axially Symmetric 3D Pots Configuration System Using Axis of Symmetry and Break Curve	257
<i>Kilho Son, Eduardo B. Almeida, and David B. Cooper</i>	
Wide-Baseline Hair Capture Using Strand-Based Refinement	265
<i>Linjie Luo, Cha Zhang, Zhengyou Zhang, and Szymon Rusinkiewicz</i>	
Dense 3D Reconstruction from Severely Blurred Images Using a Single Moving Camera	273
<i>Hee Seok Lee and Kuoung Mu Lee</i>	
Simultaneous Super-Resolution of Depth and Images Using a Single Camera	281
<i>Hee Seok Lee and Kuoung Mu Lee</i>	
Recovering Stereo Pairs from Anaglyphs	289
<i>Armand Joulin and Sing Bing Kang</i>	
Exploiting the Power of Stereo Confidences	297
<i>David Pfeiffer, Stefan Gehrig, and Nicolai Schneider</i>	
Ensemble Learning for Confidence Measures in Stereo Vision	305
<i>Ralf Haeusler, Rahul Nair, and Daniel Kondermann</i>	

Segment-Tree Based Cost Aggregation for Stereo Matching	313
<i>Xing Mei, Xun Sun, Weiming Dong, Haitao Wang, and Xiaopeng Zhang</i>	

Posters 1B: Statistics and Learning

Multi-class Video Co-segmentation with a Generative Multi-video Model	321
<i>Wei-Chen Chiu and Mario Fritz</i>	
A Bayesian Approach to Multimodal Visual Dictionary Learning	329
<i>Go Irie, Dong Liu, Zhenguo Li, and Shih-Fu Chang</i>	
A Statistical Model for Recreational Trails in Aerial Images	337
<i>Andrew Predoehl, Scott Morris, and Kobus Barnard</i>	
Beta Process Joint Dictionary Learning for Coupled Feature Spaces with Application to Single Image Super-Resolution	345
<i>Li He, Hairong Qi, and Russell Zaretzki</i>	
Dictionary Learning from Ambiguously Labeled Data	353
<i>Yi-Chen Chen, Vishal M. Patel, Jaishanker K. Pillai, Rama Chellappa, and P. Jonathon Phillips</i>	
Generalized Domain-Adaptive Dictionaries	361
<i>Sumit Shekhar, Vishal M. Patel, Hien V. Nguyen, and Rama Chellappa</i>	
Tag Taxonomy Aware Dictionary Learning for Region Tagging	369
<i>Jingjing Zheng and Zhuolin Jiang</i>	
Block and Group Regularized Sparse Modeling for Dictionary Learning	377
<i>Yu-Tseh Chi, Mohsen Ali, Ajit Rajwade, and Jeffrey Ho</i>	
Multi-level Discriminative Dictionary Learning towards Hierarchical Visual Categorization	383
<i>Li Shen, Shuhui Wang, Gang Sun, Shuqiang Jiang, and Qingming Huang</i>	
Fast Convolutional Sparse Coding	391
<i>Hilton Bristow, Anders Eriksson, and Simon Lucey</i>	
In Defense of Sparsity Based Face Recognition	399
<i>Weihong Deng, Jiani Hu, and Jun Guo</i>	
Transfer Sparse Coding for Robust Image Representation	407
<i>Mingsheng Long, Guiguang Ding, Jianmin Wang, Jianguang Sun, Yuchen Guo, and Philip S. Yu</i>	
Online Robust Dictionary Learning	415
<i>Cewu Lu, Jiaping Shi, and Jiaya Jia</i>	
Multi-task Sparse Learning with Beta Process Prior for Action Recognition	423
<i>Chunfeng Yuan, Weiming Hu, Guodong Tian, Shuang Yang, and Haoran Wang</i>	
Scalable Sparse Subspace Clustering	430
<i>Xi Peng, Lei Zhang, and Zhang Yi</i>	

Separable Dictionary Learning	438
<i>Simon Hawe, Matthias Seibert, and Martin Kleinsteuber</i>	
Compressed Hashing	446
<i>Yue Lin, Rong Jin, Deng Cai, Shuicheng Yan, and Xuelong Li</i>	
Improved Image Set Classification via Joint Sparse Approximated Nearest Subspaces	452
<i>Shaokang Chen, Conrad Sanderson, Mehrtash T. Harandi, and Brian C. Lovell</i>	
Optimizing 1-Nearest Prototype Classifiers	460
<i>Paul Wohlhart, Martin Köstinger, Michael Donoser, Peter M. Roth, and Horst Bischof</i>	
Sparse Subspace Denoising for Image Manifolds	468
<i>Bo Wang and Zhuowen Tu</i>	
Weakly Supervised Learning of Mid-Level Features with Beta-Bernoulli Process Restricted Boltzmann Machines	476
<i>Roni Mittelman, Honglak Lee, Benjamin Kuipers, and Silvio Savarese</i>	
Learning Binary Codes for High-Dimensional Data Using Bilinear Projections	484
<i>Yunchao Gong, Sanjiv Kumar, Henry A. Rowley, and Svetlana Lazebnik</i>	
Semi-supervised Node Splitting for Random Forest Construction	492
<i>Xiao Liu, Mingli Song, Dacheng Tao, Zicheng Liu, Luming Zhang, Chun Chen, and Jiajun Bu</i>	
Capturing Layers in Image Collections with Componential Models: From the Layered Epitome to the Componential Counting Grid	500
<i>Alessandro Perina and Nebojsa Jojic</i>	
Alternating Decision Forests	508
<i>Samuel Schulter, Paul Wohlhart, Christian Leistner, Amir Saffari, Peter M. Roth, and Horst Bischof</i>	
Exploring Implicit Image Statistics for Visual Representativeness Modeling	516
<i>Xiaoshuai Sun, Xin-Jing Wang, Hongxun Yao, and Lei Zhang</i>	
A Divide-and-Conquer Method for Scalable Low-Rank Latent Matrix Pursuit	524
<i>Yan Pan, Hanjiang Lai, Cong Liu, and Shuicheng Yan</i>	
Supervised Descent Method and Its Applications to Face Alignment	532
<i>Xuehan Xiong and Fernando De la Torre</i>	
Robust Canonical Time Warping for the Alignment of Grossly Corrupted Sequences	540
<i>Yannis Panagakis, Mihalis A. Nicolaou, Stefanos Zafeiriou, and Maja Pantic</i>	
Relative Hidden Markov Models for Evaluating Motion Skill	548
<i>Qiang Zhang and Baoxin Li</i>	
A Fast Approximate AIB Algorithm for Distributional Word Clustering	556
<i>Lei Wang, Jianjia Zhang, Luping Zhou, and Wanqing Li</i>	

Orals 1C: Recognition

Perceptual Organization and Recognition of Indoor Scenes from RGB-D Images	564
<i>Saurabh Gupta, Pablo Arbeláez, and Jitendra Malik</i>	
Watching Unlabeled Video Helps Learn New Human Actions from Very Few Labeled Snapshots	572
<i>Chao-Yeh Chen and Kristen Grauman</i>	
Fine-Grained Crowdsourcing for Fine-Grained Recognition	580
<i>Jia Deng, Jonathan Krause, and Li Fei-Fei</i>	
Poselet Conditioned Pictorial Structures	588
<i>Leonid Pishchulin, Mykhaylo Andriluka, Peter Gehler, and Bernt Schiele</i>	
Beyond Physical Connections: Tree Models in Human Pose Estimation	596
<i>Fang Wang and Yi Li</i>	

Orals 1D: Imaging and Segmentation

Discriminative Non-blind Deblurring	604
<i>Uwe Schmidt, Carsten Rother, Sebastian Nowozin, Jeremy Jancsary, and Stefan Roth</i>	
Handling Noise in Single Image Deblurring Using Directional Filters	612
<i>Lin Zhong, Sunghyun Cho, Dimitris Metaxas, Sylvain Paris, and Jue Wang</i>	
Jointly Aligning and Segmenting Multiple Web Photo Streams for the Inference of Collective Photo Storylines	620
<i>Gunhee Kim and Eric P. Xing</i>	
Video Object Segmentation through Spatially Accurate and Temporally Dense Extraction of Primary Object Regions	628
<i>Dong Zhang, Omar Javed, and Mubarak Shah</i>	
Improving Image Matting Using Comprehensive Sampling Sets	636
<i>Ehsan Shahrian, Deepu Rajan, Brian Price, and Scott Cohen</i>	

Posters 1C: Recognition

Simultaneous Active Learning of Classifiers & Attributes via Relative Feedback	644
<i>Arijit Biswas and Devi Parikh</i>	
Expanded Parts Model for Human Attribute and Action Recognition in Still Images	652
<i>Gaurav Sharma, Frédéric Jurie, and Cordelia Schmid</i>	
Multipath Sparse Coding Using Hierarchical Matching Pursuit	660
<i>Liefeng Bo, Xiaofeng Ren, and Dieter Fox</i>	

Semi-supervised Domain Adaptation with Instance Constraints	668
<i>Jeff Donahue, Judy Hoffman, Erik Rodner, Kate Saenko, and Trevor Darrell</i>	
Learning Structured Low-Rank Representations for Image Classification	676
<i>Yangmuzi Zhang, Zhuolin Jiang, and Larry S. Davis</i>	
MKPLS: Manifold Kernel Partial Least Squares for Lipreading and Speaker Identification	684
<i>Amr Bakry and Ahmed Elgammal</i>	
Subspace Interpolation via Dictionary Learning for Unsupervised Domain Adaptation	692
<i>Jie Ni, Qiang Qiu, and Rama Chellappa</i>	
Graph-Based Discriminative Learning for Location Recognition	700
<i>Song Cao and Noah Snavely</i>	
Learning by Associating Ambiguously Labeled Images	708
<i>Zinan Zeng, Shijie Xiao, Kui Jia, Tsung-Han Chan, Shenghua Gao, Dong Xu, and Yi Ma</i>	
HON4D: Histogram of Oriented 4D Normals for Activity Recognition from Depth Sequences	716
<i>Omar Oreifej and Zicheng Liu</i>	
3D R Transform on Spatio-temporal Interest Points for Action Recognition	724
<i>Chunfeng Yuan, Xi Li, Weiming Hu, Haibin Ling, and Stephen Maybank</i>	
Learning Cross-Domain Information Transfer for Location Recognition and Clustering	731
<i>Raghuraman Gopalan</i>	
Studying Relationships between Human Gaze, Description, and Computer Vision	739
<i>Kiwon Yun, Yifan Peng, Dimitris Samaras, Gregory J. Zelinsky, and Tamara L. Berg</i>	
BOØMeets HOG: Feature Extraction Based on Histograms of Oriented p.d.f. Gradients for Image Classification	747
<i>Takumi Kobayashi</i>	
Class Generative Models Based on Feature Regression for Pose Estimation of Object Categories	755
<i>Michele Fenzi, Laura Leal-Taixé, Bodo Rosenhahn, and Jörn Ostermann</i>	
Leveraging Structure from Motion to Learn Discriminative Codebooks for Scalable Landmark Classification	763
<i>Alessandro Bergamo, Sudipta N. Sinha, and Lorenzo Torresani</i>	
Designing Category-Level Attributes for Discriminative Visual Recognition	771
<i>Felix X. Yu, Liangliang Cao, Rogerio S. Feris, John R. Smith, and Shih-Fu Chang</i>	
Attribute-Based Detection of Unfamiliar Classes with Humans in the Loop	779
<i>Catherine Wah and Serge Belongie</i>	

Object-Centric Anomaly Detection by Attribute-Based Reasoning	787
<i>Babak Saleh, Ali Farhadi, and Ahmed Elgammal</i>	
Learning Class-to-Image Distance with Object Matchings	795
<i>Guang-Tong Zhou, Tian Lan, Weilong Yang, and Greg Mori</i>	
Sample-Specific Late Fusion for Visual Category Recognition	803
<i>Dong Liu, Kuan-Ting Lai, Guangnan Ye, Ming-Syan Chen, and Shih-Fu Chang</i>	
Efficient Object Detection and Segmentation for Fine-Grained Recognition	811
<i>Anelia Angelova and Shenghuo Zhu</i>	
Label-Embedding for Attribute-Based Classification	819
<i>Zeynep Akata, Florent Perronnin, Zaid Harchaoui, and Cordelia Schmid</i>	
Subcategory-Aware Object Classification	827
<i>Jian Dong, Wei Xia, Qiang Chen, Jianshi Feng, Zhongyang Huang, and Shuicheng Yan</i>	
Vantage Feature Frames for Fine-Grained Categorization	835
<i>Asma Rejeb Sfar, Nozha Boujemaa, and Donald Geman</i>	
Probabilistic Label Trees for Efficient Large Scale Image Classification	843
<i>Baoyuan Liu, Fereshteh Sadeghi, Marshall Tappen, Ohad Shamir, and Ce Liu</i>	
Harvesting Mid-level Visual Concepts from Large-Scale Internet Images	851
<i>Quannan Li, Jiajun Wu, and Zhuowen Tu</i>	
Adaptive Active Learning for Image Classification	859
<i>Xin Li and Yuhong Guo</i>	
SCaLE: Supervised and Cascaded Laplacian Eigenmaps for Visual Object Recognition Based on Nearest Neighbors	867
<i>Ruobing Wu, Yizhou Yu, and Wenping Wang</i>	
Adding Unlabeled Samples to Categories by Learned Attributes	875
<i>Jonghyun Choi, Mohammad Rastegari, Ali Farhadi, and Larry S. Davis</i>	
Visual Place Recognition with Repetitive Structures	883
<i>Akihiko Torii, Josef Sivic, Tomáš Pajdla, and Masatoshi Okutomi</i>	
Cross-View Image Geolocalization	891
<i>Tsung-Yi Lin, Serge Belongie, and James Hays</i>	
Efficient 2D-to-3D Correspondence Filtering for Scalable 3D Object Recognition	899
<i>Qiang Hao, Rui Cai, Zhiwei Li, Lei Zhang, Yanwei Pang, Feng Wu, and Yong Rui</i>	
Learning and Calibrating Per-Location Classifiers for Visual Place Recognition	907
<i>Petr Gronát, Guillaume Obozinski, Josef Sivic, and Tomáš Pajdla</i>	
An Approach to Pose-Based Action Recognition	915
<i>Chunyu Wang, Yizhou Wang, and Alan L. Yuille</i>	

Blocks That Shout: Distinctive Parts for Scene Classification	923
<i>Mayank Juneja, Andrea Vedaldi, C.V. Jawahar, and Andrew Zisserman</i>	
Part Discovery from Partial Correspondence	931
<i>Subhransu Maji and Gregory Shakhnarovich</i>	
Learning Collections of Part Models for Object Recognition	939
<i>Ian Endres, Kevin J. Shih, Johnston Jiaa, and Derek Hoiem</i>	
Fast Multiple-Part Based Object Detection Using KD-Ferns	947
<i>Dan Levi, Shai Silberstein, and Aharon Bar-Hillel</i>	
POOF: Part-Based One-vs.-One Features for Fine-Grained Categorization, Face Verification, and Attribute Estimation	955
<i>Thomas Berg and Peter N. Belhumeur</i>	

Posters 1D: Imaging

Non-parametric Filtering for Geometric Detail Extraction and Material Representation	963
<i>Zicheng Liao, Jason Rock, Yang Wang, and David Forsyth</i>	
Learning the Change for Automatic Image Cropping	971
<i>Jianzhou Yan, Stephen Lin, Sing Bing Kang, and Xiaoou Tang</i>	
Statistical Textural Distinctiveness for Salient Region Detection in Natural Images	979
<i>Christian Scharfenberger, Alexander Wong, Khalil Fergani, John S. Zelek, and David A. Clausi</i>	
Real-Time No-Reference Image Quality Assessment Based on Filter Learning	987
<i>Peng Ye, Jayant Kumar, Le Kang, and David Doermann</i>	
Learning without Human Scores for Blind Image Quality Assessment	995
<i>Wufeng Xue, Lei Zhang, and Xuanqin Mou</i>	
The Variational Structure of Disparity and Regularization of 4D Light Fields	1003
<i>Bastian Goldluecke and Sven Wanner</i>	
Globally Consistent Multi-label Assignment on the Ray Space of 4D Light Fields	1011
<i>Sven Wanner, Christoph Straehle, and Bastian Goldluecke</i>	
Principal Observation Ray Calibration for Tiled-Lens-Array Integral Imaging Display	1019
<i>Weiming Li, Haitao Wang, Mingcai Zhou, Shandong Wang, Shaohui Jiao, Xing Mei, Tao Hong, Hoyoung Lee, and Jiyeun Kim</i>	
Decoding, Calibration and Rectification for Lenselet-Based Plenoptic Cameras	1027
<i>Donald G. Dansereau, Oscar Pizarro, and Stefan B. Williams</i>	
Adherent Raindrop Detection and Removal in Video	1035
<i>Shaodi You, Robby T. Tan, Rei Kawakami, and Katsushi Ikeuchi</i>	

Stochastic Deconvolution	1043
<i>James Gregson, Felix Heide, Matthias B. Hullin, Mushfiqur Rouf, and Wolfgang Heidrich</i>	
Multi-image Blind Deblurring Using a Coupled Adaptive Sparse Prior	1051
<i>Haichao Zhang, David Wipf, and Yanning Zhang</i>	
Fast Image Super-Resolution Based on In-Place Example Regression	1059
<i>Jianchao Yang, Zhe Lin, and Scott Cohen</i>	
A Machine Learning Approach for Non-blind Image Deconvolution	1067
<i>Christian J. Schuler, Harold Christopher Burger, Stefan Harmeling, and Bernhard Schölkopf</i>	
Learning to Estimate and Remove Non-uniform Image Blur	1075
<i>Florent Couzinié-Devy, Jian Sun, Karteek Alahari, and Jean Ponce</i>	
On a Link Between Kernel Mean Maps and Fraunhofer Diffraction, with an Application to Super-Resolution Beyond the Diffraction Limit	1083
<i>Stefan Harmeling, Michael Hirsch, and Bernhard Schölkopf</i>	
Blur Processing Using Double Discrete Wavelet Transform	1091
<i>Yi Zhang and Keigo Hiraoka</i>	
Structured Face Hallucination	1099
<i>Chih-Yuan Yang, Sifei Liu, and Ming-Hsuan Yang</i>	
Unnatural L0 Sparse Representation for Natural Image Deblurring	1107
<i>Li Xu, Shicheng Zheng, and Jiaya Jia</i>	
Non-uniform Motion Deblurring for Bilayer Scenes	1115
<i>Chandramouli Paramanand and Ambasamudram N. Rajagopalan</i>	
Depth Super Resolution by Rigid Body Self-Similarity in 3D	1123
<i>Michael Hornáček, Christoph Rhemann, Margrit Gelautz, and Carsten Rother</i>	
Saliency Aggregation: A Data-Driven Approach	1131
<i>Long Mai, Yuzhen Niu, and Feng Liu</i>	
What Makes a Patch Distinct?	1139
<i>Ran Margolin, Ayellet Tal, and Lihi Zelnik-Manor</i>	
Learning Video Saliency from Human Gaze Using Candidate Selection	1147
<i>Dmitry Rudoy, Dan B. Goldman, Eli Shechtman, and Lihi Zelnik-Manor</i>	
Hierarchical Saliency Detection	1155
<i>Qiong Yan, Li Xu, Jianping Shi, and Jiaya Jia</i>	
HDR Deghosting: How to Deal with Saturation?	1163
<i>Jun Hu, Orazio Gallo, Kari Pulli, and Xiaobai Sun</i>	
FrameBreak: Dramatic Image Extrapolation by Guided Shift-Maps	1171
<i>Yinda Zhang, Jianxiong Xiao, James Hays, and Ping Tan</i>	

Video Enhancement of People Wearing Polarized Glasses: Darkening Reversal and Reflection Reduction	1179
<i>Mao Ye, Cha Zhang, and Ruigang Yang</i>	
Layer Depth Denoising and Completion for Structured-Light RGB-D Cameras	1187
<i>Ju Shen and Sen-Ching S. Cheung</i>	
Separating Signal from Noise Using Patch Recurrence across Scales	1195
<i>Maria Zontak, Inbar Mosseri, and Michal Irani</i>	
Texture Enhanced Image Denoising via Gradient Histogram Preservation	1203
<i>Wangmeng Zuo, Lei Zhang, Chunwei Song, and David Zhang</i>	
Fast Patch-Based Denoising Using Approximated Patch Geodesic Paths	1211
<i>Xiaogang Chen, Sing Bing Kang, Jie Yang, and Jingyi Yu</i>	
A New Model and Simple Algorithms for Multi-label Mumford-Shah Problems	1219
<i>Byung-Woo Hong, Zhaojin Lu, and Ganesh Sundaramoorthi</i>	
Computing Diffeomorphic Paths for Large Motion Interpolation	1227
<i>Dohyung Seo, Jeffrey Ho, and Baba C. Vemuri</i>	
Rotation, Scaling and Deformation Invariant Scattering for Texture Discrimination	1233
<i>Laurent Sifre and Stéphane Mallat</i>	
Sensing and Recognizing Surface Textures Using a GelSight Sensor	1241
<i>Rui Li and Edward H. Adelson</i>	
Enriching Texture Analysis with Semantic Data	1248
<i>Tim Matthews, Mark S. Nixon, and Mahesan Niranjan</i>	

Orals 2A: Motion and Reconstruction

Megastereo: Constructing High-Resolution Stereo Panoramas	1256
<i>Christian Richardt, Yael Pritch, Henning Zimmer, and Alexander Sorkine-Hornung</i>	
Dense Object Reconstruction with Semantic Priors	1264
<i>Sid Yingze Bao, Manmohan Chandraker, Yuanqing Lin, and Silvio Savarese</i>	
Dense Variational Reconstruction of Non-rigid Surfaces from Monocular Video	1272
<i>Ravi Garg, Anastasios Roussos, and Lourdes Agapito</i>	
Procrustean Normal Distribution for Non-rigid Structure from Motion	1280
<i>Minsik Lee, Jungchan Cho, Chong-Ho Choi, and Songhwai Oh</i>	
Dense Reconstruction Using 3D Object Shape Priors	1288
<i>Amaury Dame, Victor A. Prisacariu, Carl Y. Ren, and Ian Reid</i>	

Orals 2B: Optimization Methods

Gauging Association Patterns of Chromosome Territories via Chromatic Median	1296
<i>Hu Ding, Branislav Stojkovic, Ronald Berezney, and Jinhui Xu</i>	
Auxiliary Cuts for General Classes of Higher Order Functionals	1304
<i>Ismail Ben Ayed, Lena Gorelick, and Yuri Boykov</i>	
A Fast Semidefinite Approach to Solving Binary Quadratic Problems	1312
<i>Peng Wang, Chunhua Shen, and Anton van den Hengel</i>	
Diffusion Processes for Retrieval Revisited	1320
<i>Michael Donoser and Horst Bischof</i>	
A Comparative Study of Modern Inference Techniques for Discrete Energy Minimization Problems	1328
<i>Jörg H. Kappes, Bjoern Andres, Fred A. Hamprecht, Christoph Schnörr, Sebastian Nowozin, Dhruv Batra, Sungwoong Kim, Bernhard X. Kausler, Jan Lellmann, Nikos Komodakis, and Carsten Rother</i>	

Posters 2A: Pose and Photometry

A Global Approach for the Detection of Vanishing Points and Mutually Orthogonal Vanishing Directions	1336
<i>Michel Antunes and João P. Barreto</i>	
Cloud Motion as a Calibration Cue	1344
<i>Nathan Jacobs, Mohammad T. Islam, and Scott Workman</i>	
SLAM++: Simultaneous Localisation and Mapping at the Level of Objects	1352
<i>Renato F. Salas-Moreno, Richard A. Newcombe, Hauke Strasdat, Paul H.J. Kelly, and Andrew J. Davison</i>	
Rolling Shutter Camera Calibration	1360
<i>Luc Oth, Paul Furgale, Laurent Kneip, and Roland Siegwart</i>	
Radial Distortion Self-Calibration	1368
<i>José Henrique Brito, Roland Angst, Kevin Köser, and Marc Pollefeys</i>	
A Minimum Error Vanishing Point Detection Approach for Uncalibrated Monocular Images of Man-Made Environments	1376
<i>Yiliang Xu, Sangmin Oh, and Anthony Hoogs</i>	
Five Shades of Grey for Fast and Reliable Camera Pose Estimation	1384
<i>Adam Herout, István Szentandrás, Michal Zachariáš, Markéta Dubská, and Rudolf Kajan</i>	
Can a Fully Unconstrained Imaging Model Be Applied Effectively to Central Cameras?	1391
<i>Filippo Bergamasco, Andrea Albarelli, Emanuele Rodolà, and Andrea Torsello</i>	

Single Image Calibration of Multi-axial Imaging Systems	1399
<i>Amit Agrawal and Srikumar Ramalingam</i>	
The Epipolar Constraint: Monocular Shape from Shadow Correspondence	1407
<i>Austin Abrams, Kyliia Miskell, and Robert Pless</i>	
Shading-Based Shape Refinement of RGB-D Images	1415
<i>Lap-Fai Yu, Sai-Kit Yeung, Yu-Wing Tai, and Stephen Lin</i>	
Illumination Estimation Based on Bilayer Sparse Coding	1423
<i>Bing Li, Weihua Xiong, Weiming Hu, and Houwen Peng</i>	
Learning Discriminative Illumination and Filters for Raw Material Classification with Optimal Projections of Bidirectional Texture Functions	1430
<i>Chao Liu, Geifei Yang, and Jinwei Gu</i>	
A Theory of Refractive Photo-Light-Path Triangulation	1438
<i>Visesh Chari and Peter Sturm</i>	
Analytic Bilinear Appearance Subspace Construction for Modeling Image Irradiance under Natural Illumination and Non-Lambertian Reflectance	1446
<i>Shireen Y. Elhabian and Aly A. Farag</i>	
Spectral Modeling and Relighting of Reflective-Fluorescent Scenes	1452
<i>Antony Lam and Imari Sato</i>	
Specular Reflection Separation Using Dark Channel Prior	1460
<i>Hyeongwoo Kim, Hailin Jin, Sunil Hadap, and Inso Kweon</i>	
BRDF Slices: Accurate Adaptive Anisotropic Appearance Acquisition	1468
<i>Jirí Filip, Radomír Vávra, Michal Haindl, Pavel Žid, Mikuláš Krupika, and Vlastimil Havran</i>	
A New Perspective on Uncalibrated Photometric Stereo	1474
<i>Thoma Papadhimetri and Paolo Favaro</i>	
Multi-view Photometric Stereo with Spatially Varying Isotropic Materials	1482
<i>Zhenglong Zhou, Zhe Wu, and Ping Tan</i>	
Uncalibrated Photometric Stereo for Unknown Isotropic Reflectances	1490
<i>Feng Lu, Yasuyuki Matsushita, Imari Sato, Takahiro Okabe, and Yoichi Sato</i>	
Calibrating Photometric Stereo by Holistic Reflectance Symmetry Analysis	1498
<i>Zhe Wu and Ping Tan</i>	
Articulated and Restricted Motion Subspaces and Their Signatures	1506
<i>Bastien Jacquet, Roland Angst, and Marc Pollefeys</i>	
Template-Based Isometric Deformable 3D Reconstruction with Sampling-Based Focal Length Self-Calibration	1514
<i>Adrien Bartoli and Toby Collins</i>	

Monocular Template-Based 3D Reconstruction of Extensible Surfaces with Local Linear Elasticity	1522
<i>Abed Malti, Richard Hartley, Adrien Bartoli, and Jae-Hak Kim</i>	
Non-rigid Structure from Motion with Diffusion Maps Prior	1530
<i>Lili Tao and Bogdan J. Matuszewski</i>	
Joint Detection, Tracking and Mapping by Semantic Bundle Adjustment	1538
<i>Nicola Fioraio and Luigi Di Stefano</i>	
A Practical Rank-Constrained Eight-Point Algorithm for Fundamental Matrix Estimation	1546
<i>Yinqiang Zheng, Shigeki Sugimoto, and Masatoshi Okutomi</i>	
CLAM: Coupled Localization and Mapping with Efficient Outlier Handling	1554
<i>Jonathan Balzer and Stefano Soatto</i>	

Posters 2B: Methods and Retrieval

Inductive Hashing on Manifolds	1562
<i>Fumin Shen, Chunhua Shen, Qinfeng Shi, Anton van den Hengel, and Zhenmin Tang</i>	
Hash Bit Selection: A Unified Solution for Selection Problems in Hashing	1570
<i>Xianglong Liu, Junfeng He, Bo Lang, and Shih-Fu Chang</i>	
All About VLAD	1578
<i>Relja Arandjelović and Andrew Zisserman</i>	
Binary Code Ranking with Weighted Hamming Distance	1586
<i>Lei Zhang, Yongdong Zhang, Jinhu Tang, Ke Lu, and Qi Tian</i>	
Consensus of k-NNs for Robust Neighborhood Selection on Graph-Based Manifolds	1594
<i>Vittal Premachandran and Ramakrishna Kakarala</i>	
Topical Video Object Discovery from Key Frames by Modeling Word Co-occurrence Prior	1602
<i>Gangqiang Zhao, Junsong Yuan, and Gang Hua</i>	
Query Adaptive Similarity for Large Scale Object Retrieval	1610
<i>Danfeng Qin, Christian Wengert, and Luc Van Gool</i>	
Image Tag Completion via Image-Specific and Tag-Specific Linear Sparse Reconstructions	1618
<i>Zijia Lin, Guiguang Ding, Mingqing Hu, Jianmin Wang, and Xiaojun Ye</i>	
Lp-Norm IDF for Large Scale Image Search	1626
<i>Liang Zheng, Shengjin Wang, Ziqiong Liu, and Qi Tian</i>	
Constraints as Features	1634
<i>Shmuel Asafi and Daniel Cohen-Or</i>	

Learning a Manifold as an Atlas	1642
<i>Nikolaos Pitelis, Chris Russell, and Lourdes Agapito</i>	
Semi-supervised Learning of Feature Hierarchies for Object Detection in a Video	1650
<i>Yang Yang, Guang Shu, and Mubarak Shah</i>	
Fully-Connected CRFs with Non-Parametric Pairwise Potential	1658
<i>Neill D.F. Campbell, Kartic Subr, and Jan Kautz</i>	
Discriminative Sub-categorization	1666
<i>Minh Hoai and Andrew Zisserman</i>	
Whitened Expectation Propagation: Non-Lambertian Shape from Shading and Shadow	1674
<i>Brian Potetz and Mohammadreza Hajiarbabi</i>	
Fast Energy Minimization Using Learned State Filters	1682
<i>Mathieu Guillaumin, Luc Van Gool, and Vittorio Ferrari</i>	
Bilinear Programming for Human Activity Recognition with Unknown MRF Graphs	1690
<i>Zhenhua Wang, Qinfeng Shi, Chunhua Shen, and Anton van den Hengel</i>	
A Higher-Order CRF Model for Road Network Extraction	1698
<i>Jan D. Wegner, Javier A. Montoya-Zegarra, and Konrad Schindler</i>	
Nonlinearly Constrained MRFs: Exploring the Intrinsic Dimensions of Higher-Order Cliques	1706
<i>Yun Zeng, Chaohui Wang, Stefano Soatto, and Shing-Tung Yau</i>	
Fast Trust Region for Segmentation	1714
<i>Lena Gorelick, Frank R. Schmidt, and Yuri Boykov</i>	
Optimal Geometric Fitting under the Truncated L2-Norm	1722
<i>Erik Ask, Olof Enqvist, and Fredrik Kahl</i>	
In Defense of 3D-Label Stereo	1730
<i>Carl Olsson, Johannes Ullén, and Yuri Boykov</i>	
Universality of the Local Marginal Polytope	1738
<i>Daniel Průša and Tomáš Werner</i>	
Continuous Inference in Graphical Models with Polynomial Energies	1744
<i>Mathieu Salzmann</i>	
Towards Efficient and Exact MAP-Inference for Large Scale Discrete Computer Vision Problems via Combinatorial Optimization	1752
<i>Jörg Hendrik Kappes, Markus Speth, Gerhard Reinelt, and Christoph Schnörr</i>	
An Iterated L1 Algorithm for Non-smooth Non-convex Optimization in Computer Vision	1759
<i>Peter Ochs, Alexey Dosovitskiy, Thomas Brox, and Thomas Pock</i>	

A Genetic Algorithm-Based Solver for Very Large Jigsaw Puzzles	1767
<i>Dror Sholomon, Omid David, and Nathan S. Netanyahu</i>	
A Convex Regularize for Reducing Color Artifact in Color Image Recovery	1775
<i>Shunsuke Ono and Isao Yamada</i>	
Kernel Learning for Extrinsic Classification of Manifold Features	1782
<i>Raviteja Vemulapalli, Jaishanker K. Pillai, and Rama Chellappa</i>	

Orals 2C: Detection (+Medical/Curves)

Learning Structured Hough Voting for Joint Object Detection and Occlusion Reasoning	1790
<i>Tao Wang, Xuming He, and Nick Barnes</i>	
Detection Evolution with Multi-order Contextual Co-occurrence	1798
<i>Guang Chen, Yuanyuan Ding, Jing Xiao, and Tony X. Han</i>	
Efficient Large-Scale Structured Learning	1806
<i>Steve Branson, Oscar Beijbom, and Serge Belongie</i>	
Fast, Accurate Detection of 100,000 Object Classes on a Single Machine	1814
<i>Thomas Dean, Mark A. Ruzon, Mark Segal, Jonathon Shlens, Sudheendra Vijayanarasimhan, and Jay Yagnik</i>	
Reconstructing Loopy Curvilinear Structures Using Integer Programming	1822
<i>Engin Türetken, Fethallah Benmansour, Bjoern Andres, Hanspeter Pfister, and Pascal Fua</i>	

Orals 2D: Tracking and Flow

Tracking Sports Players with Context-Conditioned Motion Models	1830
<i>Jingchen Liu, Peter Carr, Robert T. Collins, and Yanxi Liu</i>	
Structure Preserving Object Tracking	1838
<i>Lu Zhang and Laurens van der Maaten</i>	
Multi-target Tracking by Lagrangian Relaxation to Min-cost Network Flow	1846
<i>Asad A. Butt and Robert T. Collins</i>	
Patch Match Filter: Efficient Edge-Aware Filtering Meets Randomized Search for Fast Correspondence Field Estimation	1854
<i>Jiangbo Lu, Hongsheng Yang, Dongbo Min, and Minh N. Do</i>	
Robust Monocular Epipolar Flow Estimation	1862
<i>Koichiro Yamaguchi, David McAllester, and Raquel Urtasun</i>	

Posters 2C: Segmentation and Shape

Deep Learning Shape Priors for Object Segmentation	1870
<i>Fei Chen, Huimin Yu, Roland Hu, and Xunxun Zeng</i>	
PDM-ENLOR: Learning Ensemble of Local PDM-Based Regressions	1878
<i>Yen H. Le, Uday Kurkure, and Ioannis A. Kakadiaris</i>	
Incorporating User Interaction and Topological Constraints within Contour Completion via Discrete Calculus	1886
<i>Jia Xu, Maxwell D. Collins, and Vikas Singh</i>	
Recovering Line-Networks in Images by Junction-Point Processes	1894
<i>Dengfeng Chai, Wolfgang Förstner, and Florent Lafarge</i>	
Image Matting with Local and Nonlocal Smooth Priors	1902
<i>Xiaowu Chen, Dongqing Zou, Steven Zhiying Zhou, Qinqing Zhao, and Ping Tan</i>	
Probabilistic Graphlet Cut: Exploiting Spatial Structure Cue for Weakly Supervised Image Segmentation	1908
<i>Luming Zhang, Mingli Song, Zicheng Liu, Xiao Liu, Jiajun Bu, and Chun Chen</i>	
Towards Fast and Accurate Segmentation	1916
<i>Camillo Jose Taylor</i>	
Discriminative Re-ranking of Diverse Segmentations	1923
<i>Payman Yadollahpour, Dhruv Batra, and Gregory Shakhnarovich</i>	
Robust Region Grouping via Internal Patch Statistics	1931
<i>Xiaobai Liu, Liang Lin, and Alan L. Yuille</i>	
Unsupervised Joint Object Discovery and Segmentation in Internet Images	1939
<i>Michael Rubinstein, Armand Joulin, Johannes Kopf, and Ce Liu</i>	
Ensemble Video Object Cut in Highly Dynamic Scenes	1947
<i>Xiaobo Ren, Tony X. Han, and Zhihai He</i>	
Graph Transduction Learning with Connectivity Constraints with Application to Multiple Foreground Cosegmentation	1955
<i>Tianyang Ma and Longin Jan Latecki</i>	
Top-Down Segmentation of Non-rigid Visual Objects Using Derivative-Based Search on Sparse Manifolds	1963
<i>Jacinto C. Nascimento and Gustavo Carneiro</i>	
A Principled Deep Random Field Model for Image Segmentation	1971
<i>Pushmeet Kohli, Anton Osokin, and Stefanie Jegelka</i>	
Background Modeling Based on Bidirectional Analysis	1979
<i>Atsushi Shimada, Hajime Nagahara, and Rin-ichiro Taniguchi</i>	
Learning for Structured Prediction Using Approximate Subgradient Descent with Working Sets	1987
<i>Aurélien Lucchi, Yunpeng Li, and Pascal Fua</i>	

A Sentence Is Worth a Thousand Pixels	1995
<i>Sanja Fidler, Abhishek Sharma, and Raquel Urtasun</i>	
GRASP Recurring Patterns from a Single View	2003
<i>Jingchen Liu and Yanxi Liu</i>	
Image Segmentation by Cascaded Region Agglomeration	2011
<i>Zhile Ren and Gregory Shakhnarovich</i>	
Augmenting CRFs with Boltzmann Machine Shape Priors for Image Labeling	2019
<i>Andrew Kae, Kihyuk Sohn, Honglak Lee, and Erik Learned-Miller</i>	
Voxel Cloud Connectivity Segmentation - Supervoxels for Point Clouds	2027
<i>Jeremie Papon, Alexey Abramov, Markus Schoeler, and Florentin Wörgötter</i>	
SCALPEL: Segmentation Cascades with Localized Priors and Efficient Learning	2035
<i>David Weiss and Ben Taskar</i>	
Submodular Salient Region Detection	2043
<i>Zhuolin Jiang and Larry S. Davis</i>	
A Video Representation Using Temporal Superpixels	2051
<i>Jason Chang, Donglai Wei, and John W. Fisher III</i>	
Pose from Flow and Flow from Pose	2059
<i>Katerina Fragkiadaki, Han Hu, and Jianbo Shi</i>	
Mesh Based Semantic Modelling for Indoor and Outdoor Scenes	2067
<i>Julien P.C. Valentin, Sunando Sengupta, Jonathan Warrell, Ali Shahrokni, and Philip H.S. Torr</i>	
Weakly-Supervised Dual Clustering for Image Semantic Segmentation	2075
<i>Yang Liu, Jing Liu, Zechao Li, Jinhui Tang, and Hanqing Lu</i>	
Salient Object Detection: A Discriminative Regional Feature Integration Approach	2083
<i>Huaizu Jiang, Jingdong Wang, Zejian Yuan, Yang Wu, Nanning Zheng, and Shipeng Li</i>	
Revisiting Depth Layers from Occlusions	2091
<i>Adarsh Kowdle, Andrew Gallagher, and Tsuhan Chen</i>	
Hierarchical Video Representation with Trajectory Binary Partition Tree	2099
<i>Guillem Palou and Philippe Salembier</i>	
Discriminative Subspace Clustering	2107
<i>Vasileios Zografos, Liam Ellis, and Rudolf Mester</i>	
PISA: Pixelwise Image Saliency by Aggregating Complementary Appearance Contrast Measures with Spatial Priors	2115
<i>Keyang Shi, Keze Wang, Jiangbo Lu, and Liang Lin</i>	

Boundary Detection Benchmarking: Beyond F-Measures	2123
<i>Xiaodi Hou, Alan Yuille, and Christof Koch</i>	
Measures and Meta-Measures for the Supervised Evaluation of Image Segmentation	2131
<i>Jordi Pont-Tuset and Ferran Marques</i>	
Multi-resolution Shape Analysis via Non-Euclidean Wavelets: Applications to Mesh Segmentation and Surface Alignment Problems	2139
<i>Won Hwa Kim, Moo K. Chung, and Vikas Singh</i>	
Robust Estimation of Nonrigid Transformation for Point Set Registration	2147
<i>Jiayi Ma, Ji Zhao, Jinwen Tian, Zhuowen Tu, and Alan L. Yuille</i>	
Efficient Computation of Shortest Path-Concavity for 3D Meshes	2155
<i>Henrik Zimmer, Marcel Campen, and Leif Kobbelt</i>	
Boundary Cues for 3D Object Shape Recovery	2163
<i>Kevin Karsch, Zicheng Liao, Jason Rock, Jonathan T. Barron, and Derek Hoiem</i>	
A Linear Approach to Matching Cuboids in RGBD Images	2171
<i>Hao Jiang and Jianxiang Xiao</i>	

Posters 2D: Motion and Medical Imaging

Blind Deconvolution of Widefield Fluorescence Microscopic Data by Regularization of the Optical Transfer Function (OTF)	2179
<i>Margret Keuper, Thorsten Schmidt, Maja Temerinac-Ott, Jan Padeken, Patrick Heun, Olaf Ronneberger, and Thomas Brox</i>	
Image Understanding from Experts' Eyes by Modeling Perceptual Skill of Diagnostic Reasoning Processes	2187
<i>Rui Li, Pengcheng Shi, and Anne R. Haake</i>	
Adaptive Compressed Tomography Sensing	2195
<i>Oren Barkan, Jonathan Weill, Amir Averbuch, and Shai Dekel</i>	
Classification of Tumor Histology via Morphometric Context	2203
<i>Hang Chang, Alexander Borowsky, Paul Spellman, and Bahram Parvin</i>	
Efficient 3D Endfiring TRUS Prostate Segmentation with Globally Optimized Rotational Symmetry	2211
<i>Jing Yuan, Wu Qiu, Eranga Ukwatta, Martin Rajchl, Xue-Cheng Tai, and Aaron Fenster</i>	
Graph-Based Optimization with Tubularity Markov Tree for 3D Vessel Segmentation	2219
<i>Ning Zhu and Albert C.S. Chung</i>	

Prostate Segmentation in CT Images via Spatial-Constrained Transductive Lasso	2227
<i>Yinghuan Shi, Shu Liao, Yaozong Gao, Daoqiang Zhang, Yang Gao, and Dinggang Shen</i>	
Area Preserving Brain Mapping	2235
<i>Zhengyu Su, Wei Zeng, Rui Shi, Yalin Wang, Jian Sun, and Xianfeng Gu</i>	
Discriminative Brain Effective Connectivity Analysis for Alzheimer's Disease: A Kernel Learning Approach upon Sparse Gaussian Bayesian Network	2243
<i>Luping Zhou, Lei Wang, Lingqiao Liu, Philip Ogunbona, and Dinggang Shen</i>	
Compressible Motion Fields	2251
<i>Giuseppe Ottaviano and Pushmeet Kohli</i>	
Fast Rigid Motion Segmentation via Incrementally-Complex Local Models	2259
<i>Fernando Flores-Mangas and Allan D. Jepson</i>	
Determining Motion Directly from Normal Flows Upon the Use of a Spherical Eye Platform	2267
<i>Tak-Wai Hui and Ronald Chung</i>	
Correspondence-Less Non-rigid Registration of Triangular Surface Meshes	2275
<i>Zsolt Sánta and Zoltan Kato</i>	
Video Editing with Temporal, Spatial and Appearance Consistency	2283
<i>Xiaojie Guo, Xiaochun Cao, Xiaowu Chen, and Yi Ma</i>	
Correlation Filters for Object Alignment	2291
<i>Vishnu Naresh Boddeti, Takeo Kanade, and B.V.K. Vijaya Kumar</i>	
Plane-Based Content Preserving Warps for Video Stabilization	2299
<i>Zihan Zhou, Hailin Jin, and Yi Ma</i>	
Deformable Spatial Pyramid Matching for Fast Dense Correspondences	2307
<i>Jaechul Kim, Ce Liu, Fei Sha, and Kristen Grauman</i>	
The Generalized Laplacian Distance and Its Applications for Visual Matching	2315
<i>Elhanan Elboer, Michael Werman, and Yacov Hel-Or</i>	
Groupwise Registration via Graph Shrinkage on the Image Manifold	2323
<i>Shihui Ying, Guorong Wu, Qian Wang, and Dinggang Shen</i>	
FasT-Match: Fast Affine Template Matching	2331
<i>Simon Korman, Daniel Reichman, Gilad Tsur, and Shai Avidan</i>	
As-Projective-As-Possible Image Stitching with Moving DLT	2339
<i>Julio Zaragoza, Tat-Jun Chin, Michael S. Brown, and David Suter</i>	
Real-Time Model-Based Rigid Object Pose Estimation and Tracking Combining Dense and Sparse Visual Cues	2347
<i>Karl Pauwels, Leonardo Rubio, Javier Díaz, and Eduardo Ros</i>	

Minimum Uncertainty Gap for Robust Visual Tracking	2355
<i>Junseok Kwon and Kyoung Mu Lee</i>	
Part-Based Visual Tracking with Online Latent Structural Learning	2363
<i>Rui Yao, Qinfeng Shi, Chunhua Shen, Yanning Zhang, and Anton van den Hengel</i>	
Least Soft-Threshold Squares Tracking	2371
<i>Dong Wang, Huchuan Lu, and Ming-Hsuan Yang</i>	
Self-Paced Learning for Long-Term Tracking	2379
<i>James Steven Supančič III and Deva Ramanan</i>	
Multi-target Tracking by Rank-1 Tensor Approximation	2387
<i>Xinchu Shi, Haibin Ling, Junling Xing, and Weiming Hu</i>	
Robust Real-Time Tracking of Multiple Objects by Volumetric Mass Densities	2395
<i>Horst Possegger, Sabine Sternig, Thomas Mauthner, Peter M. Roth, and Horst Bischof</i>	
Information Consensus for Distributed Multi-target Tracking	2403
<i>Ahmed T. Kamal, Jay A. Farrell, and Amit K. Roy-Chowdhury</i>	
Online Object Tracking: A Benchmark	2411
<i>Yi Wu, Jongwoo Lim, and Ming-Hsuan Yang</i>	
Learning Compact Binary Codes for Visual Tracking	2419
<i>Xi Li, Chunhua Shen, Anthony Dick, and Anton van den Hengel</i>	
Visual Tracking via Locality Sensitive Histograms	2427
<i>Shengfeng He, Qingxiong Yang, Rynson W.H. Lau, Jiang Wang, and Ming-Hsuan Yang</i>	
Optical Flow Estimation Using Laplacian Mesh Energy	2435
<i>Wenbin Li, Darren Cosker, Matthew Brown, and Rui Tang</i>	
Large Displacement Optical Flow from Nearest Neighbor Fields	2443
<i>Zhuoyuan Chen, Hailin Jin, Zhe Lin, Scott Cohen, and Ying Wu</i>	
A Fully-Connected Layered Model of Foreground and Background Flow	2451
<i>Deqing Sun, Jonas Wulff, Erik B. Sudderth, Hanspeter Pfister, and Michael J. Black</i>	

Orals 3A: Video

Event Retrieval in Large Video Collections with Circulant Temporal Encoding	2459
<i>Jérôme Revaud, Matthijs Douze, Cordelia Schmid, and Hervé Jégou</i>	
Cumulative Attribute Space for Age and Crowd Density Estimation	2467
<i>Ke Chen, Shaogang Gong, Tao Xiang, and Chen Change Loy</i>	
Social Role Discovery in Human Events	2475
<i>Vignesh Ramanathan, Bangpeng Yao, and Li Fei-Fei</i>	
Discriminative Segment Annotation in Weakly Labeled Video	2483
<i>Kevin Tang, Rahul Sukthankar, Jay Yagnik, and Li Fei-Fei</i>	

Context-Aware Modeling and Recognition of Activities in Video	2491
<i>Yingying Zhu, Nandita M. Nayak, and Amit K. Roy-Chowdhury</i>	

Orals 3B: Geometry and Physics (+Medical)

Underwater Camera Calibration Using Wavelength Triangulation	2499
<i>Timothy Yau, Minglun Gong, and Yee-Hong Yang</i>	
Reconstructing Gas Flows Using Light-Path Approximation	2507
<i>Yu Ji, Jinwei Ye, and Jingyi Yu</i>	
Photometric Ambient Occlusion	2515
<i>Daniel Hauage, Scott Wehrwein, Kavita Bala, and Noah Snavely</i>	
What Object Motion Reveals about Shape with Unknown BRDF and Lighting	2523
<i>Manmohan Chandraker, Dikpal Reddy, Yizhou Wang, and Ravi Ramamoorthi</i>	
Hyperbolic Harmonic Mapping for Constrained Brain Surface Registration	2531
<i>Rui Shi, Wei Zeng, Zhengyu Su, Hanna Damasio, Zhonglin Lu, Yalin Wang, Shing-Tung Yau, and Xianfeng Gu</i>	

Posters 3A: Video Analysis

Crossing the Line: Crowd Counting by Integer Programming with Local Features	2539
<i>Zheng Ma and Antoni B. Chan</i>	
Multi-source Multi-scale Counting in Extremely Dense Crowd Images	2547
<i>Haroon Idrees, Imran Saleemi, Cody Seibert, and Mubarak Shah</i>	
Better Exploiting Motion for Better Action Recognition	2555
<i>Mihir Jain, Hervé Jégou, and Patrick Bouthemy</i>	
Detection of Manipulation Action Consequences (MAC)	2563
<i>Yezhou Yang, Cornelia Fermüller, and Yiannis Aloimonos</i>	
Representing Videos Using Mid-level Discriminative Patches	2571
<i>Arpit Jain, Abhinav Gupta, Mikel Rodriguez, and Larry S. Davis</i>	
Modeling Actions through State Changes	2579
<i>Alireza Fathi and James M. Rehg</i>	
Recognizing Activities via Bag of Words for Attribute Dynamics	2587
<i>Weixin Li, Qian Yu, Harpreet Sawhney, and Nuno Vasconcelos</i>	
Sampling Strategies for Real-Time Action Recognition	2595
<i>Feng Shi, Emil Petriu, and Robert Laganière</i>	
Dynamic Scene Classification: Learning Motion Descriptors with Slow Features Analysis	2603
<i>Christian Thériault, Nicolas Thome, and Matthieu Cord</i>	

Online Dominant and Anomalous Behavior Detection in Videos	2611
<i>Mehrsan Javan Roshtkhari and Martin D. Levine</i>	
Augmenting Bag-of-Words: Data-Driven Discovery of Temporal and Structural Information for Activity Recognition	2619
<i>Vinay Bettadapura, Grant Schindler, Thomas Ploetz, and Irfan Essa</i>	
Complex Event Detection via Multi-source Video Attributes	2627
<i>Zhigang Ma, Yi Yang, Zhongwen Xu, Shuicheng Yan, Nicu Sebe, and Alexander G. Hauptmann</i>	
A Thousand Frames in Just a Few Words: Lingual Description of Videos through Latent Topics and Sparse Object Stitching	2634
<i>Pradipto Das, Chenliang Xu, Richard F. Doell, and Jason J. Corso</i>	
Spatiotemporal Deformable Part Models for Action Detection	2642
<i>Yicong Tian, Rahul Sukthankar, and Mubarak Shah</i>	
Poselet Key-Framing: A Model for Human Activity Recognition	2650
<i>Michalis Raptis and Leonid Sigal</i>	
Recognize Human Activities from Partially Observed Videos	2658
<i>Yu Cao, Daniel Barrett, Andrei Barbu, Siddharth Narayanaswamy, Haonan Yu, Aaron Michaux, Yuewei Lin, Sven Dickinson, Jeffrey Mark Siskind, and Song Wang</i>	
Event Recognition in Videos by Learning from Heterogeneous Web Sources	2666
<i>Lin Chen, Lixin Duan, and Dong Xu</i>	
Motionlets: Mid-level 3D Parts for Human Motion Recognition	2674
<i>LiMin Wang, Yu Qiao, and Xiaoou Tang</i>	
Multi-agent Event Detection: Localization and Role Assignment	2682
<i>Suha Kwak, Bohyung Han, and Joon Hee Han</i>	
Cross-View Action Recognition via a Continuous Virtual Path	2690
<i>Zhong Zhang, Chunheng Wang, Baihua Xiao, Wen Zhou, Shuang Liu, and Cunzhao Shi</i>	
Large-Scale Video Summarization Using Web-Image Priors	2698
<i>Aditya Khosla, Raffay Hamid, Chih-Jen Lin, and Neel Sundaresan</i>	
Representing and Discovering Adversarial Team Behaviors Using Player Roles	2706
<i>Patrick Lucey, Alina Bialkowski, Peter Carr, Stuart Morgan, Iain Matthews, and Yaser Sheikh</i>	
Story-Driven Summarization for Egocentric Video	2714
<i>Zheng Lu and Kristen Grauman</i>	
Finding Group Interactions in Social Clutter	2722
<i>Ruonan Li, Parker Porfilio, and Todd Zickler</i>	

First-Person Activity Recognition: What Are They Doing to Me?	2730
<i>Michael S. Ryoo and Larry Matthies</i>	
Joint Sparsity-Based Representation and Analysis of Unconstrained Activities	2738
<i>Raghuraman Gopalan</i>	
Motion Estimation for Self-Driving Cars with a Generalized Camera	2746
<i>Gim Hee Lee, Friedrich Faundorfer, and Marc Pollefeys</i>	

Posters 3B: Features and Contours

Learning Separable Filters	2754
<i>Roberto Rigamonti, Amos Sironi, Vincent Lepetit, and Pascal Fua</i>	
Robust Feature Matching with Alternate Hough and Inverted Hough Transforms	2762
<i>Hsin-Yi Chen, Yen-Yu Lin, and Bing-Yu Chen</i>	
SWIGS: A Swift Guided Sampling Method	2770
<i>Victor Fragoso and Matthew Turk</i>	
Learning Multiple Non-linear Sub-spaces Using K-RBMs	2778
<i>Siddhartha Chandra, Shailesh Kumar, and C.V. Jawahar</i>	
Light Field Distortion Feature for Transparent Object Recognition	2786
<i>Kazuki Maeno, Hajime Nagahara, Atsushi Shimada, and Rin-Ichiro Taniguchi</i>	
From Local Similarity to Global Coding: An Application to Image Classification	2794
<i>Amirreza Shaban, Hamid R. Rabiee, Mehrdad Farajtabar, and Marjan Ghazvininejad</i>	
Joint Spectral Correspondence for Disparate Image Matching	2802
<i>Mayank Bansal and Kostas Daniilidis</i>	
Efficient Color Boundary Detection with Color-Opponent Mechanisms	2810
<i>Kaifu Yang, Shaobing Gao, Chaoyi Li, and Yongjie Li</i>	
Winding Number for Region-Boundary Consistent Salient Contour Extraction	2818
<i>Yansheng Ming, Hongdong Li, and Xuming He</i>	
Supervised Semantic Gradient Extraction Using Linear-Time Optimization	2826
<i>Shulin (Lynn) Yang, Jue Wang, and Linda Shapiro</i>	
Spatio-temporal Depth Cuboid Similarity Feature for Activity Recognition Using Depth Camera	2834
<i>Lu Xia and J.K. Aggarwal</i>	
Sparse Quantization for Patch Description	2842
<i>Xavier Boix, Michael Gygli, Gemma Roig, and Luc Van Gool</i>	
Evaluation of Color STIPs for Human Action Recognition	2850
<i>Ivo Everts, Jan C. van Gemert, and Theo Gevers</i>	

Supervised Kernel Descriptors for Visual Recognition	2858
<i>Peng Wang, Jingdong Wang, Gang Zeng, Weiwei Xu, Hongbin Zha, and Shipeng Li</i>	
Discriminative Color Descriptors	2866
<i>Rahat Khan, Joost van de Weijer, Fahad Shahbaz Khan, Damien Muselet, Christophe Ducottet, and Cecile Barat</i>	
Boosting Binary Keypoint Descriptors	2874
<i>Tomasz Trzcinski, Mario Christoudias, Pascal Fua, and Vincent Lepetit</i>	
Exploring Weak Stabilization for Motion Feature Extraction	2882
<i>Dennis Park, C. Lawrence Zitnick, Deva Ramanan, and Piotr Dollár</i>	
Dense Segmentation-Aware Descriptors	2890
<i>Eduard Trulls, Iasonas Kokkinos, Alberto Sanfeliu, and Francesc Moreno-Noguer</i>	
Keypoints from Symmetries by Wave Propagation	2898
<i>Samuele Salti, Alessandro Lanza, and Luigi Di Stefano</i>	
Graph Matching with Anchor Nodes: A Learning Approach	2906
<i>Nan Hu, Raif M. Rustamov, and Leonidas Guibas</i>	
Dense Non-rigid Point-Matching Using Random Projections	2914
<i>Raffay Hamid, Dennis Decoste, and Chih-Jen Lin</i>	
Deformable Graph Matching	2922
<i>Feng Zhou and Fernando De la Torre</i>	
Scene Coordinate Regression Forests for Camera Relocalization in RGB-D Images	2930
<i>Jamie Shotton, Ben Glocker, Christopher Zach, Shahram Izadi, Antonio Criminisi, and Andrew Fitzgibbon</i>	
K-Means Hashing: An Affinity-Preserving Quantization Method for Learning Binary Compact Codes	2938
<i>Kaiming He, Fang Wen, and Jian Sun</i>	
Optimized Product Quantization for Approximate Nearest Neighbor Search	2946
<i>Tiezheng Ge, Kaiming He, Qifa Ke, and Jian Sun</i>	
A Non-parametric Framework for Document Bleed-through Removal	2954
<i>Róisín Rowley-Brooke, François Pitié, and Anil Kokaram</i>	
Scene Text Recognition Using Part-Based Tree-Structured Character Detection	2961
<i>Cunzhao Shi, Chunheng Wang, Baihua Xiao, Yang Zhang, Song Gao, and Zhong Zhang</i>	
Active Contours with Group Similarity	2969
<i>Xiaowei Zhou, Xiaojie Huang, James S. Duncan, and Weichuan Yu</i>	

Accurate and Robust Registration of Nonrigid Surface Using Hierarchical Statistical Shape Model	2977
<i>Hidekata Hontani, Yuto Tsunekawa, and Yoshihide Sawada</i>	

Orals 3C: Context and Scenes (+ANN)

Spatial Inference Machines	2985
<i>Roman Shapovalov, Dmitry Vetrov, and Pushmeet Kohli</i>	
Hallucinated Humans as the Hidden Context for Labeling 3D Scenes	2993
<i>Yun Jiang, Hema Koppula, and Ashutosh Saxena</i>	
Finding Things: Image Parsing with Regions and Per-Exemplar Detectors	3001
<i>Joseph Tighe and Svetlana Lazebnik</i>	
Bringing Semantics into Focus Using Visual Abstraction	3009
<i>C. Lawrence Zitnick and Devi Parikh</i>	
Cartesian K-Means	3017
<i>Mohammad Norouzi and David J. Fleet</i>	

Orals 3D: Faces, People, and Crowds

Blessing of Dimensionality: High-Dimensional Feature and Its Efficient Compression for Face Verification	3025
<i>Dong Chen, Xudong Cao, Fang Wen, and Jian Sun</i>	
Robust Multi-resolution Pedestrian Detection in Traffic Scenes	3033
<i>Junjie Yan, Xucong Zhang, Zhen Lei, Shengcai Liao, and Stan Z. Li</i>	
Human Pose Estimation Using Body Parts Dependent Joint Regressors	3041
<i>Matthias Dantone, Juergen Gall, Christian Leistner, and Luc Van Gool</i>	
Measuring Crowd Collectiveness	3049
<i>Bolei Zhou, Xiaoou Tang, and Xiaogang Wang</i>	
Lost! Leveraging the Crowd for Probabilistic Visual Self-Localization	3057
<i>Marcus A. Brubaker, Andreas Geiger, and Raquel Urtasun</i>	

Posters 3C: Objects and Scenes

Manhattan Junction Catalogue for Spatial Reasoning of Indoor Scenes	3065
<i>Srikumar Ramalingam, Jaishanker K. Pillai, Arpit Jain, and Yuichi Taguchi</i>	
Tensor-Based High-Order Semantic Relation Transfer for Semantic Scene Segmentation	3073
<i>Heesoo Myeong and Kyoung Mu Lee</i>	
Geometric Context from Videos	3081
<i>S. Hussain Raza, Matthias Grundmann, and Irfan Essa</i>	

It's Not Polite to Point: Describing People with Uncertain Attributes	3089
<i>Amir Sadovnik, Andrew Gallagher, and Tsuhan Chen</i>	
Heterogeneous Visual Features Fusion via Sparse Multimodal Machine	3097
<i>Hua Wang, Feiping Nie, Heng Huang, and Chris Ding</i>	
A Max-Margin Riffled Independence Model for Image Tag Ranking	3103
<i>Tian Lan and Greg Mori</i>	
Weakly Supervised Learning for Attribute Localization in Outdoor Scenes	3111
<i>Shuo Wang, Jungseock Joo, Yizhou Wang, and Song-Chun Zhu</i>	
Scene Parsing by Integrating Function, Geometry and Appearance Models	3119
<i>Yibiao Zhao and Song-Chun Zhu</i>	
Beyond Point Clouds: Scene Understanding by Reasoning Geometry and Physics	3127
<i>Bo Zheng, Yibiao Zhao, Joey C. Yu, Katsushi Ikeuchi, and Song-Chun Zhu</i>	
Label Propagation from ImageNet to 3D Point Clouds	3135
<i>Yan Wang, Rongrong Ji, and Shih-Fu Chang</i>	
Analyzing Semantic Segmentation Using Hybrid Human-Machine CRFs	3143
<i>Roozbeh Mottaghi, Sanja Fidler, Jian Yao, Raquel Urtasun, and Devi Parikh</i>	
Nonparametric Scene Parsing with Adaptive Feature Relevance and Semantic Context	3151
<i>Gautam Singh and Jana Kosecka</i>	
Sketch Tokens: A Learned Mid-level Representation for Contour and Object Detection	3158
<i>Joseph J. Lim, C. Lawrence Zitnick, and Piotr Dollár</i>	
Saliency Detection via Graph-Based Manifold Ranking	3166
<i>Chuan Yang, Lihe Zhang, Huchuan Lu, Xiang Ruan, and Ming-Hsuan Yang</i>	
Maximum Cohesive Grid of Superpixels for Fast Object Localization	3174
<i>Liang Li, Wei Feng, Liang Wan, and Jiawan Zhang</i>	
Accurate Localization of 3D Objects from RGB-D Data Using Segmentation Hypotheses	3182
<i>Byung-soo Kim, Shili Xu, and Silvio Savarese</i>	
Efficient Maximum Appearance Search for Large-Scale Object Detection	3190
<i>Qiang Chen, Zheng Song, Rogerio Feris, Ankur Datta, Liangliang Cao, Zhongyang Huang, and Shuicheng Yan</i>	
Single-Pedestrian Detection Aided by Multi-pedestrian Detection	3198
<i>Wanli Ouyang and Xiaogang Wang</i>	
Robust Object Co-detection	3206
<i>Xin Guo, Dong Liu, Brendan Jou, Mojun Zhu, Anni Cai, and Shih-Fu Chang</i>	

Integrating Grammar and Segmentation for Human Pose Estimation	3214
<i>Brandon Rothrock, Seyoung Park, and Song-Chun Zhu</i>	
Modeling Mutual Visibility Relationship in Pedestrian Detection	3222
<i>Wanli Ouyang, Xingyu Zeng, and Xiaogang Wang</i>	
Learning to Detect Partially Overlapping Instances	3230
<i>Carlos Arteta, Victor Lempitsky, J. Alison Noble, and Andrew Zisserman</i>	
Looking Beyond the Image: Unsupervised Learning for Object Saliency and Detection	3238
<i>Parthipan Siva, Chris Russell, Tao Xiang, and Lourdes Agapito</i>	
Histograms of Sparse Codes for Object Detection	3246
<i>Xiaofeng Ren and Deva Ramanan</i>	
Efficient Detector Adaptation for Object Detection in a Video	3254
<i>Pramod Sharma and Ram Nevatia</i>	
A Lazy Man's Approach to Benchmarking: Semisupervised Classifier Evaluation and Recalibration	3262
<i>Peter Welinder, Max Welling, and Pietro Perona</i>	
Fast Object Detection with Entropy-Driven Evaluation	3270
<i>Raphael Sznitman, Carlos Becker, François Fleuret, and Pascal Fua</i>	
Discriminatively Trained And-Or Tree Models for Object Detection	3278
<i>Xi Song, Tianfu Wu, Yunde Jia, and Song-Chun Zhu</i>	
Occlusion Patterns for Object Class Detection	3286
<i>Bojan Pepikj, Michael Stark, Peter Gehler, and Bernt Schiele</i>	
Bottom-Up Segmentation for Top-Down Detection	3294
<i>Sanja Fidler, Roozbeh Mottaghi, Alan Yuille, and Raquel Urtasun</i>	
Composite Statistical Inference for Semantic Segmentation	3302
<i>Fuxin Li, Joao Carreira, Guy Lebanon, and Cristian Sminchisescu</i>	
Multi-attribute Queries: To Merge or Not to Merge?	3310
<i>Mohammad Rastegari, Ali Diba, Devi Parikh, and Ali Farhadi</i>	
Local Fisher Discriminant Analysis for Pedestrian Re-identification	3318
<i>Sateesh Pedagadi, James Orwell, Sergio Velastin, and Boghos Boghossian</i>	
Explicit Occlusion Modeling for 3D Object Class Representations	3326
<i>M. Zeeshan Zia, Michael Stark, and Konrad Schindler</i>	
Incorporating Structural Alternatives and Sharing into Hierarchy for Multiclass Object Recognition and Detection	3334
<i>Xiaolong Wang, Liang Lin, Lichao Huang, and Shuicheng Yan</i>	
Articulated Pose Estimation Using Discriminative Armlet Classifiers	3342
<i>Georgia Gkioxari, Pablo Arbeláez, Lubomir Bourdev, and Jitendra Malik</i>	

Sparse Output Coding for Large-Scale Visual Recognition	3350
<i>Bin Zhao and Eric P. Xing</i>	
From N to N+1: Multiclass Transfer Incremental Learning	3358
<i>Ilja Kuzborskij, Francesco Orabona, and Barbara Caputo</i>	
What's in a Name? First Names as Facial Attributes	3366
<i>Huizhong Chen, Andrew C. Gallagher, and Bernd Girod</i>	
Kernel Null Space Methods for Novelty Detection	3374
<i>Paul Bodesheim, Alexander Freytag, Erik Rodner, Michael Kemmler, and Joachim Denzler</i>	

Posters 3D: People and Faces

Expressive Visual Text-to-Speech Using Active Appearance Models	3382
<i>Robert Anderson, Björn Stenger, Vincent Wan, and Roberto Cipolla</i>	
Computationally Efficient Regression on a Dependency Graph for Human Pose Estimation	3390
<i>Kota Hara and Rama Chellappa</i>	
Hollywood 3D: Recognizing Actions in 3D Natural Scenes	3398
<i>Simon Hadfield and Richard Bowden</i>	
3D Visual Proxemics: Recognizing Human Interactions in 3D from a Single Image	3406
<i>Ishani Chakraborty, Hui Cheng, and Omar Javed</i>	
Decoding Children's Social Behavior	3414
<i>James M. Rehg, Gregory D. Abowd, Agata Rozga, Mario Romero, Mark A. Clements, Stan Sclaroff, Irfan Essa, Opal Y. Ousley, Yin Li, Chanh Kim, Hrishikesh Rao, Jonathan C. Kim, Liliana Lo Presti, Jianming Zhang, Denis Lantsman, Jonathan Bidwell, and Zhefan Ye</i>	
Capturing Complex Spatio-temporal Relations among Facial Muscles for Facial Expression Recognition	3422
<i>Ziheng Wang, Shangfei Wang, and Qiang Ji</i>	
Detecting Pulse from Head Motions in Video	3430
<i>Guha Balakrishnan, Fredo Durand, and John Guttag</i>	
Towards Contactless, Low-Cost and Accurate 3D Fingerprint Identification	3438
<i>Ajay Kumar and Cyril Kwong</i>	
Robust Discriminative Response Map Fitting with Constrained Local Models	3444
<i>Akshay Asthana, Stefanos Zafeiriou, Shiyang Cheng, and Maja Pantic</i>	
Facial Feature Tracking Under Varying Facial Expressions and Face Poses Based on Restricted Boltzmann Machines	3452
<i>Yue Wu, Zuoguan Wang, and Qiang Ji</i>	

Detecting and Aligning Faces by Image Retrieval	3460
<i>Xiaohui Shen, Zhe Lin, Jonathan Brandt, and Ying Wu</i>	
Learning SURF Cascade for Fast and Accurate Object Detection	3468
<i>Jianguo Li and Yimin Zhang</i>	
Deep Convolutional Network Cascade for Facial Point Detection	3476
<i>Yi Sun, Xiaogang Wang, and Xiaoou Tang</i>	
Exemplar-Based Face Parsing	3484
<i>Brandon M. Smith, Li Zhang, Jonathan Brandt, Zhe Lin, and Jianchao Yang</i>	
Graph-Laplacian PCA: Closed-Form Solution and Robustness	3492
<i>Bo Jiang, Chris Ding, Bio Luo, and Jin Tang</i>	
Probabilistic Elastic Matching for Pose Variant Face Verification	3499
<i>Haoxiang Li, Gang Hua, Zhe Lin, Jonathan Brandt, and Jianchao Yang</i>	
Constrained Clustering and Its Application to Face Clustering in Videos	3507
<i>Baoyuan Wu, Yifan Zhang, Bao-Gang Hu, and Qiang Ji</i>	
Selective Transfer Machine for Personalized Facial Action Unit Detection	3515
<i>Wen-Sheng Chu, Fernando De La Torre, and Jeffery F. Cohn</i>	
The SVM-Minus Similarity Score for Video Face Recognition	3523
<i>Lior Wolf and Noga Levy</i>	
Face Recognition in Movie Trailers via Mean Sequence Sparse Representation-Based Classification	3531
<i>Enrique G. Ortiz, Alan Wright, and Mubarak Shah</i>	
Towards Pose Robust Face Recognition	3539
<i>Dong Yi, Zhen Lei, and Stan Z. Li</i>	
Single-Sample Face Recognition with Image Corruption and Misalignment via Sparse Illumination Transfer	3546
<i>Liansheng Zhuang, Allen Y. Yang, Zihan Zhou, S. Shankar Sastry, and Yi Ma</i>	
Fusing Robust Face Region Descriptors via Multiple Metric Learning for Face Recognition in the Wild	3554
<i>Zhen Cui, Wen Li, Dong Xu, Shiguang Shan, and Xilin Chen</i>	
Action Recognition by Hierarchical Sequence Summarization	3562
<i>Yale Song, Louis-Philippe Morency, and Randall Davis</i>	
Pixel-Level Hand Detection in Ego-centric Videos	3570
<i>Cheng Li and Kris M. Kitani</i>	
Human Pose Estimation Using a Joint Pixel-wise and Part-wise Formulation	3578
<i>Lubor Ladický, Philip H.S. Torr, and Andrew Zisserman</i>	
Unsupervised Saliency Learning for Person Re-identification	3586
<i>Rui Zhao, Wanli Ouyang, and Xiaogang Wang</i>	

Locally Aligned Feature Transforms across Views	3594
<i>Wei Li and Xiaogang Wang</i>	
Semi-supervised Learning with Constraints for Person Identification in Multimedia Data	3602
<i>Martin Bäuml, Makarand Tapaswi, and Rainer Stiefelhagen</i>	
Learning Locally-Adaptive Decision Functions for Person Verification	3610
<i>Zhen Li, Shiyu Chang, Feng Liang, Thomas S. Huang, Liangliang Cao, and John R. Smith</i>	
3D Pictorial Structures for Multiple View Articulated Pose Estimation	3618
<i>Magnus Burenius, Josephine Sullivan, and Stefan Carlsson</i>	
Pedestrian Detection with Unsupervised Multi-stage Feature Learning	3626
<i>Pierre Sermanet, Koray Kavukcuoglu, Soumith Chintala, and Yann Lecun</i>	
A Joint Model for 2D and 3D Pose Estimation from a Single Image	3634
<i>Edgar Simo-Serra, Ariadna Quattoni, Carme Torras, and Francesc Moreno-Noguer</i>	
Unconstrained Monocular 3D Human Pose Estimation by Action Detection and Cross-Modality Regression Forest	3642
<i>Tsz-Ho Yu, Tae-Kyun Kim, and Roberto Cipolla</i>	
Hypergraphs for Joint Multi-view Reconstruction and Multi-object Tracking	3650
<i>Martin Hofmann, Daniel Wolf, and Gerhard Rigoll</i>	
Tracking People and Their Objects	3658
<i>Tobias Baumgartner, Dennis Mitzel, and Bastian Leibe</i>	
Seeking the Strongest Rigid Detector	3666
<i>Rodrigo Benenson, Markus Mathias, Tinne Tuytelaars, and Luc Van Gool</i>	
MODEC: Multimodal Decomposable Models for Human Pose Estimation	3674
<i>Ben Sapp and Ben Taskar</i>	
Detection- and Trajectory-Level Exclusion in Multiple Object Tracking	3682
<i>Anton Milan, Konrad Schindler, and Stefan Roth</i>	
Optimized Pedestrian Detection for Multiple and Occluded People	3690
<i>Sitapa Rujikietgumjorn and Robert T. Collins</i>	
Long-Term Occupancy Analysis Using Graph-Based Optimisation in Thermal Imagery	3698
<i>Rikke Gade, Anders Jørgensen, and Thomas B. Moeslund</i>	
Detecting and Naming Actors in Movies Using Generative Appearance Models	3706
<i>Vineet Gandhi and Remi Ronfard</i>	
Harry Potter's Marauder's Map: Localizing and Tracking Multiple Persons-of-Interest by Nonnegative Discretization	3714
<i>Shoou-I Yu, Yi Yang, and Alexander Hauptmann</i>	

Improving an Object Detector and Extracting Regions Using Superpixels	3721
<i>Guang Shu, Afshin Dehghan, and Mubarak Shah</i>	
Tracking Human Pose by Tracking Symmetric Parts	3728
<i>Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh</i>	

Author Index