2015 IEEE International Conference on Computer Vision (ICCV 2015)

Santiago, Chile
7-13 December 2015

Pages 1-791

IEEE Catalog Number: CFP15198-POD
ISBN: 978-1-4673-8392-9

Copyright © 2015 by the Institute of Electrical and Electronic Engineers, Inc
All Rights Reserved

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

For other copying, reprint or republication permission, write to IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854. All rights reserved.

******This publication is a representation of what appears in the IEEE Digital Libraries. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number:	CFP15198-POD
ISBN (Print-On-Demand):	978-1-4673-8392-9
ISBN (Online):	978-1-4673-8391-2
ISSN:	1550-5499

Additional Copies of This Publication Are Available From:

Curran Associates, Inc
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: (845) 758-0400
Fax: (845) 758-2633
E-mail: curran@proceedings.com
Web: www.proceedings.com

2015 IEEE International Conference on Computer Vision

ICCV 2015

Table of Contents

Message from the General Chair and Program Chairs.....xxxvi

Organizing Committee.....xxxviii

ICCV 2015 Area Chairs.....xl

ICCV 2015 Reviewers.....xli

Sponsors.....xlii

O-1A: Vision and Language

Ask Your Neurons: A Neural-Based Approach to Answering Questions about Images 1
Mateusz Malinowski, Marcus Rohrbach, and Mario Fritz

Segment-Phrase Table for Semantic Segmentation, Visual Entailment and Paraphrasing 10
Hamid Izadinia, Fereshteh Sadeghi, Santosh K. Divvala, Hannaneh Hajishirzi, Yejin Choi, and Ali Farhadi

Aligning Books and Movies: Towards Story-Like Visual Explanations by Watching Movies and Reading Books 19
Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler

Learning Query and Image Similarities with Ranking Canonical Correlation Analysis 28
Ting Yao, Tao Mei, and Chong-Wah Ngo

P-1A: Recognition, Low-Level Vision, and Biomedical Image Analysis

Learning to See by Moving 37
Pulkit Agrawal, João Carreira, and Jitendra Malik

Object Detection Using Generalization and Efficiency Balanced Co-Occurrence Features 46
Haoyu Ren and Ze-Nian Li

Mining And-Or Graphs for Graph Matching and Object Discovery 55
Quanshi Zhang, Ying Nian Wu, and Song-Chun Zhu

Pose Induction for Novel Object Categories 64
Shubham Tulsiani, João Carreira, and Jitendra Malik

Dynamic Texture Recognition via Orthogonal Tensor Dictionary Learning 73
Yuhui Quan, Yan Huang, and Hui Ji

Convolutional Channel Features 82
Bin Yang, Junjie Yan, Zhen Lei, and Stan Z. Li

Local Convolutional Features with Unsupervised Training for Image Retrieval 91
Mattis Paulin, Matthijs Douze, Zaid Harchaoui, Julien Mairal, Florent Perronin, and Cordelia Schmid

RIDE: Reversal Invariant Descriptor Enhancement 100
Lingxi Xie, Jingdong Wang, Weiyao Lin, Bo Zhang, and Qi Tian

Discrete Tabu Search for Graph Matching 109
Kamil Adamczewski, Yumin Suh, and Kyoung Mu Lee

Discriminative Learning of Deep Convolutional Feature Point Descriptors 118
Edgar Simo-Serra, Eduard Trulls, Luis Ferraz, Iasonas Kokkinos, Pascal Fua, and Francesc Moreno-Noguer

Amodal Completion and Size Constancy in Natural Scenes 127
Abhishek Kar, Shubham Tulsiani, João Carreira, and Jitendra Malik

Learning Where to Position Parts in 3D 136
Marco Pedersoli and Tinne Tuytelaars

Query Adaptive Similarity Measure for RGB-D Object Recognition 145
Yanhua Cheng, Rui Cai, Chi Zhang, Zhiwei Li, Xin Zhao, Kaiqi Huang, and Yong Rui

Listening with Your Eyes: Towards a Practical Visual Speech Recognition System Using Deep Boltzmann Machines 154
Chao Sui, Mohammed Bennamoun, and Roberto Togneri

Cluster-Based Point Set Saliency 163
Flora Ponjou Tasse, Jiri Kosinka, and Neil Dodgson

A Comprehensive Multi-Illuminant Dataset for Benchmarking of the Intrinsic Image Algorithms 172
Shida Beigpour, Andreas Kolb, and Sven Kunz

PatchMatch-Based Automatic Lattice Detection for Near-Regular Textures 181
Siying Liu, Tian-Tsong Ng, Kalyan Sunkavalli, Minh N. Do, Eli Shechtman, and Nathan Carr

A Data-Driven Metric for Comprehensive Evaluation of Saliency Models 190
Jia Li, Changqun Xia, Yafei Song, Shu Fang, and Xiaowu Chen

A Matrix Decomposition Perspective to Multiple Graph Matching 199
Junchi Yan, Hongteng Xu, Hongyuan Zha, Xiaokang Yang, Huanxi Liu, and Stephen Chu

Fast and Effective L0 Gradient Minimization by Region Fusion 208
Rang M. H. Nguyen and Michael S. Brown

Generic Promotion of Diffusion-Based Salient Object Detection 217
Peng Jiang, Nuno Vasconcelos, and Jingliang Peng

Nighttime Haze Removal with Glow and Multiple Light Colors 226
Yu Li, Robby T. Tan, and Michael S. Brown

Conformal and Low-Rank Sparse Representation for Image Restoration 235
Jianwei Li, Xiaowu Chen, Dongqing Zou, Bo Gao, and Wei Teng

Patch Group Based Nonlocal Self-Similarity Prior Learning for Image Denoising 244
Jun Xu, Lei Zhang, Wangmeng Zuo, David Zhang, and Xiangchu Feng

Automatic Thumbnail Generation Based on Visual Representativeness and Foreground Recognizability 253
Jingwei Huang, Huarong Chen, Bin Wang, and Stephen Lin

SALICON: Reducing the Semantic Gap in Saliency Prediction by Adapting Deep Neural Networks 262
Xun Huang, Chengyao Shen, Xavier Boix, and Qi Zhao

A Novel Sparsity Measure for Tensor Recovery 271
Qian Zhao, Deyu Meng, Xu Kong, Qi Xie, Wenfei Cao, Yao Wang, and Zongben Xu

Oriented Object Proposals 280
Shengfeng He and Rynson W. H. Lau

Learning Nonlinear Spectral Filters for Color Image Reconstruction 289
Michael Moeller, Julia Diebold, Guy Gilboa, and Daniel Cremers

Beyond White: Ground Truth Colors for Color Constancy Correction 298
Dongliang Cheng, Brian Price, Scott Cohen, and Michael S. Brown

RGB-Guided Hyperspectral Image Upsampling 307
Hyeokhyen Kwon and Yu-Wing Tai

Projection onto the Manifold of Elongated Structures for Accurate Extraction 316
Amos Sironi, Vincent Lepetit, and Pascal Fua

Naive Bayes Super-Resolution Forest 325
Jordi Salvador and Eduardo Pérez-Pellitero

POP Image Fusion — Derivative Domain Image Fusion without Reintegration 334
Graham D. Finlayson and Alex E. Hayes

Adaptive Spatial-Spectral Dictionary Learning for Hyperspectral Image Denoising 343
Ying Fu, Antony Lam, Imari Sato, and Yoichi Sato

Fully Connected Guided Image Filtering 352
Longquan Dai, Mengke Yuan, Feihu Zhang, and Xiaopeng Zhang

Segment Graph Based Image Filtering: Fast Structure-Preserving Smoothing 361
Feihu Zhang, Longquan Dai, Shiming Xiang, and Xiaopeng Zhang

Deep Networks for Image Super-Resolution with Sparse Prior 370
Zhaowen Wang, Ding Liu, Jianchao Yang, Wei Han, and Thomas Huang

Convolutional Color Constancy 379
Jonathan T. Barron

Learning Ordinal Relationships for Mid-Level Vision 388
Daniel Zoran, Phillip Isola, Dilip Krishnan, and William T. Freeman

Thin Structure Estimation with Curvature Regularization 397
Dmitrii Marin, Yuchen Zhong, Maria Drangova, and Yuri Boykov

HARF: Hierarchy-Associated Rich Features for Salient Object Detection 406
Wenbin Zou and Nikos Komodakis

Deep Colorization 415
Zezhou Cheng, Qingxiong Yang, and Bin Sheng

Image Matting with KL-Divergence Based Sparse Sampling 424
Levent Karacan, Aykut Erdem, and Erkut Erdem

Intrinsic Decomposition of Image Sequences from Local Temporal Variations 433
Pierre-Yves Laffont and Jean-Charles Bazin

Low-Rank Tensor Approximation with Laplacian Scale Mixture Modeling for Multiframe Image Denoising 442
Weisheng Dong, Guangyu Li, Guangming Shi, Xin Li, and Yi Ma

Learning Parametric Distributions for Image Super-Resolution: Where Patch Matching Meets Sparse Coding 450
Yongbo Li, Weisheng Dong, Guangming Shi, and Xuemei Xie

Improving Image Restoration with Soft-Rounding 459
Xing Mei, Honggang Qi, Bao-Gang Hu, and Siwei Lyu

See the Difference: Direct Pre-Image Reconstruction and Pose Estimation by Differentiating HOG 468
Wei-Chen Chiu and Mario Fritz

An Efficient Statistical Method for Image Noise Level Estimation 477
Guangyong Chen, Fengyuan Zhu, and Pheng Ann Heng

Contour Detection and Characterization for Asynchronous Event Sensors 486
Francisco Barranco, Ching L. Teo, Cornelia Fermüller, and Yiannis Aloimonos

Class-Specific Image Deblurring 495
Saeed Anwar, Cong Phuoc Huynh, and Fatih Porikli

High-for-Low and Low-for-High: Efficient Boundary Detection from Deep Object Features and Its Applications to High-Level Vision 504
Gedas Bertasius, Jianbo Shi, and Lorenzo Torresani

Variational Depth Superresolution Using Example-Based Edge Representations 513
David Ferstl, Matthias Rüther, and Horst Bischof

Conditioned Regression Models for Non-blind Single Image Super-Resolution 522
Gernot Riegler, Samuel Schulter, Matthias Rüther, and Horst Bischof

Video Super-Resolution via Deep Draft-Ensemble Learning 531
Renjie Liao, Xin Tao, Ruiyu Li, Ziyang Ma, and Jiaya Jia

Pan-Sharpening with a Hyper-Laplacian Penalty 540
Yiyong Jiang, Xinghao Ding, Delu Zeng, Yue Huang, and John Paisley

Video Restoration Against Yin-Yang Phasing 549
Xiaolin Wu, Zhenhao Li, and Xiaowei Deng

Rolling Shutter Super-Resolution 558
Abhijith Punnappurath, Vijay Rengarajan, and A. N. Rajagopalan

Learning Large-Scale Automatic Image Colorization 567
Aditya Deshpande, Jason Rock, and David Forsyth

Compression Artifacts Reduction by a Deep Convolutional Network 576
Chao Dong, Yubin Deng, Chen Change Loy, and Xiaoou Tang

Multiple-Hypothesis Affine Region Estimation with Anisotropic LoG Filters 585
Takahiro Hasegawa, Mitsuru Ambai, Kohta Ishikawa, Gou Koutaki, Yuji Yamauchi, Takayoshi Yamashita, and Hironobu Fujiyoshi

A Self-Paced Multiple-Instance Learning Framework for Co-Saliency Detection 594
Dingwen Zhang, Deyu Meng, Chao Li, Lu Jiang, Qian Zhao, and Junwei Han

External Patch Prior Guided Internal Clustering for Image Denoising 603
Fei Chen, Lei Zhang, and Huimin Yu

Self-Calibration of Optical Lenses 612
Michael Hirsch and Bernhard Schölkopf

Illumination Robust Color Naming via Label Propagation 621
Yuanliu Liu, Zejian Yuan, Badong Chen, Jianru Xue, and Nanning Zheng

Unsupervised Cross-Modal Synthesis of Subject-Specific Scans 630
Raviteja Vemulapalli, Hien Van Nguyen, and Shaohua Kevin Zhou

Learning to Boost Filamentary Structure Segmentation 639
Lin Gu and Li Cheng

Weakly-Supervised Structured Output Learning with Flexible and Latent Graphs Using High-Order Loss Functions 648
Gustavo Carneiro, Tingying Peng, Christine Bayer, and Nassir Navab

Efficient Classifier Training to Minimize False Merges in Electron Microscopy Segmentation 657
Toufiq Parag, Dan C. Ciresan, and Alessandro Giusti

On Statistical Analysis of Neuroimages with Imperfect Registration 666
Won Hwa Kim, Sathya N. Ravi, Sterling C. Johnson, Ozioma C. Okonkwo, and Vikas Singh

S-1A: Plenary Session

Convex Optimization with Abstract Linear Operators 675
Steven Diamond and Stephen Boyd

P-1B: Recognition and 3D Computer Vision I

Building Dynamic Cloud Maps from the Ground Up 684
Calvin Murdock, Nathan Jacobs, and Robert Pless

A Versatile Learning-Based 3D Temporal Tracker: Scalable, Robust, Online 693
David Joseph Tan, Federico Tombari, Slobodan Ilic, and Nassir Navab

Realtime Edge-Based Visual Odometry for a Monocular Camera 702
Juan José Tarrio and Sol Pedre

Fill and Transfer: A Simple Physics-Based Approach for Containability Reasoning 711
Lap-Fai Yu, Noah Duncan, and Sai-Kit Yeung

On Linear Structure from Motion for Light Field Cameras 720
Ole Johannsen, Antonin Sulc, and Bastian Goldluecke

3D Object Reconstruction from Hand-Object Interactions 729
Dimitrios Tzionas and Juergen Gall

Minimal Solvers for 3D Geometry from Satellite Imagery 738
Enliang Zheng, Ke Wang, Enrique Dunn, and Jan-Michael Frahm

An Efficient Minimal Solution for Multi-camera Motion 747
Jonathan Ventura, Clemens Arth, and Vincent Lepetit

Learning Shape, Motion and Elastic Models in Force Space 756
Antonio Agudo and Francesc Moreno-Noguer

A Versatile Scene Model with Differentiable Visibility Applied to Generative Pose Estimation 765
Helge Rhodin, Nadia Robertini, Christian Richardt, Hans-Peter Seidel, and Christian Theobalt

Semantic Pose Using Deep Networks Trained on Synthetic RGB-D 774
Jeremie Papon and Markus Schoeler

Exploiting High Level Scene Cues in Stereo Reconstruction 783
Simon Hadfield and Richard Bowden

Point Triangulation through Polyhedron Collapse Using the ℓ_∞ Norm 792
Simon Donné, Bart Goossens, and Wilfried Philips

Optimizing the Viewing Graph for Structure-from-Motion 801
Chris Sweeney, Torsten Sattler, Tobias Höllerer, Matthew Turk, and Marc Pollefeys

Intrinsic Scene Decomposition from RGB-D Images 810
Mohammed Hachama, Bernard Ghanem, and Peter Wonka

3D Hand Pose Estimation Using Randomized Decision Forest with Segmentation Index Points 819
Peiyi Li, Haibin Ling, Xi Li, and Chunyuan Liao

Accurate Camera Calibration Robust to Defocus Using a Smartphone 828
Hyowon Ha, Yunsu Bok, Kyungdon Joo, Jiyoung Jung, and In So Kweon

High Quality Structure from Small Motion for Rolling Shutter Cameras 837
Sunghoon Im, Hyowon Ha, Gyeongmin Choe, Hae-Gon Jeon, Kyungdon Joo, and In So Kweon

Photogeometric Scene Flow for High-Detail Dynamic 3D Reconstruction 846
Paulo F. U. Gotardo, Tomas Simon, Yaser Sheikh, and Iain Matthews

Blur-Aware Disparity Estimation from Defocus Stereo Images 855
Ching-Hui Chen, Hui Zhou, and Timo Ahonen

Global Structure-from-Motion by Similarity Averaging 864
Zhaopeng Cui and Ping Tan

Massively Parallel Multiview Stereopsis by Surface Normal Diffusion 873
Silvano Galliani, Katrin Lasinger, and Konrad Schindler

Variational PatchMatch MultiView Reconstruction and Refinement 882
Philipp Heise, Brian Jensen, Sebastian Klose, and Alois Knoll

As-Rigid-as-Possible Volumetric Shape-from-Template 891
Shaifali Parashar, Daniel Pizarro, Adrien Bartoli, and Toby Collins

General Dynamic Scene Reconstruction from Multiple View Video 900
Armin Mustafa, Hansung Kim, Jean-Yves Guillemaut, and Adrian Hilton

The Joint Image Handbook 909
Matthew Trager, Martial Hebert, and Jean Ponce

Direct, Dense, and Deformable: Template-Based Non-rigid 3D Reconstruction from RGB Video 918
Rui Yu, Chris Russell, Neill D. F. Campbell, and Lourdes Agapito

Single Image Pop-Up from Discriminatively Learned Parts 927
Menglong Zhu, Xiaowei Zhou, and Kostas Daniilidis

Learning Informative Edge Maps for Indoor Scene Layout Prediction 936
Arun Mallya and Svetlana Lazebnik

Multi-view Convolutional Neural Networks for 3D Shape Recognition 945
Hang Su, Subhransu Maji, Evangelos Kalogerakis, and Erik Learned-Miller

Learning Analysis-by-Synthesis for 6D Pose Estimation in RGB-D Images 954
Alexander Krull, Eric Brachmann, Frank Michel, Michael Ying Yang, Stefan Gumhold, and Carsten Rother

3D Surface Profilometry Using Phase Shifting of De Bruijn Pattern 963
Matea Đonlić, Tomislav Petković, and Tomislav Pribanić

A Deep Visual Correspondence Embedding Model for Stereo Matching Costs 972
Zhuoyuan Chen, Xun Sun, Liang Wang, Yinan Yu, and Chang Huang

Learning Concept Embeddings with Combined Human-Machine Expertise 981
Michael J. Wilber, Iljung S. Kwak, David Kriegman, and Serge Belongie

Deep Multi-patch Aggregation Network for Image Style, Aesthetics, and Quality Estimation 990
Xin Lu, Zhe Lin, Xiaohui Shen, Radomír Měch, and James Z. Wang

Towards Computational Baby Learning: A Weakly-Supervised Approach for Object Detection 999
Xiaodan Liang, Si Liu, Yunchao Wei, Luoqi Liu, Liang Lin, and Shuicheng Yan

Improving Image Classification with Location Context 1008
Kevin Tang, Manohar Paluri, Li Fei-Fei, Rob Fergus, and Lubomir Bourdev

HICO: A Benchmark for Recognizing Human-Object Interactions in Images 1017
Yu-Wei Chao, Zhan Wang, Yugeng He, Jiaxuan Wang, and Jia Deng

Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification 1026
Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun

Continuous Pose Estimation with a Spatial Ensemble of Fisher Regressors 1035
Michele Fenzi, Laura Leal-Taixé, Jörn Ostermann, and Tinne Tuytelaars

Adaptive Hashing for Fast Similarity Search 1044
Fatih Cakir and Stan Sclaroff

Single Image 3D without a Single 3D Image 1053
David F. Fouhey, Wajahat Hussain, Abhinav Gupta, and Martial Hebert

Cross-Domain Image Retrieval with a Dual Attribute-Aware Ranking Network 1062
Junshi Huang, Rogerio Feris, Qiang Chen, and Shuicheng Yan

Attribute-Graph: A Graph Based Approach to Image Ranking 1071
Nikita Prabhu and R. Venkatesh Babu

Contextual Action Recognition with R*CNN 1080
Georgia Gkioxari, Ross Girshick, and Jitendra Malik

What Makes an Object Memorable? 1089
Rachit Dubey, Joshua Peterson, Aditya Khosla, Ming-Hsuan Yang, and Bernard Ghanem

kNN Hashing with Factorized Neighborhood Representation 1098
Kun Ding, Chunlei Huo, Bin Fan, and Chunhong Pan

Multi-View Complementary Hash Tables for Nearest Neighbor Search 1107
Xianglong Liu, Lei Huang, Cheng Deng, Jiwen Lu, and Bo Lang

Scalable Person Re-identification: A Benchmark 1116
Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian

MMSS: Multi-modal Sharable and Specific Feature Learning for RGB-D Object Recognition 1125
Anran Wang, Jianfei Cai, Jiwen Lu, and Tat-Jen Cham

Object Detection via a Multi-region and Semantic Segmentation-Aware CNN Model 1134
Spyros Gidaris and Nikos Komodakis

Neural Activation Constellations: Unsupervised Part Model Discovery with Convolutional Networks 1143
Marcel Simon and Erik Rodner

Cascaded Sparse Spatial Bins for Efficient and Effective Generic Object Detection 1152
David Novotny and Jiri Matas

Probabilistic Label Relation Graphs with Ising Models 1161
Nan Ding, Jia Deng, Kevin P. Murphy, and Hartmut Neven

Predicting Good Features for Image Geo-Localization Using Per-Bundle VLAD 1170
Hyo Jin Kim, Enrique Dunn, and Jan-Michael Frahm

Task-Driven Feature Pooling for Image Classification 1179
Guo-Sen Xie, Xu-Yao Zhang, Xiangbo Shu, Shuicheng Yan, and Cheng-Lin Liu

Cutting Edge: Soft Correspondences in Multimodal Scene Parsing 1188
Sarah Taghavi Namin, Mohammad Najafi, Mathieu Salzmann, and Lars Petersson

One Shot Learning via Compositions of Meaningful Patches 1197
Alex Wong and Alan Yuille

FASText: Efficient Unconstrained Scene Text Detector 1206
Michal Bušta, Lukáš Neumann, and Jiří Matas

Multi-scale Recognition with DAG-CNNs 1215
Songfan Yang and Deva Ramanan

Relaxed Multiple-Instance SVM with Application to Object Discovery 1224
Xinggang Wang, Zhuotun Zhu, Cong Yao, and Xiang Bai

Im2Calories: Towards an Automated Mobile Vision Food Diary 1233
Austin Myers, Nick Johnston, Vivek Rathod, Anoop Korattikara, Alex Gorban, Nathan Silberman, Sergio Guadarrama, George Papandreou, Jonathan Huang, and Kevin Murphy

LEWIS: Latent Embeddings for Word Images and Their Semantics 1242
Albert Gordo, Jon Almazán, Naila Murray, and Florent Perronin

Per-Sample Kernel Adaptation for Visual Recognition and Grouping 1251
Borislav Antic and Björn Ommer

Fine-Grained Change Detection of Misaligned Scenes with Varied Illuminations 1260
Wei Feng, Fei-Peng Tian, Qian Zhang, Nan Zhang, Liang Wan, and Jizhou Sun

Aggregating Local Deep Features for Image Retrieval 1269
Artem Babenko Yandex and Victor Lempitsky

Learning Deep Object Detectors from 3D Models 1278
Xingchao Peng, Baochen Sun, Karim Ali, and Kate Saenko

Harvesting Discriminative Meta Objects with Deep CNN Features for Scene Classification 1287
Ruobing Wu, Baoyuan Wang, Wenping Wang, and Yizhou Yu

Scalable Nonlinear Embeddings for Semantic Category-Based Image Retrieval 1296
Gaurav Sharma and Bernt Schiele

Person Re-Identification Ranking Optimisation by Discriminant Context Information Analysis 1305
Jorge García, Niki Martinel, Christian Micheloni, and Alfredo Gardel

Unsupervised Generation of a View Point Annotated Car Dataset from Videos 1314
Nima Sedaghat and Thomas Brox

O-1B: 3D Vision

Structured Indoor Modeling 1323
Satoshi Ikehata, Hang Yang, and Yasutaka Furukawa

3D Time-Lapse Reconstruction from Internet Photos 1332
Ricardo Martin-Brualla, David Gallup, and Steven M. Seitz

Global, Dense Multiscale Reconstruction for a Billion Points 1341
Benjamin Ummenhofer and Thomas Brox

On the Visibility of Point Clouds 1350
Sagi Katz and Ayellet Tal

O-2A: Segmentation, Edges and Saliency

Weakly Supervised Graph Based Semantic Segmentation by Learning Communities of Image-Parts 1359
Niloufar Pourian, S. Karthikeyan, and B. S. Manjunath

Piecewise Flat Embedding for Image Segmentation 1368
Yizhou Yu, Chaowei Fang, and Zicheng Liao

Semantic Image Segmentation via Deep Parsing Network 1377
Ziwei Liu, Xiaoxiao Li, Ping Luo, Chen-Change Loy, and Xiaoou Tang

Human Parsing with Contextualized Convolutional Neural Network 1386
Xiaodan Liang, Chunyan Xu, Xiaohui Shen, Jianchao Yang, Si Liu, Jinhui Tang, Liang Lin, and Shuicheng Yan

Holistically-Nested Edge Detection 1395
Saining Xie and Zhuowen Tu

Minimum Barrier Salient Object Detection at 80 FPS 1404
Jianming Zhang, Stan Sclaroff, Zhe Lin, Xiaohui Shen, Brian Price, and Radomír Měch

O-2B: Learning Representations and Attributes

Learning Image Representations Tied to Ego-Motion 1413
Dinesh Jayaraman and Kristen Grauman

Unsupervised Visual Representation Learning by Context Prediction 1422
Carl Doersch, Abhinav Gupta, and Alexei A. Efros

Webly Supervised Learning of Convolutional Networks 1431
Xinlei Chen and Abhinav Gupta

Fast R-CNN 1440
Ross Girshick

Bilinear CNN Models for Fine-Grained Visual Recognition 1449
Tsung-Yu Lin, Aruni RoyChowdhury, and Subhransu Maji

Discovering the Spatial Extent of Relative Attributes 1458
Fanyi Xiao and Yong Jae Lee

O-2C: Statistical Methods and Learning

Deep Neural Decision Forests 1467
Peter Kontschieder, Madalina Fiterau, Antonio Criminisi, and Samuel Rota Bulò

Deep Fried Convnets 1476
Zichao Yang, Marcin Moczulski, Misha Denil, Nando de Freitas, Alex Smola, Le Song, and Ziyu Wang

Semantic Component Analysis 1484
Calvin Murdock and Fernando De la Torre

Low-Rank Matrix Factorization under General Mixture Noise Distributions 1493
Xiangyong Cao, Yang Chen, Qian Zhao, Deyu Meng, Yao Wang, Dong Wang, and Zongben Xu

Web-Scale Image Clustering Revisited 1502
Yannis Avrithis, Yannis Kalantidis, Evangelos Anagnostopoulos, and Ioannis Z. Emiris

Learning Discriminative Reconstructions for Unsupervised Outlier Removal 1511
Yan Xia, Xudong Cao, Fang Wen, Gang Hua, and Jian Sun

P-2A: Optimization, Segmentation, and Recognition

Learning Deconvolution Network for Semantic Segmentation 1520
Hyeonwoo Noh, Seunghoon Hong, and Bohyung Han

Conditional Random Fields as Recurrent Neural Networks 1529
Shuai Zheng, Sadeep Jayasumana, Bernardino Romera-Paredes, Vibhav Vineet, Zhizhong Su, Dalong Du, Chang Huang, and Philip H. S. Torr

The One Triangle Three Parallelograms Sampling Strategy and Its Application in Shape Regression 1538
Mikael Nilsson

Boosting Object Proposals: From Pascal to COCO 1546
Jordi Pont-Tuset and Luc Van Gool

Secrets of GrabCut and Kernel K-Means 1555
Meng Tang, Ismail Ben Ayed, Dmitrii Marin, and Yuri Boykov

Video Matting via Sparse and Low-Rank Representation 1564
Dongqing Zou, Xiaowu Chen, Guangying Cao, and Xiaogang Wang

Joint Object and Part Segmentation Using Deep Learned Potentials 1573
Peng Wang, Xiaohui Shen, Zhe Lin, Scott Cohen, Brian Price, and Alan Yuille

Low-Rank Tensor Constrained Multiview Subspace Clustering 1582
Changqing Zhang, Huazhu Fu, Si Liu, Guangcan Liu, and Xiaochun Cao

BodyPrint: Pose Invariant 3D Shape Matching of Human Bodies 1591
Jiangping Wang, Kai Ma, Vivek Kumar Singh, Thomas Huang, and Terrence Chen

The Middle Child Problem: Revisiting Parametric Min-Cut and Seeds for Object Proposals 1600
Ahmad Humayun, Fuxin Li, and James M. Rehg

Contour Guided Hierarchical Model for Shape Matching 1609
Yuanqi Su, Yuehu Liu, Bonan Cuan, and Nanning Zheng

Robust Image Segmentation Using Contour-Guided Color Palettes 1618
Xiang Fu, Chien-Yi Wang, Chen Chen, Changhu Wang, and C.-C. Jay Kuo

Joint Optimization of Segmentation and Color Clustering 1626
Ekaterina Lobacheva, Olga Veksler, and Yuri Boykov

BoxSup: Exploiting Bounding Boxes to Supervise Convolutional Networks for Semantic Segmentation 1635
Jifeng Dai, Kaiming He, and Jian Sun

Detection and Segmentation of 2D Curved Reflection Symmetric Structures 1644
Ching L. Teo, Cornelia Fermüller, and Yiannis Aloimonos

Unsupervised Tube Extraction Using Transductive Learning and Dense Trajectories 1653
Mihai Marian Puscas, Enver Sangineto, Dubravko Culibrk, and Nicu Sebe

Compositional Hierarchical Representation of Shape Manifolds for Classification of Non-manifold Shapes 1662
Mete Ozay, Umit Rusen Aktas, Jeremy L. Wyatt, and Aleš Leonardis

Shell PCA: Statistical Shape Modelling in Shell Space 1671
Chao Zhang, Behrend Heeren, Martin Rumpf, and William A. P. Smith

Learning to Combine Mid-Level Cues for Object Proposal Generation 1680
Tom Lee, Sanja Fidler, and Sven Dickinson

Enhancing Road Maps by Parsing Aerial Images Around the World 1689
Gellért Máttyus, Shenlong Wang, Sanja Fidler, and Raquel Urtasun

Probabilistic Appearance Models for Segmentation and Classification 1698
Julia Krüger, Jan Ehrhardt, and Heinz Handels

A Randomized Ensemble Approach to Industrial CT Segmentation 1707
Hyojin Kim, Jayaraman J. Thiagarajan, and Peer-Timo Bremer

Semi-Supervised Normalized Cuts for Image Segmentation 1716
Selene E. Chew and Nathan D. Cahill

StereoSnakes: Contour Based Consistent Object Extraction for Stereo Images 1724
Ran Ju, Tongwei Ren, and Gangshan Wu

Semantic Segmentation of RGBD Images with Mutex Constraints 1733
Zhuo Deng, Sinisa Todorovic, and Longin Jan Latecki

Weakly-and Semi-Supervised Learning of a Deep Convolutional Network for Semantic Image Segmentation 1742
George Papandreou, Liang-Chieh Chen, Kevin P. Murphy, and Alan L. Yuille

Efficient Decomposition of Image and Mesh Graphs by Lifted Multicuts 1751
M. Keuper, E. Levinkov, N. Bonneel, G. Lavoué, T. Brox, and B. Andres

Parsimonious Labeling 1760
Puneet K. Dokania and M. Pawan Kumar

Volumetric Bias in Segmentation and Reconstruction: Secrets and Solutions 1769
Yuri Boykov, Hossam Isack, Carl Olsson, and Ismail Ben Ayed

Entropy Minimization for Convex Relaxation Approaches 1778
Mohamed Souiai, Martin R. Oswald, Youngwook Kee, Junmo Kim, Marc Pollefeys, and Daniel Cremers

Adaptively Unified Semi-Supervised Dictionary Learning with Active Points 1787
Xiaobo Wang, Xiaojie Guo, and Stan Z. Li

Constrained Convolutional Neural Networks for Weakly Supervised Segmentation 1796
Deepak Pathak, Philipp Krähenbühl, and Trevor Darrell

A Multiscale Variable-Grouping Framework for MRF Energy Minimization 1805
Omer Meir, Meirav Galun, Stav Yagev, Ronen Basri, and Irad Yavneh

Inferring M-Best Diverse Labelings in a Single One 1814
Alexander Kirillov, Bogdan Savchynskyy, Dmitrij Schlesinger, Dmitry Vetrov, and Carsten Rother

Convolutional Sparse Coding for Image Super-Resolution 1823
Shuhang Gu, Wangmeng Zuo, Qi Xie, Deyu Meng, Xiangchu Feng, and Lei Zhang

A Wavefront Marching Method for Solving the Eikonal Equation on Cartesian Grids 1832
B. Cancela, M. Ortega, and M. G. Penedo

A Projection Free Method for Generalized Eigenvalue Problem with a Nonsmooth Regularizer 1841
Seong Jae Hwang, Maxwell D. Collins, Sathya N. Ravi, Vamsi K. Ithapu, Nagesh Adluru, Sterling C. Johnson, and Vikas Singh

Optimizing Expected Intersection-Over-Union with Candidate-Constrained CRFs 1850
Faruk Ahmed, Dany Tarlow, and Dhruv Batra

Higher-Order Inference for Multi-class Log-Supermodular Models 1859
Jian Zhang, Josip Djolonga, and Andreas Krause

Depth-Based Hand Pose Estimation: Data, Methods, and Challenges 1868
James S. Supančič III, Grégory Rogez, Yi Yang, Jamie Shotton, and Deva Ramanan

Adaptive Dither Voting for Robust Spatial Verification 1877
Xiaomeng Wu and Kunio Kashino

Alternating Co-Quantization for Cross-Modal Hashing 1886
Go Irie, Hiroyuki Arai, and Yukinobu Taniguchi

Learning Deep Representation with Large-Scale Attributes 1895
Wanli Ouyang, Hongyang Li, Xingyu Zeng, and Xiaogang Wang

Deep Learning Strong Parts for Pedestrian Detection 1904
Yonglong Tian, Ping Luo, Xiaogang Wang, and Xiaoou Tang

Flowing ConvNets for Human Pose Estimation in Videos 1913
Tomas Pfister, James Charles, and Andrew Zisserman

Top Rank Supervised Binary Coding for Visual Search 1922
Dongjin Song, Wei Liu, Rongrong Ji, David A. Meyer, and John R. Smith

BubbLeNet: Foveated Imaging for Visual Discovery 1931
Kevin Matzen and Noah Snavely

PQTable: Fast Exact Asymmetric Distance Neighbor Search for Product Quantization Using Hash Tables 1940
Yusuke Matsui, Toshihiko Yamasaki, and Kiyoharu Aizawa

Lending A Hand: Detecting Hands and Recognizing Activities in Complex Egocentric Interactions 1949
Sven Bambach, Stefan Lee, David J. Crandall, and Chen Yu

Fast and Accurate Head Pose Estimation via Random Projection Forests 1958
Donghoon Lee, Ming-Hsuan Yang, and Songhwai Oh

An MRF-Poselets Model for Detecting Highly Articulated Humans 1967
Duc Thanh Nguyen, Minh-Khoi Tran, and Sai-Kit Yeung

Beyond Tree Structure Models: A New Occlusion Aware Graphical Model for Human Pose Estimation 1976
Lianrui Fu, Junge Zhang, and Kaiqi Huang

Relaxing from Vocabulary: Robust Weakly-Supervised Deep Learning for Vocabulary-Free Image Tagging 1985
Jianlong Fu, Yue Wu, Tao Mei, Jinqiao Wang, Hanqing Lu, and Yong Rui

Visual Phrases for Exemplar Face Detection 1994
Vijay Kumar, Anoop Namboodiri, and C. V. Jawahar

Spatial Semantic Regularisation for Large Scale Object Detection 2003
Damian Mrowca, Marcus Rohrbach, Judy Hoffman, Ronghang Hu, Kate Saenko, and Trevor Darrell

Human Pose Estimation in Videos 2012
Dong Zhang and Mubarak Shah

Contour Box: Rejecting Object Proposals without Explicit Closed Contours 2021
Cewu Lu, Shu Liu, Jiaya Jia, and Chi-Keung Tang

O-3A: Registration, Alignment and Stereo

Registering Images to Untextured Geometry Using Average Shading Gradients 2030
Tobias Plötz and Stefan Roth

Robust Nonrigid Registration by Convex Optimization 2039
Qifeng Chen and Vladlen Koltun

Robust and Optimal Sum-of-Squares-Based Point-to-Plane Registration of Image Sets and Structured Scenes 2048
Danda Pani Paudel, Adlane Habed, Cédric Demonceaux, and Pascal Vasseur

MeshStereo: A Global Stereo Model with Mesh Alignment Regularization for View Interpolation 2057
Chi Zhang, Zhiwei Li, Yanhua Cheng, Rui Cai, Hongyang Chao, and Yong Rui

CV-HAZOP: Introducing Test Data Validation for Computer Vision 2066
Oliver Zendel, Markus Murschitz, Martin Humenberger, and Wolfgang Herzner

P-3A: Recognition and 3D Computer Vision II

Structure from Motion Using Structure-Less Resection 2075
Enliang Zheng and Changchang Wu

Joint Camera Clustering and Surface Segmentation for Large-Scale Multi-view Stereo 2084
Runze Zhang, Shiwei Li, Tian Fang, Siyu Zhu, and Long Quan

Higher-Order CRF Structural Segmentation of 3D Reconstructed Surfaces 2093
Jingbo Liu, Jinglu Wang, Tian Fang, Chiew-Lan Tai, and Long Quan

Hyperpoints and Fine Vocabularies for Large-Scale Location Recognition 2102
Torsten Sattler, Michal Havlena, Filip Radenović, Konrad Schindler, and Marc Pollefeys

Globally Optimal 2D-3D Registration from Points or Lines without Correspondences 2111
Mark Brown, David Windridge, and Jean-Yves Guillemaut

The HCI Stereo Metrics: Geometry-Aware Performance Analysis of Stereo Algorithms 2120
Katrin Honauer, Lena Maier-Hein, and Daniel Kondermann

Merging the Unmatchable: Stitching Visually Disconnected SfM Models 2129
Andrea Cohen, Torsten Sattler, and Marc Pollefeys

3D Fragment Reassembly Using Integrated Template Guidance and Fracture-Region Matching 2138
Kang Zhang, Wuyi Yu, Mary Manhein, Warren Waggenspack, and Xin Li

Procedural Editing of 3D Building Point Clouds 2147
İlke Demir, Daniel G. Aliaga, and Bedrich Benes

Semantically-Aware Aerial Reconstruction from Multi-modal Data 2156
Randi Cabezas, Julian Straub, and John W. Fisher III

Guaranteed Outlier Removal for Rotation Search 2165
Álvaro Parra Bustos and Tat-Jun Chin

Peeking Template Matching for Depth Extension 2174
Simon Korman, Eyal Ofek, and Shai Avidan

Deformable 3D Fusion: From Partial Dynamic 3D Observations to Complete 4D Models 2183
Weipeng Xu, Mathieu Salzmann, Yongtian Wang, and Yue Liu

Non-parametric Structure-Based Calibration of Radially Symmetric Cameras 2192
Federico Camposeco, Torsten Sattler, and Marc Pollefeys

Exploiting Object Similarity in 3D Reconstruction 2201
Chen Zhou, Fatma Güney, Yizhou Wang, and Andreas Geiger

You are Here: Mimicking the Human Thinking Process in Reading Floor-Plans 2210
Hang Chu, Dong Ki Kim, and Tsuhan Chen

MAP Disparity Estimation Using Hidden Markov Trees 2219
Eric T. Psota, Jędrzej Kowalczuk, Mateusz Mittek, and Lance C. Pérez

Wide Baseline Stereo Matching with Convex Bounded Distortion Constraints 2228
Meirav Galun, Tal Amir, Tal Hassner, Ronen Basri, and Yaron Lipman

Interactive Visual Hull Refinement for Specular and Transparent Object Surface Reconstruction 2237
Xinxin Zuo, Chao Du, Sen Wang, Jiangbin Zheng, and Ruigang Yang

Hierarchical Higher-Order Regression Forest Fields: An Application to 3D Indoor Scene Labelling 2246
Trung T. Pham, Ian Reid, Yasir Latif, and Stephen Gould

Classical Scaling Revisited 2255
Gil Shamai, Yonathan Aflalo, Michael Zibulevsky, and Ron Kimmel

Dense Continuous-Time Tracking and Mapping with Rolling Shutter RGB-D Cameras 2264
Christian Kerl, Jörg Stückler, and Daniel Cremers

Dense Image Registration and Deformable Surface Reconstruction in Presence of Occlusions and Minimal Texture 2273
Dat Tien Ngo, Sanghyuk Park, Anne Jorstad, Alberto Crivellaro, Chang D. Yoo, and Pascal Fua

The Likelihood-Ratio Test and Efficient Robust Estimation 2282
Andrea Cohen and Christopher Zach

Reflection Modeling for Passive Stereo 2291
Rahul Nair, Andrew Fitzgibbon, Daniel Kondermann, and Carsten Rother

Detailed Full-Body Reconstructions of Moving People from Monocular RGB-D Sequences 2300
Federica Bogo, Michael J. Black, Matthew Loper, and Javier Romero

Efficient Solution to the Epipolar Geometry for Radially Distorted Cameras 2309
Zuzana Kukelova, Jan Heller, Martin Bujnak, Andrew Fitzgibbon, and Tomas Pajdla

Learning a Descriptor-Specific 3D Keypoint Detector 2318
Samuele Salti, Federico Tombari, Riccardo Spezialetti, and Luigi Di Stefano

Component-Wise Modeling of Articulated Objects 2327
Valsamis Ntouskos, Marta Sanzari, Bruno Cafaro, Federico Nardi, Fabrizio Natola, Fiora Pirri, and Manuel Ruiz

A Collaborative Filtering Approach to Real-Time Hand Pose Estimation 2336
Chiho Choi, Ayan Sinha, Joon Hee Choi, Sujin Jang, and Karthik Ramani

On the Equivalence of Moving Entrance Pupil and Radial Distortion for Camera Calibration 2345
Avinash Kumar and Narendra Ahuja

A Linear Generalized Camera Calibration from Three Intersecting Reference Planes 2354
Mai Nishimura, Shohei Nobuhara, Takashi Matsuyama, Shinya Shimizu, and Kensaku Fujii

Towards Pointless Structure from Motion: 3D Reconstruction and Camera Parameters from General 3D Curves 2363
Irina Nurutdinova and Andrew Fitzgibbon

Attributed Grammars for Joint Estimation of Human Attributes, Part and Pose 2372
Seyoung Park and Song-Chun Zhu

Real-Time Pose Estimation Piggybacked on Object Detection 2381
Roman Juránek, Adam Herout, Markéta Dubská, and Pavel Zemčík

Understanding and Predicting Image Memorability at a Large Scale 2390
Aditya Khosla, Akhil S. Raju, Antonio Torralba, and Aude Oliva

Multiple Granularity Descriptors for Fine-Grained Categorization 2399
Dequan Wang, Zhiqiang Shen, Jie Shao, Wei Zhang, Xiangyang Xue, and Zheng Zhang

Guiding the Long-Short Term Memory Model for Image Caption Generation 2407
Xu Jia, Efstratios Gavves, Basura Fernando, and Tinne Tuytelaars

Just Noticeable Differences in Visual Attributes 2416
Aron Yu and Kristen Grauman

VQA: Visual Question Answering 2425
Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh

Localize Me Anywhere, Anytime: A Multi-task Point-Retrieval Approach 2434
Guoyu Lu, Yan Yan, Li Ren, Jingkuan Song, Nicu Sebe, and Chandra Kambhamettu

Dense Optical Flow Prediction from a Static Image 2443
Jacob Walker, Abhinav Gupta, and Martial Hebert

Unsupervised Domain Adaptation for Zero-Shot Learning 2452
Elyor Kodirov, Tao Xiang, Zhenyong Fu, and Shaogang Gong

Visual Madlibs: Fill in the Blank Description Generation and Question Answering 2461
Licheng Yu, Eunbyung Park, Alexander C. Berg, and Tamara L. Berg

Actions and Attributes from Wholes and Parts 2470
Georgia Gkioxari, Ross Girshick, and Jitendra Malik

DeepBox: Learning Objectness with Convolutional Networks 2479
Weicheng Kuo, Bharath Hariharan, and Jitendra Malik

Active Object Localization with Deep Reinforcement Learning 2488
Juan C. Caicedo and Svetlana Lazebnik

Scene-Domain Active Part Models for Object Representation 2497
Zhou Ren, Chaohui Wang, and Alan Yuille

A Unified Multiplicative Framework for Attribute Learning 2506
Kongming Liang, Hong Chang, Shiguang Shan, and Xilin Chen

Contractive Rectifier Networks for Nonlinear Maximum Margin Classification 2515
Senjian An, Munawar Hayat, Salman H. Khan, Mohammed Bennamoun, Farid Boussaid, and Ferdous Sohel

Augmenting Strong Supervision Using Web Data for Fine-Grained Categorization 2524
Zhe Xu, Shaoli Huang, Ya Zhang, and Dacheng Tao

Learning Like a Child: Fast Novel Visual Concept Learning from Sentence Descriptions of Images 2533
Junhua Mao, Xu Wei, Yi Yang, Jiang Wang, Zhiheng Huang, and Alan L. Yuille

Learning Common Sense through Visual Abstraction 2542
Ramakrishna Vedantam, Xiao Lin, Tanmay Batra, C. Lawrence Zitnick, and Devi Parikh

Domain Generalization for Object Recognition with Multi-task Autoencoders 2551
Muhammad Ghifary, W. Bastiaan Kleijn, Mengjie Zhang, and David Balduzzi

Square Localization for Efficient and Accurate Object Detection 2560
Cewu Lu, Yongyi Lu, Hao Chen, and Chi-Keung Tang

Box Aggregation for Proposal Decimation: Last Mile of Object Detection 2569
Shu Liu, Cewu Lu, and Jiaya Jia

DeepProposal: Hunting Objects by Cascading Deep Convolutional Layers 2578
Amir Ghodrati, Ali Diba, Marco Pedersoli, Tinne Tuytelaars, and Luc Van Gool

Semantic Segmentation with Object Clique Potential 2587
Xiaojuan Qi, Jianping Shi, Shu Liu, Renjie Liao, and Jiaya Jia

Automatic Concept Discovery from Parallel Text and Visual Corpora 2596
Chen Sun, Chuang Gan, and Ram Nevatia

Simpler Non-Parametric Methods Provide as Good or Better Results to Multiple-Instance Learning 2605
Ragav Venkatesan, Parag Shridhar Chandakkar, and Baoxin Li

Monocular Object Instance Segmentation and Depth Ordering with CNNs 2614
Ziyu Zhang, Alexander G. Schwing, Sanja Fidler, and Raquel Urtasun

Multimodal Convolutional Neural Networks for Matching Image and Sentence 2623
Lin Ma, Zhengdong Lu, Lifeng Shang, and Hang Li

Structural Kernel Learning for Large Scale Multiclass Object Co-detection 2632
Zeeshan Hayder, Xuming He, and Mathieu Salzmann

Flickr30k Entities: Collecting Region-to-Phrase Correspondences for Richer Image-to-Sentence Models 2641
Bryan A. Plummer, Liwei Wang, Chris M. Cervantes, Juan C. Caicedo, Julia Hockenmaier, and Svetlana Lazebnik

Predicting Depth, Surface Normals and Semantic Labels with a Common Multi-scale Convolutional Architecture 2650
David Eigen and Rob Fergus

AttentionNet: Aggregating Weak Directions for Accurate Object Detection 2659
Donggeun Yoo, Sunggyun Park, Joon-Young Lee, Anthony S. Paek, and In So Kweon

Common Subspace for Model and Similarity: Phrase Learning for Caption Generation from Images 2668
Yoshitaka Ushiku, Masataka Yamaguchi, Yusuke Mukuta, and Tatsuya Harada

O-3B: Representations for Recognition and Localization

3D-Assisted Feature Synthesis for Novel Views of an Object 2677
Hao Su, Fan Wang, Eric Yi, and Leonidas Guibas

Render for CNN: Viewpoint Estimation in Images Using CNNs Trained with Rendered 3D Model Views 2686
Hao Su, Charles R. Qi, Yangyan Li, and Leonidas J. Guibas

Lost Shopping! Monocular Localization in Large Indoor Spaces 2695
Shenlong Wang, Sanja Fidler, and Raquel Urtasun

Camera Pose Voting for Large-Scale Image-Based Localization 2704
Bernhard Zeisl, Torsten Sattler, and Marc Pollefeys

P-3B: Statistical Methods and Learning, Motion and Tracking, and Video Analysis I

MANTRA: Minimum Maximum Latent Structural SVM for Image Classification and Ranking 2713
Thibaut Durand, Nicolas Thome, and Matthieu Cord

DeepDriving: Learning Affordance for Direct Perception in Autonomous Driving 2722
Chenyi Chen, Ari Seff, Alain Kornhauser, and Jianxiong Xiao

Active Transfer Learning with Zero-Shot Priors: Reusing Past Datasets for Future Tasks 2731
E. Gavves, T. Mensink, T. Tommasi, C. G. M. Snoek, and T. Tuytelaars

HD-CNN: Hierarchical Deep Convolutional Neural Networks for Large Scale Visual Recognition 2740
Zhicheng Yan, Hao Zhang, Robinson Piramuthu, Vignesh Jagadeesh, Dennis DeCoste, Wei Di, and Yizhou Yu

Learning the Structure of Deep Convolutional Networks 2749
Jiashi Feng and Trevor Darrell

FlowNet: Learning Optical Flow with Convolutional Networks 2758
Alexey Dosovitskiy, Philipp Fischery, Eddy Ilg, Philip Häusser, Caner Hazırbaş, Vladimir Golkov, Patrick van der Smagt, Daniel Cremers, and Thomas Brox

Learning Semi-Supervised Representation Towards a Unified Optimization Framework for Semi-Supervised Learning 2767
Chun-Guang Li, Zhouchen Lin, Honggang Zhang, and Jun Guo

Context-Guided Diffusion for Label Propagation on Graphs 2776
Kwang In Kim, James Tompkin, Hanspeter Pfister, and Christian Theobalt

Learning to Rank Based on Subsequences 2785
Basura Fernando, Efstratios Gavves, Damien Muselet, and Tinne Tuytelaars

Unsupervised Learning of Visual Representations Using Videos 2794
Xiaolong Wang and Abhinav Gupta

A Nonparametric Bayesian Approach toward Stacked Convolutional Independent Component Analysis 2803
Sotirios P. Chatzis and Dimitrios Kosmopoulos

Robust Principal Component Analysis on Graphs 2812
Nauman Shahid, Vassilis Kalofolias, Xavier Bresson, Michael Bronstein, and Pierre Vandergheynst

Projection Bank: From High-Dimensional Data to Medium-Length Binary Codes 2821
Li Liu, Mengyang Yu, and Ling Shao

Robust Optimization for Deep Regression 2830
Vasileios Belagiannis, Christian Rupprecht, Gustavo Carneiro, and Nassir Navab

Multi-class Multi-annotator Active Learning with Robust Gaussian Process for Visual Recognition 2839
Chengjiang Long and Gang Hua

Maximum-Margin Structured Learning with Deep Networks for 3D Human Pose Estimation 2848
Sijin Li, Weichen Zhang, and Antoni B. Chan

An Exploration of Parameter Redundancy in Deep Networks with Circulant Projections 2857
Yu Cheng, Felix X. Yu, Rogerio S. Feris, Sanjiv Kumar, Alok Choudhary, and Shi-Fu Chang

Additive Nearest Neighbor Feature Maps 2866
Zhenzhen Wang, Xiao-Tong Yuan, Qingshan Liu, and Shuicheng Yan

Understanding Deep Features with Computer-Generated Imagery 2875
Mathieu Aubry and Bryan C. Russell

Interpolation on the Manifold of K Component GMMs 2884
Hyunwoo J. Kim, Nagesh Adluru, Monami Banerjee, Baba C. Vemuri, and Vikas Singh

Context-Aware CNNs for Person Head Detection 2893
Tuan-Hung Vu, Anton Osokin, and Ivan Laptev

Mode-Seeking on Hypergraphs for Robust Geometric Model Fitting 2902
Hanzi Wang, Guobao Xiao, Yan Yan, and David Suter

Highly-Expressive Spaces of Well-Behaved Transformations: Keeping it Simple 2911
Oren Freifeld, Søren Hauberg, Kayhan Batmanghelich, and John W. Fisher III

Entropy-Based Latent Structured Output Prediction 2920
Diane Bouchacourt, Sebastian Nowozin, and M. Pawan Kumar

Fast Orthogonal Projection Based on Kronecker Product 2929
Xu Zhang, Felix X. Yu, Ruiqi Guo, Sanjiv Kumar, Shengjin Wang, and Shi-Fu Chang

PoseNet: A Convolutional Network for Real-Time 6-DOF Camera Relocalization 2938
Alex Kendall, Matthew Grimes, and Roberto Cipolla

Predicting Multiple Structured Visual Interpretations 2947
Debadeepta Dey, Varun Ramakrishna, Martial Hebert, and J. Andrew Bagnell

Look and Think Twice: Capturing Top-Down Visual Attention with Feedback Convolutional Neural Networks 2956
Chunshui Cao, Xianming Liu, Yi Yang, Yinan Yu, Jiang Wang, Zilei Wang, Yongzhen Huang, Liang Wang, Chang Huang, Wei Xu, Deva Ramanan, and Thomas S. Huang

Matrix Backpropagation for Deep Networks with Structured Layers 2965
Catalin Ionescu, Orestis Vantzos, and Cristian Sminchisescu

Introducing Geometry in Active Learning for Image Segmentation 2974
Ksenia Konyushkova, Raphael Sznitman, and Pascal Fua

Joint Fine-Tuning in Deep Neural Networks for Facial Expression Recognition 2983
Heechul Jung, Sihaeng Lee, Junho Yim, Sunjeong Park, and Junmo Kim

Direct Intrinsics: Learning Albedo-Shading Decomposition by Convolutional Regression 2992
Takuya Narihira, Michael Maire, and Stella X. Yu

Face Flow 2993
Patrick Snape, Anastasios Roussos, Yannis Panagakis, and Stefanos Zafeiriou

Discriminative Low-Rank Tracking 3002
Yao Sui, Yafei Tang, and Li Zhang

SOWP: Spatially Ordered and Weighted Patch Descriptor for Visual Tracking 3011
Han-Ul Kim, Dae-Youn Lee, Jae-Young Sim, and Chang-Su Kim

Live Repetition Counting 3020
Ofir Levy and Lior Wolf

Near-Online Multi-target Tracking with Aggregated Local Flow Descriptor 3029
Wongun Choi

Multi-kernel Correlation Filter for Visual Tracking 3038
Ming Tang and Jiayi Feng

Joint Probabilistic Data Association Revisited 3047
Seyed Hamid Rezatofighi, Anton Milan, Zhen Zhang, Qinfeng Shi, Anthony Dick, and Ian Reid

Tracking-by-Segmentation with Online Gradient Boosting Decision Tree 3056
Jeany Son, Ilchae Jung, Kayoung Park, and Bohyung Han

Exploring Causal Relationships in Visual Object Tracking 3065
Karel Lebeda, Simon Hadfield, and Richard Bowden

Hierarchical Convolutional Features for Visual Tracking 3074
Chao Ma, Jia-Bin Huang, Xiaokang Yang, and Ming-Hsuan Yang

Robust Non-rigid Motion Tracking and Surface Reconstruction Using L0 Regularization 3083
Kaiwen Guo, Feng Xu, Yangang Wang, Yebin Liu, and Qionghai Dai

Online Object Tracking with Proposal Selection 3092
Yang Hua, Karteek Alahari, and Cordelia Schmid

Understanding and Diagnosing Visual Tracking Systems 3101
Naiyan Wang, Jianping Shi, Dit-Yan Yeung, and Jiaya Jia

Integrating Dashcam Views through Inter-Video Mapping 3110
Hsin-I Chen, Yi-Ling Chen, Wei-Tse Lee, Fan Wang, and Bing-Yu Chen

Visual Tracking with Fully Convolutional Networks 3119
Lijun Wang, Wanli Ouyang, Xiaogang Wang, and Huchuan Lu

Multiple Feature Fusion via Weighted Entropy for Visual Tracking 3128
Lin Ma, Jiwen Lu, Jianjiang Feng, and Jie Zhou

Pedestrian Travel Time Estimation in Crowded Scenes 3137
Shuai Yi, Hongsheng Li, and Xiaogang Wang

Unsupervised Synchrony Discovery in Human Interaction 3146
Wen-Sheng Chu, Jiabei Zeng, Fernando De la Torre, Jeffrey F. Cohn, and Daniel S. Messinger

Efficient Video Segmentation Using Parametric Graph Partitioning 3155
Chen-Ping Yu, Hieu Le, Gregory Zelinsky, and Dimitris Samaras

Learning to Track for Spatio-Temporal Action Localization 3164
Philippe Weinzaepfel, Zaid Harchaoui, and Cordelia Schmid

Unsupervised Object Discovery and Tracking in Video Collections 3173
Suha Kwak, Minsu Cho, Ivan Laptev, Jean Ponce, and Cordelia Schmid

Car that Knows Before You Do: Anticipating Maneuvers via Learning Temporal Driving Models 3182
Ashesh Jain, Hema S. Koppula, Bharad Raghavan, Shane Soh, and Ashutosh Saxena

Activity Auto-Completion: Predicting Human Activities from Partial Videos 3191
Zhen Xu, Laiyun Qing, and Jun Miao

Person Re-Identification with Correspondence Structure Learning 3200
Yang Shen, Weiyao Lin, Junchi Yan, Mingliang Xu, Jianxin Wu, and Jingdong Wang

Adaptive Exponential Smoothing for Online Filtering of Pixel Prediction Maps 3209
Kang Dang, Jiong Yang, and Junsong Yuan

P-CNN: Pose-Based CNN Features for Action Recognition 3218
Guilhem Chéron, Ivan Laptev, and Cordelia Schmid

Fully Connected Object Proposals for Video Segmentation 3227
Federico Perazzi, Oliver Wang, Markus Gross, and Alexander Sorkine-Hornung

Video Segmentation with Just a Few Strokes 3235
Naveen Shankar Nagaraja, Frank R. Schmidt, and Thomas Brox

Actionness-Assisted Recognition of Actions 3244
Ye Luo, Loong-Fah Cheong, and An Tran

COUNT Forest: CO-Voting Uncertain Number of Targets Using Random Forest for Crowd Density Estimation 3253
Viet-Quoc Pham, Tatsuo Kozakaya, Osamu Yamaguchi, and Ryuzo Okada

Multi-cue Structure Preserving MRF for Unconstrained Video Segmentation 3262
Saehoon Yi and Vladimir Pavlovic

Motion Trajectory Segmentation via Minimum Cost Multicuts 3271
Margret Keuper, Bjoern Andres, and Thomas Brox

Action Localization in Videos through Context Walk 3280
Khurram Soomro, Haroon Idrees, and Mubarak Shah

RGB-W: When Vision Meets Wireless 3289
Alexandre Alahi, Albert Haque, and Li Fei-Fei

Action Detection by Implicit Intentional Motion Clustering 3298
Wei Chen and Jason J. Corso

Simultaneous Foreground Detection and Classification with Hybrid Features 3307
Jaemyun Kim, Adín Ramírez Rivera, Byungyong Ryu, and Oksam Chae

O-3C: Vision and People

Training a Feedback Loop for Hand Pose Estimation 3316
Markus Oberweger, Paul Wohlhart, and Vincent Lepetit

Opening the Black Box: Hierarchical Sampling Optimization for Estimating Human Hand Pose 3325
Danhang Tang, Jonathan Taylor, Pushmeet Kohli, Cem Keskin, Tae-Kyun Kim, and Jamie Shotton

Panoptic Studio: A Massively Multiview System for Social Motion Capture 3334
Hanbyul Joo, Hao Liu, Lei Tan, Lin Gui, Bart Nabbe, Iain Matthews, Takeo Kanade, Shohei Nobuhara, and Yaser Sheikh

Where to Buy It: Matching Street Clothing Photos in Online Shops 3343
M. Hadi Kiapour, Xufeng Han, Svetlana Lazebnik, Alexander C. Berg, and Tamara L. Berg

Multi-task Recurrent Neural Network for Immediacy Prediction 3352
Xiao Chu, Wanli Ouyang, Wei Yang, and Xiaogang Wang

Learning Complexity-Aware Cascades for Deep Pedestrian Detection 3361
Zhaowei Cai, Mohammad Saberian, and Nuno Vasconcelos

O-4A: Computational Photography and Image Enhancement

Polarized 3D: High-Quality Depth Sensing with Polarization Cues 3370
Achuta Kadambi, Vage Taamazyan, Boxin Shi, and Ramesh Raskar

Airborne Three-Dimensional Cloud Tomography 3379
Aviad Levis, Yoav Y. Schechner, Amit Aides, and Anthony B. Davis

Leave-One-Out Kernel Optimization for Shadow Detection 3388
Tomás F. Yago Vicente, Minh Hoai, and Dimitris Samaras

Removing Rain from a Single Image via Discriminative Sparse Coding 3397
Yu Luo, Yong Xu, and Hui Ji

Mutual-Structure for Joint Filtering 3406
Xiaoyong Shen, Chao Zhou, Li Xu, and Jiaya Jia

P-4A: Computational Photography, Face and Gesture, and Vision for X

Photometric Stereo in a Scattering Medium 3415
Zak Murez, Tali Treibitz, Ravi Ramamoorthi, and David Kriegman

Resolving Scale Ambiguity via XSlit Aspect Ratio Analysis 3424
Wei Yang, Haiting Lin, Sing Bing Kang, and Jingyi Yu

Single-Shot Specular Surface Reconstruction with Gonio-Plenoptic Imaging 3433
Lingfei Meng, Liyang Lu, Noah Bedard, and Kathrin Berkner

TransCut: Transparent Object Segmentation from a Light-Field Image 3442
Yichao Xu, Hajime Nagahara, Atsushi Shimada, and Rin-ichiro Taniguchi

Depth Recovery from Light Field Using Focal Stack Symmetry 3451
Haiting Lin, Can Chen, Sing Bing Kang, and Jingyi Yu

Depth Map Estimation and Colorization of Anaglyph Images Using Local Color Prior and Reverse Intensity Distribution 3460
Williem, Ramesh Raskar, and In Kyu Park

Learning Data-Driven Reflectance Priors for Intrinsic Image Decomposition 3469
Tinghui Zhou, Philipp Krähenbühl, and Alexei A. Efros

Photometric Stereo with Small Angular Variations 3478
Jian Wang, Yasuyuki Matsushita, Boxin Shi, and Aswin C. Sankaranarayanan

Occlusion-Aware Depth Estimation Using Light-Field Cameras 3487
Ting-Chun Wang, Alexei A. Efros, and Ravi Ramamoorthi

Oriented Light-Field Windows for Scene Flow 3496
Pratul P. Srinivasan, Michael W. Tao, Ren Ng, and Ravi Ramamoorthi

Extended Depth of Field Catadioptric Imaging Using Focal Sweep 3505
Ryunosuke Yokoya and Shree K. Nayar

Intrinsic Depth: Improving Depth Transfer with Intrinsic Images 3514
Naejin Kong and Michael J. Black

Separating Fluorescent and Reflective Components by Using a Single Hyperspectral Image 3523
Yinqiang Zheng, Ying Fu, Antony Lam, Imari Sato, and Yoichi Sato

Frequency-Based Environment Matting by Compressive Sensing 3532
Yiming Qian, Minglun Gong, and Yee-Hong Yang

Complementary Sets of Shutter Sequences for Motion Deblurring 3541
Hae-Gon Jeon, Joon-Young Lee, Yudeog Han, Seon Joo Kim, and In So Kweon

Hyperspectral Compressive Sensing Using Manifold-Structured Sparsity Prior 3550
Lei Zhang, Wei Wei, Yanning Zhang, Fei Li, Chunhua Shen, and Qinfeng Shi

A Gaussian Process Latent Variable Model for BRDF Inference 3559
Stamatios Georgoulis, Vincent Vanweddingen, Marc Proesmans, and Luc Van Gool

Active One-Shot Scan for Wide Depth Range Using a Light Field Projector Based on Coded Aperture 3568
Hiroshi Kawasaki, Satoshi Ono, Yuki Horita, Yuki Shiba, Ryo Furukawa, and Shinsaku Hiura

Model-Based Tracking at 300Hz Using Raw Time-of-Flight Observations 3577
Jan Stühmer, Sebastian Nowozin, Andrew Fitzgibbon, Richard Szeliski, Travis Perry, Sunil Acharya, Daniel Cremers, and Jamie Shotton

Hyperspectral Super-Resolution by Coupled Spectral Unmixing 3586
Charis Lanaras, Emmanuel Baltsavias, and Konrad Schindler

Depth Selective Camera: A Direct, On-Chip, Programmable Technique for Depth Selectivity in Photography 3595
Ryuichi Tadano, Adithya Kumar Pediredla, and Ashok Veeraraghavan

A Groupwise Multilinear Correspondence Optimization for 3D Faces 3604
Timo Bolkart and Stefanie Wuhrer

Selective Encoding for Recognizing Unreliably Localized Faces 3613
Ang Li, Vlad I. Morariu, and Larry S. Davis

Confidence Preserving Machine for Facial Action Unit Detection 3622
Jiabei Zeng, Wen-Sheng Chu, Fernando De la Torre, Jeffrey F. Cohn, and Zhang Xiong

Learning Social Relation Traits from Face Images 3631
Zhanpeng Zhang, Ping Luo, Chen-Change Loy, and Xiaoou Tang

Robust Heart Rate Measurement from Video Using Select Random Patches 3640
Antony Lam and Yoshinori Kuno

Robust Model-Based 3D Head Pose Estimation 3649
Gregory P. Meyer, Shalini Gupta, Iuri Frosio, Dikpal Reddy, and Jan Kautz

Robust Facial Landmark Detection Under Significant Head Poses and Occlusion 3658
Yue Wu and Qiang Ji

Conditional Convolutional Neural Network for Modality-Aware Face Recognition 3667
Chao Xiong, Xiaowei Zhao, Danhang Tang, Karlekar Jayashree, Shuicheng Yan, and Tae-Kyun Kim

From Facial Parts Responses to Face Detection: A Deep Learning Approach 3676
Shuo Yang, Ping Luo, Chen-Change Loy, and Xiaoou Tang

Efficient PSD Constrained Asymmetric Metric Learning for Person Re-Identification 3685
Shengcai Liao and Stan Z. Li

Pose-Invariant 3D Face Alignment 3694
Amin Jourabloo and Xiaoming Liu

From Emotions to Action Units with Hidden and Semi-Hidden-Task Learning 3703
Adria Ruiz, Joost Van de Weijer, and Xavier Binefa

Automated Facial Trait Judgment and Election Outcome Prediction: Social Dimensions of Face 3712
Jungseock Joo, Francis F. Steen, and Song-Chun Zhu

Simultaneous Local Binary Feature Learning and Encoding for Face Recognition 3721
Jiwen Lu, Venice Erin Liong, and Jie Zhou

Deep Learning Face Attributes in the Wild 3730
Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang

Multi-Task Learning with Low Rank Attribute Embedding for Person Re-Identification 3739
Chi Su, Fan Yang, Shiliang Zhang, Qi Tian, Larry S. Davis, and Wen Gao

Regressing a 3D Face Shape from a Single Image 3748
Sergey Tulyakov and Nicu Sebe

Rendering of Eyes for Eye-Shape Registration and Gaze Estimation 3756
Erroll Wood, Tadas Baltrušaitis, Xucong Zhang, Yusuke Sugano, Peter Robinson, and Andreas Bulling

Multi-Scale Learning for Low-Resolution Person Re-Identification 3765
Xiang Li, Wei-Shi Zheng, Xiaojuan Wang, Tao Xiang, and Shaogang Gong

Learning to Transfer: Transferring Latent Task Structures and Its Application to Person-Specific Facial Action Unit Detection 3774
Timur Almaev, Brais Martinez, and Michel Valstar

Pairwise Conditional Random Forests for Facial Expression Recognition 3783
Arnaud Dapogny, Kevin Bailly, and Séverine Dubuisson

Multi-conditional Latent Variable Model for Joint Facial Action Unit Detection 3792
Stefanos Eleftheriadis, Ognjen Rudovic, and Maja Pantic

Leveraging Datasets with Varying Annotations for Face Alignment via Deep Regression Network 3801
Jie Zhang, Meina Kan, Shiguang Shan, and Xilin Chen

A Spatio-Temporal Appearance Representation for Viceo-Based Pedestrian Re-Identification 3810
Kan Liu, Bingpeng Ma, Wei Zhang, and Rui Huang

Two Birds, One Stone: Jointly Learning Binary Code for Large-Scale Face Image Retrieval and Attributes Prediction 3819
Yan Li, Ruiping Wang, Haomiao Liu, Huajie Jiang, Shiguang Shan, and Xilin Chen

An Accurate Iris Segmentation Framework Under Relaxed Imaging Constraints Using Total Variation Model 3828
Zijing Zhao and Ajay Kumar

Discriminative Pose-Free Descriptors for Face and Object Matching 3837
Soubhik Sanyal, Sivaram Prasad Mudunuri, and Soma Biswas

Bi-Shifting Auto-Encoder for Unsupervised Domain Adaptation 3846
Meina Kan, Shiguang Shan, and Xilin Chen

Regressive Tree Structured Model for Facial Landmark Localization 3855
Gee-Sern (Jison) Hsu, Kai-Hsiang Chang, and Shih-Chieh Huang

Person Recognition in Personal Photo Collections 3862
Seong Joon Oh, Rodrigo Benenson, Mario Fritz, and Bernt Schiele

Robust Statistical Face Frontalization 3871
Christos Sagonas, Yannis Panagakis, Stefanos Zafeiriou, and Maja Pantic

PIEFA: Personalized Incremental and Ensemble Face Alignment 3880
Xi Peng, Shaoting Zhang, Yu Yang, and Dimitris N. Metaxas

Understanding Everyday Hands in Action from RGB-D Images 3889
Grégory Rogez, James S. Supančič III, and Deva Ramanan

Example-Based Modeling of Facial Texture from Deficient Data 3898
Arnaud Dessein, William A. P. Smith, Richard C. Wilson, and Edwin R. Hancock

Learning to Predict Saliency on Face Images 3907
Mai Xu, Yun Ren, and Zulin Wang

Group Membership Prediction 3916
Ziming Zhang, Yuting Chen, and Venkatesh Saligrama

Extraction of Virtual Baselines from Distorted Document Images Using Curvilinear Projection 3925
Gaofeng Meng, Zuming Huang, Yonghong Song, Shiming Xiang, and Chunhong Pan

Robust RGB-D Odometry Using Point and Line Features 3934
Yan Lu and Dezhen Song

Learning a Discriminative Model for the Perception of Realism in Composite Images 3943
Jun-Yan Zhu, Philipp Krähenbühl, Eli Shechtman, and Alexei A. Efros

What Makes Tom Hanks Look Like Tom Hanks 3952
Supasorn Suwajanakorn, Steven M. Seitz, and Ira Kemelmacher-Shlizerman

Wide-Area Image Geolocalization with Aerial Reference Imagery 3961
Scott Workman, Richard Souvenir, and Nathan Jacobs

Personalized Age Progression with Aging Dictionary 3970
Xiangbo Shu, Jinhui Tang, Hanjiang Lai, Luoqi Liu, and Shuicheng Yan

FaceDirector: Continuous Control of Facial Performance in Video 3979
Charles Malleson, Jean-Charles Bazin, Oliver Wang, Derek Bradley, Thabo Beeler, Adrian Hilton, and Alexander Sorkine-Hornung

Synthesizing Illumination Mosaics from Internet Photo-Collections 3988
Dinghuang Ji, Enrique Dunn, and Jan-Michael Frahm

Hot or Not: Exploring Correlations between Appearance and Temperature 3997
Daniel Glasner, Pascal Fua, Todd Zickler, and Lihi Zelnik-Manor

O-4B: Motion and Correspondence

SPM-BP: Sped-Up PatchMatch Belief Propagation for Continuous MRFs 4006
Yu Li, Dongbo Min, Michael S. Brown, Minh N. Do, and Jiangbo Lu

Flow Fields: Dense Correspondence Fields for Highly Accurate Large Displacement Optical Flow Estimation 4015
Christian Bailer, Bertram Taetz, and Didier Stricker

Dense Semantic Correspondence Where Every Pixel is a Classifier 4024
Hilton Bristow, Jack Valmadre, and Simon Lucey

Multi-image Matching via Fast Alternating Minimization 4032
Xiaowei Zhou, Menglong Zhu, and Kostas Daniilidis

P-4B: Statistical Methods and Learning, Motion and Tracking, and Video Analysis II

Differential Recurrent Neural Networks for Action Recognition 4041
Vivek Veeriah, Naifan Zhuang, and Guo-Jun Qi

Similarity Gaussian Process Latent Variable Model for Multi-modal Data Analysis 4050
Guoli Song, Shuhui Wang, Qingming Huang, and Qi Tian

Learning Ensembles of Potential Functions for Structured Prediction with Latent Variables 4059
Hossein Hajimirsadeghi and Greg Mori

Simultaneous Deep Transfer Across Domains and Tasks 4068
Eric Tzeng, Judy Hoffman, Trevor Darrell, and Kate Saenko

Low Dimensional Explicit Feature Maps 4077
Ondřej Chum

Unsupervised Learning of Spatiotemporally Coherent Metrics 4086
Ross Goroshin, Joan Bruna, Jonathan Tompson, David Eigen, and Yann LeCun

Multi-label Cross-Modal Retrieval 4094
Viresh Ranjan, Nikhil Rasiwasia, and C. V. Jawahar

Improving Ferns Ensembles by Sparsifying and Quantising Posterior Probabilities 4103
Antonio L. Rodriguez and Vitor Sequeira

Beyond Gauss: Image-Set Matching on the Riemannian Manifold of PDFs 4112
Mehrtash Harandi, Mathieu Salzmann, and Mahsa Baktashmotlagh

Unsupervised Domain Adaptation with Imbalanced Cross-Domain Data 4121
Tzu Ming Harry Hsu, Wei Yu Chen, Cheng-An Hou, Yao-Hung Hubert Tsai, Yi-Ren Yeh, and Yu-Chiang Frank Wang

Secrets of Matrix Factorization: Approximations, Numerics, Manifold Optimization and Random Restarts 4130
Je Hyeong Hong and Andrew Fitzgibbon

Geometry-Aware Deep Transform 4139
Jiaji Huang, Qiang Qiu, Robert Calderbank, and Guillermo Sapiro

Learning Binary Codes for Maximum Inner Product Search 4148
Fumin Shen, Wei Liu, Shaoting Zhang, Yang Yang, and Heng Tao Shen

ML-MG: Multi-label Learning with Missing Labels Using a Mixed Graph 4157
Baoyuan Wu, Siwei Lyu, and Bernard Ghanem

Zero-Shot Learning via Semantic Similarity Embedding 4166
Ziming Zhang and Venkatesh Saligrama

Bayesian Model Adaptation for Crowd Counts 4175
Bo Liu and Nuno Vasconcelos

An NMF Perspective on Binary Hashing 4184
Lopamudra Mukherjee, Sathya N. Ravi, Vamsi K. Ithapu, Tyler Holmes, and Vikas Singh

Multi-view Domain Generalization for Visual Recognition 4193
Li Niu, Wen Li, and Dong Xu

Infinite Feature Selection 4202
Giorgio Roffo, Simone Melzi, and Marco Cristani

Semi-Supervised Zero-Shot Classification with Label Representation Learning 4211
Xin Li, Yuhong Guo, and Dale Schuurmans

A Supervised Low-Rank Method for Learning Invariant Subspaces 4220
Farzad Siyahjani, Ranya Almohsen, Sinan Sabri, and Gianfranco Doretto

Recursive Fréchet Mean Computation on the Grassmannian and Its Applications to Computer Vision 4229
Rudrasis Chakraborty and Baba C. Vemuri

Multi-view Subspace Clustering 4238
Hongchang Gao, Feiping Nie, Xuelong Li, and Heng Huang

Predicting Deep Zero-Shot Convolutional Neural Networks Using Textual Descriptions 4247
Jimmy Lei Ba, Kevin Swersky, Sanja Fidler, and Ruslan Salakhutdinov

Structured Feature Selection 4256
Tian Gao, Ziheng Wang, and Qiang Ji

Conditional High-Order Boltzmann Machine: A Supervised Learning Model for Relation Learning 4265
Yan Huang, Wei Wang, and Liang Wang

Learning Image and User Features for Recommendation in Social Networks 4274
Xue Geng, Hanwang Zhang, Jingwen Bian, and Tat-Seng Chua

Dual-Feature Warping-Based Motion Model Estimation 4283
Shiwei Li, Lu Yuan, Jian Sun, and Long Quan

An Adaptive Data Representation for Robust Point-Set Registration and Merging 4292
Dylan Campbell and Lars Petersson

Local Subspace Collaborative Tracking 4301
Lin Ma, Xiaoqin Zhang, Weiming Hu, Junliang Xing, Jiwen Lu, and Jie Zhou

Learning Spatially Regularized Correlation Filters for Visual Tracking 4310
Martin Danelljan, Gustav Häger, Fahad Shahbaz Khan, and Michael Felsberg

SpeDo: 6 DOF Ego-Motion Sensor Using Speckle Defocus Imaging 4319
Kensei Jo, Mohit Gupta, and Shree K. Nayar

Unsupervised Trajectory Clustering via Adaptive Multi-kernel-Based Shrinkage 4328
Hongteng Xu, Yang Zhou, Weiyao Lin, and Hongyuan Zha

TRIC-track: Tracking by Regression with Incrementally Learned Cascades 4337
Xiaomeng Wang, Michel Valstar, Brais Martinez, Muhammad Haris Khan, and Tony Pridmore

Recurrent Network Models for Human Dynamics 4346
Katerina Fragkiadaki, Sergey Levine, Panna Felsen, and Jitendra Malik

Contour Flow: Middle-Level Motion Estimation by Combining Motion Segmentation and Contour Alignment 4355
Huijun Di, Qingxuan Shi, Feng Lv, Ming Qin, and Yao Lu

FollowMe: Efficient Online Min-Cost Flow Tracking with Bounded Memory and Computation 4364
Philip Lenz, Andreas Geiger, and Raquel Urtasun

Learning to Divide and Conquer for Online Multi-target Tracking 4373
Francesco Solera, Simone Calderara, and Rita Cucchiara

Minimizing Human Effort in Interactive Tracking by Incremental Learning of Model Parameters 4382
Arridhana Ciptadi and James M. Rehg

A Novel Representation of Parts for Accurate 3D Object Detection and Tracking in Monocular Images 4391
Alberto Crivellaro, Mahdi Rad, Yannick Verdie, Kwang Moo Yi, Pascal Fua, and Vincent Lepetit

Linearization to Nonlinear Learning for Visual Tracking 4400
Bo Ma, Hongwei Hu, Jianbing Shen, Yuping Zhang, and Fatih Porikli

Self-Occlusions and Disocclusions in Causal Video Object Segmentation 4408
Yanchao Yang, Ganesh Sundaramoorthi, and Stefano Soatto

Large Displacement 3D Scene Flow with Occlusion Reasoning 4417
Andrei Zanfir and Cristian Sminchisescu

Co-Interest Person Detection from Multiple Wearable Camera Videos 4426
Yuewei Lin, Kareem Abdelfatah, Youjie Zhou, Xiaochuan Fan, Hongkai Yu, Hui Qian, and Song Wang

Sparse Dynamic 3D Reconstruction from Unsynchronized Videos 4435
Enliang Zheng, Dinghuang Ji, Enrique Dunn, and Jan-Michael Frahm

Category-Blind Human Action Recognition: A Practical Recognition System 4444
Wenbo Li, Longyin Wen, Mooi Choo Chuah, and Siwei Lyu

Temporal Subspace Clustering for Human Motion Segmentation 4453
Sheng Li, Kang Li, and Yun Fu

Weakly-Supervised Alignment of Video with Text 4462
P. Bojanowski, R. Lajugie, E. Grave, F. Bach, I. Laptev, J. Ponce, and C. Schmid

Learning Temporal Embeddings for Complex Video Analysis 4471
Vignesh Ramanathan, Kevin Tang, Greg Mori, and Li Fei-Fei

Unsupervised Semantic Parsing of Video Collections 4480
Ozan Sener, Amir R. Zamir, Silvio Savarese, and Ashutosh Saxena

Learning Spatiotemporal Features with 3D Convolutional Networks 4489
Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri

Temporal Perception and Prediction in Ego-Centric Video 4498
Yipin Zhou and Tamara L. Berg

Describing Videos by Exploiting Temporal Structure 4507
Li Yao, Atousa Torabi, Kyunghyun Cho, Nicolas Ballas, Christopher Pal, Hugo Larochelle, and Aaron Courville

Person Re-Identification with Discriminatively Trained Viewpoint Invariant Dictionaries 4516
Srikrishna Karanam, Yang Li, and Richard J. Radke

Storyline Representation of Egocentric Videos with an Applications to Story-Based Search 4525
Bo Xiong, Gunhee Kim, and Leonid Sigal

Sequence to Sequence — Video to Text 4534
Subhashini Venugopalan, Marcus Rohrbach, Jeffrey Donahue, Raymond Mooney, Trevor Darrell, and Kate Saenko

Context Aware Active Learning of Activity Recognition Models 4543
Mahmudul Hasan and Amit K. Roy-Chowdhury

Action Recognition by Hierarchical Mid-Level Action Elements 4552
Tian Lan, Yuke Zhu, Amir Roshan Zamir, and Silvio Savarese

Selecting Relevant Web Trained Concepts for Automated Event Retrieval 4561
Bharat Singh, Xintong Han, Zhe Wu, Vlad I. Morariu, and Larry S. Davis

Beyond Covariance: Feature Representation with Nonlinear Kernel Matrices 4570
Lei Wang, Jianjia Zhang, Luping Zhou, Chang Tang, and Wanqing Li

Multiresolution Hierarchy Co-Clustering for Semantic Segmentation in Sequences with Small Variations 4579
David Varas, Mónica Alfaro, and Ferran Marques

Objects2action: Classifying and Localizing Actions without Any Video Example 4588
Mihir Jain, Jan C. van Gemert, Thomas Mensink, and Cees G. M. Snoek

Human Action Recognition Using Factorized Spatio-Temporal Convolutional Networks 4597
Lin Sun, Kui Jia, Dit-Yan Yeung, and Bertram E. Shi

Bayesian Non-parametric Inference for Manifold Based MoCap Representation 4606
Fabrizio Natola, Valsamis Ntouskos, Marta Sanzari, and Fiora Pirri

Semantic Video Entity Linking Based on Visual Content and Metadata 4615
Yuncheng Li, Xitong Yang, and Jiebo Luo

Love Thy Neighbors: Image Annotation by Exploiting Image Metadata 4624
Justin Johnson, Lamberto Ballan, and Li Fei-Fei

Unsupervised Extraction of Video Highlights via Robust Recurrent Auto-Encoders 4633
Huan Yang, Baoyuan Wang, Stephen Lin, David Wipf, Minyi Guo, and Baining Guo

Learning Visual Clothing Style with Heterogeneous Dyadic Co-Occurrences 4642
Andreas Veit, Balazs Kovacs, Sean Bell, Julian McAuley, Kavita Bala, and Serge Belongie

Text Flow: A Unified Text Detection System in Natural Scene Images 4651
Shangxuan Tian, Yifeng Pan, Chang Huang, Shijian Lu, Kai Yu, and Chew Lim Tan

O-4C: Video: Actions, Surveillance and Tracking

Uncovering Interactions and Interactors: Joint Estimation of Head, Body Orientation and F-Formations from Surveillance Videos 4660
Elisa Ricci, Jagannadan Varadarajan, Ramanathan Subramanian, Samuel Rota Bulò, Narendra Ahuja, and Oswald Lanz

Generating Notifications for Missing Actions: Don't Forget to Turn the Lights Off! 4669
Bilge Soran, Ali Farhadi, and Linda Shapiro

Partial Person Re-Identification 4678
Wei-Shi Zheng, Xiang Li, Tao Xiang, Shengcai Liao, Jianhuang Lai, and Shaogang Gong

Shape Interaction Matrix Revisited and Robustified: Efficient Subspace Clustering with Corrupted and Incomplete Data 4687
Pan Ji, Mathieu Salzmann, and Hongdong Li

Multiple Hypothesis Tracking Revisited 4696
Chanho Kim, Fuxin Li, Arridhana Ciptadi, and James M. Rehg

Learning to Track: Online Multi-object Tracking by Decision Making 4705
Yu Xiang, Alexandre Alahi, and Silvio Savarese

Author Index