

# Table of Contents

## Learning and Inference (*continued*)

Coarse-to-Fine Auto-Encoder Networks (CFAN) for Real-Time Face Alignment . . . . .	1
<i>Jie Zhang, Shiguang Shan, Meina Kan, and Xilin Chen</i>	
From Manifold to Manifold: Geometry-Aware Dimensionality Reduction for SPD Matrices . . . . .	17
<i>Mehrtash T. Harandi, Mathieu Salzmann, and Richard Hartley</i>	
Pose Machines: Articulated Pose Estimation via Inference Machines . . . .	33
<i>Varun Ramakrishna, Daniel Munoz, Martial Hebert, James Andrew Bagnell, and Yaser Sheikh</i>	

## Poster Session 2

Piecewise-Planar StereoScan: Structure and Motion from Plane Primitives . . . . .	48
<i>Carolina Raposo, Michel Antunes, and Joao P. Barreto</i>	
Nonrigid Surface Registration and Completion from RGBD Images . . . . .	64
<i>Weipeng Xu, Mathieu Salzmann, Yongtian Wang, and Yue Liu</i>	
Unsupervised Dense Object Discovery, Detection, Tracking and Reconstruction . . . . .	80
<i>Lu Ma and Gabe Sibley</i>	
Know Your Limits: Accuracy of Long Range Stereoscopic Object Measurements in Practice . . . . .	96
<i>Peter Pinggera, David Pfeiffer, Uwe Franke, and Rudolf Mester</i>	
As-Rigid-As-Possible Stereo under Second Order Smoothness Priors . . . .	112
<i>Chi Zhang, Zhiwei Li, Rui Cai, Hongyang Chao, and Yong Rui</i>	
Real-Time Minimization of the Piecewise Smooth Mumford-Shah Functional . . . . .	127
<i>Evgeny Strelakovskiy and Daniel Cremers</i>	
A MAP-Estimation Framework for Blind Deblurring Using High-Level Edge Priors . . . . .	142
<i>Yipin Zhou and Nikos Komodakis</i>	
Efficient Color Constancy with Local Surface Reflectance Statistics . . . . .	158
<i>Shaobing Gao, Wangwang Han, Kaifu Yang, Chaoyi Li, and Yongjie Li</i>	

A Contrast Enhancement Framework with JPEG Artifacts Suppression . . . . .	174
<i>Yu Li, Fangfang Guo, Robby T. Tan, and Michael S. Brown</i>	
Radial Bright Channel Prior for Single Image Vignetting Correction . . . . .	189
<i>Hojin Cho, Hyunjoon Lee, and Seungyong Lee</i>	
Tubular Structure Filtering by Ranking Orientation Responses of Path Operators . . . . .	203
<i>Odyssée Merveille, Hugues Talbot, Laurent Najman, and Nicolas Passat</i>	
Optimization-Based Artifact Correction for Electron Microscopy Image Stacks . . . . .	219
<i>Samaneh Azadi, Jeremy Maitin-Shepard, and Pieter Abbeel</i>	
Metric-Based Pairwise and Multiple Image Registration . . . . .	236
<i>Qian Xie, Sebastian Kurtek, Eric Klassen, Gary E. Christensen, and Anuj Srivastava</i>	
Canonical Correlation Analysis on Riemannian Manifolds and Its Applications . . . . .	251
<i>Hyunwoo J. Kim, Nagesh Adluru, Barbara B. Bendlin, Sterling C. Johnson, Baba C. Vemuri, and Vikas Singh</i>	
Scalable 6-DOF Localization on Mobile Devices . . . . .	268
<i>Sven Middelberg, Torsten Sattler, Ole Untzelmann, and Leif Kobbelt</i>	
On Mean Pose and Variability of 3D Deformable Models . . . . .	284
<i>Benjamin Allain, Jean-Sébastien Franco, Edmond Boyer, and Tony Tung</i>	
Hybrid Stochastic / Deterministic Optimization for Tracking Sports Players and Pedestrians . . . . .	298
<i>Robert T. Collins and Peter Carr</i>	
What Do I See? Modeling Human Visual Perception for Multi-person Tracking . . . . .	314
<i>Xu Yan, Ioannis A. Kakadiaris, and Shishir K. Shah</i>	
Consistent Re-identification in a Camera Network . . . . .	330
<i>Abir Das, Anirban Chakraborty, and Amit K. Roy-Chowdhury</i>	
Surface Normal Deconvolution: Photometric Stereo for Optically Thick Translucent Objects . . . . .	346
<i>Chika Inoshita, Yasuhiro Mukaigawa, Yasuyuki Matsushita, and Yasushi Yagi</i>	
Intrinsic Video . . . . .	360
<i>Naejin Kong, Peter V. Gehler, and Michael J. Black</i>	

Robust and Accurate Non-parametric Estimation of Reflectance Using Basis Decomposition and Correction Functions . . . . .	376
<i>Tobias Nöll, Johannes Köhler, and Didier Stricker</i>	
Intrinsic Textures for Relightable Free-Viewpoint Video . . . . .	392
<i>James Imber, Jean-Yves Guillemaut, and Adrian Hilton</i>	
Reasoning about Object Affordances in a Knowledge Base Representation . . . . .	408
<i>Yuke Zhu, Alireza Fathi, and Li Fei-Fei</i>	
Binary Codes Embedding for Fast Image Tagging with Incomplete Labels . . . . .	425
<i>Qifan Wang, Bin Shen, Shumiao Wang, Liang Li, and Luo Si</i>	
Recognizing Products: A Per-exemplar Multi-label Image Classification Approach . . . . .	440
<i>Marian George and Christian Floerkemeier</i>	
Part-Pair Representation for Part Localization . . . . .	456
<i>Jiongxin Liu, Yinxiao Li, and Peter N. Belhumeur</i>	
Weakly Supervised Learning of Objects, Attributes and Their Associations . . . . .	472
<i>Zhiyuan Shi, Yongxin Yang, Timothy M. Hospedales, and Tao Xiang</i>	
Interestingness Prediction by Robust Learning to Rank . . . . .	488
<i>Yanwei Fu, Timothy M. Hospedales, Tao Xiang, Shaogang Gong, and Yuan Yao</i>	
Pairwise Probabilistic Voting: Fast Place Recognition without RANSAC . . . . .	504
<i>Edward David Johns and Guang-Zhong Yang</i>	
Robust Instance Recognition in Presence of Occlusion and Clutter . . . . .	520
<i>Ujwal Bonde, Vijay Badrinarayanan, and Roberto Cipolla</i>	
Learning 6D Object Pose Estimation Using 3D Object Coordinates . . . . .	536
<i>Eric Brachmann, Alexander Krull, Frank Michel, Stefan Gumhold, Jamie Shotton, and Carsten Rother</i>	
Growing Regression Forests by Classification: Applications to Object Pose Estimation . . . . .	552
<i>Kota Hara and Rama Chellappa</i>	
Stacked Deformable Part Model with Shape Regression for Object Part Localization . . . . .	568
<i>Junjie Yan, Zhen Lei, Yang Yang, and Stan Z. Li</i>	

Transductive Multi-view Embedding for Zero-Shot Recognition and Annotation . . . . .	584
<i>Yanwei Fu, Timothy M. Hospedales, Tao Xiang, Zhenyong Fu, and Shaogang Gong</i>	
Self-explanatory Sparse Representation for Image Classification . . . . .	600
<i>Bao-Di Liu, Yu-Xiong Wang, Bin Shen, Yu-Jin Zhang, and Martial Hebert</i>	
Efficient $k$ -Support Matrix Pursuit . . . . .	617
<i>Hanjiang Lai, Yan Pan, Canyi Lu, Yong Tang, and Shuicheng Yan</i>	
Geodesic Regression on the Grassmannian . . . . .	632
<i>Yi Hong, Roland Kwitt, Nikhil Singh, Brad Davis, Nuno Vasconcelos, and Marc Niethammer</i>	
Model Selection by Linear Programming . . . . .	647
<i>Joseph Wang, Tolga Bolukbasi, Kirill Trapeznikov, and Venkatesh Saligrama</i>	
Perceptually Inspired Layout-Aware Losses for Image Segmentation . . . .	663
<i>Anton Osokin and Pushmeet Kohli</i>	
Large Margin Local Metric Learning . . . . .	679
<i>Julien Bohné, Yiming Ying, Stéphane Gencic, and Massimiliano Pontil</i>	
Movement Pattern Histogram for Action Recognition and Retrieval . . . .	695
<i>Arridhana Ciptadi, Matthew S. Goodwin, and James M. Rehg</i>	
Pose Filter Based Hidden-CRF Models for Activity Detection . . . . .	711
<i>Prithviraj Banerjee and Ram Nevatia</i>	
Action Recognition Using Super Sparse Coding Vector with Spatio-temporal Awareness . . . . .	727
<i>Xiaodong Yang and YingLi Tian</i>	
HOPC: Histogram of Oriented Principal Components of 3D Pointclouds for Action Recognition . . . . .	742
<i>Hossein Rahmani, Arif Mahmood, Du Q Huynh, and Ajmal Mian</i>	
Natural Action Recognition Using Invariant 3D Motion Encoding . . . . .	758
<i>Simon Hadfield, Karel Lebeda, and Richard Bowden</i>	
Detecting Social Actions of Fruit Flies . . . . .	772
<i>Eyrun Eyjolfsdottir, Steve Branson, Xavier P. Burgos-Artizzu, Eric D. Hoopfer, Jonathan Schor, David J. Anderson, and Pietro Perona</i>	

## Structure from Motion and Feature Matching

Progressive Mode-Seeking on Graphs for Sparse Feature Matching . . . . .	788
<i>Chao Wang, Lei Wang, and Lingqiao Liu</i>	
Globally Optimal Inlier Set Maximization with Unknown Rotation and Focal Length . . . . .	803
<i>Jean-Charles Bazin, Yongduek Seo, Richard Hartley, and Marc Pollefeys</i>	
Match Selection and Refinement for Highly Accurate Two-View Structure from Motion . . . . .	818
<i>Zhe Liu, Pascal Monasse, and Renaud Marlet</i>	
LSD-SLAM: Large-Scale Direct Monocular SLAM . . . . .	834
<i>Jakob Engel, Thomas Schöps, and Daniel Cremers</i>	
<b>Author Index</b> . . . . .	851