

2019 IEEE/CVF International Conference on Computer Vision (ICCV) ICCV 2019

Table of Contents

Message from the General and Program Chairs	cx
2019 Organizing Committee	cxi
2019 Area Chairs	cxii
2019 Reviewers	cxiii

Poster 1.1

Deep Learning

FaceForensics++: Learning to Detect Manipulated Facial Images	1
<i>Andreas Rössler (Technical University of Munich), Davide Cozzolino (University Federico II of Naples), Luisa Verdoliva (University Federico II of Naples), Christian Riess (Friedrich-Alexander University Erlangen-Nuremberg), Justus Thies (Technical University of Munich), and Matthias Niessner (Technical University of Munich)</i>	
DeepVCP: An End-to-End Deep Neural Network for Point Cloud Registration	12
<i>Weixin Lu (Baidu ADU), Guowei Wan (Baidu Company), Yao Zhou (Baidu ADT), Xiangyu Fu (Baidu), Pengfei Yuan (Baidu), and Shiyu Song (Baidu ADU)</i>	
Shape Reconstruction Using Differentiable Projections and Deep Priors	22
<i>Matheus Gadelha (University of Massachusetts Amherst), Rui Wang (University of Massachusetts), and Subhransu Maji (University of Massachusetts Amherst)</i>	
Fine-Grained Segmentation Networks: Self-Supervised Segmentation for Improved Long-Term Visual Localization	31
<i>Måns Larsson (Chalmers), Erik Stenborg (Chalmers University), Carl Toft (Chalmers), Lars Hammarstrand (Chalmers university of technology), Torsten Sattler (Chalmers University of Technology), and Fredrik Kahl (Chalmers)</i>	
SANet: Scene Agnostic Network for Camera Localization	42
<i>Luwei Yang (Simon Fraser University), Ziqian Bai (Simon Fraser University), Chengzhou Tang (Simon Fraser University), Honghua Li (Alibaba A.I. Labs), Yasutaka Furukawa (Simon Fraser University), and Ping Tan (Simon Fraser University)</i>	

Total Denoising: Unsupervised Learning of 3D Point Cloud Cleaning	52
<i>Pedro Hermosilla Casajus (Ulm University), Tobias Ritschel (UCL), and Timo Ropinski (Ulm University)</i>	
Hierarchical Self-Attention Network for Action Localization in Videos	61
<i>Rizard Renanda Adhi Pramono (National Taiwan University of Science and Technology), Yie-Tarnng Chen (National Taiwan University of Science and Technology), and Wen-Hsien Fang (National Taiwan University of Science and Technology)</i>	
Goal-Driven Sequential Data Abstraction	71
<i>Umar Riaz Muhammad (Queen Mary University of London), Yongxin Yang (University of Surrey), Timothy Hospedales (Edinburgh University), Tao Xiang (University of Surrey), and Yi-Zhe Song (University of Surrey)</i>	
Jointly Aligning Millions of Images With Deep Penalised Reconstruction Congealing	81
<i>Roberto Annunziata (Onfido), Christos Sagonas (Onfido), and Jacques Cali (Blue Prism)</i>	
Drop to Adapt: Learning Discriminative Features for Unsupervised Domain Adaptation	91
<i>Seungmin Lee (Seoul National University), Dongwan Kim (Seoul National University), Namil Kim (NAVER LABS), and Seong-Gyun Jeong (code42.ai)</i>	
NLNL: Negative Learning for Noisy Labels	101
<i>Youngdong Kim (KAIST), Junho Yim (KAIST), Juseung Yun (KAIST), and Junmo Kim (KAIST)</i>	
Adversarial Robustness vs. Model Compression, or Both?	111
<i>Shaokai Ye (Tsinghua University), Kaidi Xu (Northeastern University), Sijia Liu (IBM), Hao Cheng (Xi'an Jiaotong University), Jan-Henrik Lambrechts (Tsinghua University), Huan Zhang (UCLA), Aojun Zhou (SenseTime), Kaisheng Ma (Tsinghua University), Yanzhi Wang (Northeastern University), and Xue Lin (Northeastern University)</i>	
On the Design of Black-Box Adversarial Examples by Leveraging Gradient-Free Optimization and Operator Splitting Method	121
<i>Pu Zhao (Northeastern University), Sijia Liu (IBM), Pin-Yu Chen (IBM Research), Nghia Hoang (IBM Research), Kaidi Xu (Northeastern University), Bhavya Kailkhura (LLNL), and Xue Lin (Northeastern University)</i>	
DewarpNet: Single-Image Document Unwarping With Stacked 3D and 2D Regression Networks	131
<i>Sagnik Das (Stony Brook University), Ke Ma (Stony Brook University), Zhixin Shu (Stony Brook University), Dimitris Samaras (Stony Brook University), and Roy Shilkrot (Stony Brook University)</i>	
Learning Robust Facial Landmark Detection via Hierarchical Structured Ensemble	141
<i>Xu Zou (Huazhong University of Science and Technology), Sheng Zhong (Huazhong University of Science and Technology), Luxin Yan (Huazhong University of Science and Technology), Xiangyun Zhao (Northwestern University), Jiahuan Zhou (Northwestern University), and Ying Wu (Northwestern University)</i>	
Remote Heart Rate Measurement From Highly Compressed Facial Videos: An End-to-End Deep Learning Solution With Video Enhancement	151
<i>Zitong Yu (CMVS. University of Oulu), Wei Peng (CMVS. University of Oulu), Xiaobai Li (University of Oulu), Xiaopeng Hong (Xi'an Jiaotong University), and Guoying Zhao (University of Oulu)</i>	

Face-to-Parameter Translation for Game Character Auto-Creation	161
<i>Tianyang Shi (NetEase Fuxi AI Lab), Yi Yuan (NetEase Fuxi AI Lab), Changjie Fan (NetEase Fuxi AI Lab), Zhengxia Zou (University of Michigan), Zhenwei Shi (Beihang University), and Yong Liu (Zhejiang University)</i>	
Visual Deprojection: Probabilistic Recovery of Collapsed Dimensions	171
<i>Guha Balakrishnan (MIT), Adrian Dalca (MIT), Amy Zhao (MIT), John Guttag (MIT), Fredo Durand (MIT), and William Freeman (MIT)</i>	
StructureFlow: Image Inpainting via Structure-Aware Appearance Flow	181
<i>Yurui Ren (Shenzhen Graduate School. Peking University), Xiaoming Yu (Peking University), Ruonan Zhang (Peng Cheng LaboratoryArtificial intelligence Research Center), Thomas H Li (Advanced Institute of Information Technology. Peking University), Shan Liu (Tencent America), and Ge Li (SECE. Shenzhen Graduate School. Peking University)</i>	
Learning Fixed Points in Generative Adversarial Networks: From Image-to-Image Translation to Disease Detection and Localization	191
<i>Md Mahfuzur Rahman Siddiquee (Arizona State University), Zongwei Zhou (Arizona State University), Nima Tajbakhsh (Arizona State University), Ruibin Feng (Arizona State University), Michael Gotway (Mayo Clinic), Yoshua Bengio (Mila), and Jianming Liang (Arizona State University. USA)</i>	
Generative Adversarial Training for Weakly Supervised Cloud Matting	201
<i>Zhengxia Zou (University of Michigan), Wenyuan Li (Beihang University), Tianyang Shi (NetEase Fuxi AI Lab), Zhenwei Shi (Beihang University), and Jieping Ye (Didi Chuxing & University of Michigan)</i>	
PAMTRI: Pose-Aware Multi-Task Learning for Vehicle Re-Identification Using Highly Randomized Synthetic Data	211
<i>Zheng Tang (University of Washington), Milind Naphade (NVidia), Stan Birchfield (NVIDIA), Jonathan Tremblay (NVIDIA), William Hodge (NVIDIA), Ratnesh Kumar (NVIDIA), Shuo Wang (NVidia), and Xiaodong Yang (NVIDIA Research)</i>	
Generative Adversarial Networks for Extreme Learned Image Compression	221
<i>Eirikur Agustsson (ETH Zurich), Michael Tschannen (Google Brain), Fabian Mentzer (ETH Zurich), Radu Timofte (ETH Zurich), and Luc Van Gool (ETH Zurich)</i>	
Instance-Guided Context Rendering for Cross-Domain Person Re-Identification	232
<i>Yanbei Chen (Queen Mary University of London), Xiatian Zhu (Vision Semantics Limited), and Shaogang Gong (Queen Mary University of London)</i>	
What Else Can Fool Deep Learning? Addressing Color Constancy Errors on Deep Neural Network Performance	243
<i>Mahmoud Afifi (York University) and Michael Brown (York University)</i>	
Beyond Cartesian Representations for Local Descriptors	253
<i>Patrick Ebel (EPFL), Eduard Trulls (Google), Kwang Moo Yi (University of Victoria), Pascal Fua (EPFL. Switzerland), and Anastasiia Mishchuk (EPFL)</i>	

Distilling Knowledge From a Deep Pose Regressor Network	263
<i>Muhamad Risqi U. Saputra (University of Oxford), Pedro Gusmao (University of Oxford), Yasin Almalioglu (University of Oxford. Department of Computer Science), Andrew Markham (University of Oxford), and Niki Trigoni (University of Oxford)</i>	
Instance-Level Future Motion Estimation in a Single Image Based on Ordinal Regression	273
<i>Kyung-Rae Kim (Korea University), Whan Choi (Korea University), Yeong Jun Koh (Chungnam National University), Seong-Gyun Jeong (code42.ai), and Chang-Su Kim (Korea university)</i>	
Vision-Infused Deep Audio Inpainting	283
<i>Hang Zhou (The Chinese University of Hong Kong), Ziwei Liu (The Chinese University of Hong Kong), Xudong Xu (the Chinese University of Hong Kong), Ping Luo (The Chinese University of Hong Kong), and Xiaogang Wang (Chinese University of Hong Kong. Hong Kong)</i>	
HAWQ: Hessian AWARE Quantization of Neural Networks With Mixed-Precision	293
<i>Zhen Dong (UC Berkeley), Zhewei Yao (University of California. Berkeley), Amir Gholami (UC Berkeley), Michael Mahoney (University of California. Berkeley), and Kurt Keutzer (EECS. UC Berkeley)</i>	
Evaluating Robustness of Deep Image Super-Resolution Against Adversarial Attacks	303
<i>Jun-Ho Choi (Yonsei University), Huan Zhang (UCLA), Jun-Hyuk Kim (Yonsei University), Cho-Jui Hsieh (UCLA), and Jong-Seok Lee (Yonsei University. Korea)</i>	
Overcoming Catastrophic Forgetting With Unlabeled Data in the Wild	312
<i>Kibok Lee (University of Michigan), Kimin Lee (KAIST), Jinwoo Shin (KAIST), and Honglak Lee (University of Michigan Ann Arbor. USA)</i>	
Symmetric Cross Entropy for Robust Learning With Noisy Labels	322
<i>Yisen Wang (JD AI Research), Xingjun Ma (The University of Melbourne), Zaiyi Chen (University of Science and Technology of China), Yuan Luo (Shanghai Jiao Tong University), Jinfeng Yi (JD AI Research), and James Bailey (THE UNIVERSITY OF MELBOURNE)</i>	
Few-Shot Learning With Embedded Class Models and Shot-Free Meta Training	331
<i>Avinash Ravichandran (Amazon), Rahul Bhotika (Amazon), and Stefano Soatto (UCLA)</i>	
Dual Directed Capsule Network for Very Low Resolution Image Recognition	340
<i>Maneet Singh (IIIT-Delhi. India), Shruti Nagpal (IIIT-Delhi), Richa Singh (IIIT-Delhi), and Mayank Vatsa (IIIT-Delhi)</i>	
Recognizing Part Attributes With Insufficient Data	350
<i>Xiangyun Zhao (Northwestern University), Yi Yang (Baidu), Feng Zhou (Baidu Inc.), Xiao Tan (Baidu Inc.), Yuchen Yuan (Baidu Inc.), Yingze Bao (Baidu), and Ying Wu (Northwestern University)</i>	
USIP: Unsupervised Stable Interest Point Detection From 3D Point Clouds	361
<i>Jiixin Li (National University of Singapore) and Gim Hee Lee (National University of Singapore)</i>	
Mixed High-Order Attention Network for Person Re-Identification	371
<i>Binghui Chen (BUPT), Weihong Deng (Beijing University of Posts and Telecommunications), and Jiani Hu (Beijing University of Posts and Telecommunications)</i>	

Budget-Aware Adapters for Multi-Domain Learning	382
<i>Rodrigo Berriel (UFES), Stephane Lathuillere (University of Trento), Moin Nabi (SAP), Tassilo Klein (SAP), Thiago Oliveira-Santos (UFES), Nicu Sebe (University of Trento), and Elisa Ricci (FBK - Technologies of Vision)</i>	
Compact Trilinear Interaction for Visual Question Answering	392
<i>Tuong Do (AIOZ), Huy Tran (AIOZ), Thanh-Toan Do (The University of Liverpool), Erman Tjiputra (AIOZ), and Quang Tran (AIOZ)</i>	
Towards Latent Attribute Discovery From Triplet Similarities	402
<i>Ishan Nigam (CMU), Pavel Tokmakov (CMU), and Deva Ramanan (Carnegie Mellon University)</i>	
GeoStyle: Discovering Fashion Trends and Events	411
<i>Utkarsh Mall (Cornell University), Kevin Matzen (Facebook), Bharath Hariharan (Cornell University), Noah Snavely (Cornell University and Google AI), and Kavita Bala (Cornell)</i>	
Towards Adversarially Robust Object Detection	421
<i>Haichao Zhang (Baidu Research) and Jianyu Wang (Baidu Research USA)</i>	

Recognition

Automatic and Robust Skull Registration Based on Discrete Uniformization	431
<i>Junli Zhao (Qingdao University), Xin Qi (Stony Brook University), Chengfeng Wen (Stony Brook University), Na Lei (Dalian University of Technology), and Xianfeng Gu (Stony Brook University)</i>	
Few-Shot Image Recognition With Knowledge Transfer	441
<i>Zhimao Peng (Nanjing University of Science and Technology), Zechao Li (Nanjing University of Science and Technology), Junge Zhang (CASIA), Yan Li (Institute of Automation, Chinese Academy of Sciences), Guo-Jun Qi (Futurewei Technologies), and Jinhui Tang (Nanjing University of Science and Technology)</i>	
Fine-Grained Action Retrieval Through Multiple Parts-of-Speech Embeddings	450
<i>Michael Wray (University of Bristol), Gabriela Csurka (Naver Labs Europe), Diane Larlus (Naver Labs Europe), and Dima Damen (University of Bristol)</i>	
Vehicle Re-Identification in Aerial Imagery: Dataset and Approach	460
<i>Peng Wang (Northwestern Polytechnical University), Bingliang Jiao (Northwestern Polytechnical University), Lu Yang (Northwestern Polytechnical University), Yifei Yang (Northwestern Polytechnical University), Shizhou Zhang (NorthWestern Polytechnical University), Wei Wei (Northwestern Polytechnical University), and Yanning Zhang (Northwestern Polytechnical University)</i>	
Bridging the Domain Gap for Ground-to-Aerial Image Matching	470
<i>Krishna Regmi (University of Central Florida) and Mubarak Shah (University of Central Florida)</i>	
A Robust Learning Approach to Domain Adaptive Object Detection	480
<i>Mehran Khodabandeh (Simon Fraser University), Arash Vahdat (NVIDIA), Mani Ranjbar (Quadrant AI), and William Macready (Quadrant AI)</i>	

Graph-Based Object Classification for Neuromorphic Vision Sensing	491
<i>Yin Bi (University College London), Aaron Chadha (University College London), Alhabib Abbas (University College London), Eirina Bourtsoulatzé (University College London), and Yiannis Andreopoulos (University College London)</i>	
Gaussian YOLOv3: An Accurate and Fast Object Detector Using Localization Uncertainty for Autonomous Driving	502
<i>Jiwoong Choi (Seoul National University), Dayoung Chun (Seoul National University), Hyun Kim (Seoul National University of Science and Technology), and Hyuk-Jae Lee (SNU)</i>	
Sharpen Focus: Learning With Attention Separability and Consistency	512
<i>Lezi Wang (Rutgers), Ziyang Wu (Siemens Corporation), Srikrishna Karanam (Siemens Corporate Technology, Princeton), Kuan-Chuan Peng (Siemens Corporation), Rajat Vikram Singh (Siemens Corporation), Bo Liu (Rutgers University), and Dimitris Metaxas (Rutgers)</i>	
Learning Semantic-Specific Graph Representation for Multi-Label Image Recognition	522
<i>Tianshui Chen (Sun Yat-Sen University), Muxin Xu (Sun Yat-sen University), Xiaolu Hui (Sun Yat-sen University), Hefeng Wu (Sun Yat-sen University; Guangdong University of Foreign Studies), and Liang Lin (Sun Yat-sen University)</i>	
DeceptionNet: Network-Driven Domain Randomization	532
<i>Sergey Zakharov (TU Munich / Siemens Corporate Technology), Wadim Kehl (Toyota Research Institute), and Slobodan Ilic (TUM)</i>	
Pose-Guided Feature Alignment for Occluded Person Re-Identification	542
<i>Jiaxu Miao (University of Technology Sydney), Yu Wu (University of Technology Sydney), Ping Liu (UTS), Yuhang Ding (University of Technology, Sydney), and Yi Yang (UTS)</i>	
Robust Person Re-Identification by Modelling Feature Uncertainty	552
<i>Tianyuan Yu (University of Surrey), Da Li (QMUL), Yongxin Yang (University of Edinburgh), Timothy Hospedales (Edinburgh University), and Tao Xiang (University of Surrey)</i>	
Co-Segmentation Inspired Attention Networks for Video-Based Person Re-Identification	562
<i>Arulkumar Subramaniam (Indian Institute of Technology Madras), Athira Nambiar (Indian Institute of Technology Madras), and Anurag Mittal (Indian Institute of Technology Madras)</i>	
A Delay Metric for Video Object Detection: What Average Precision Fails to Tell	573
<i>Huizi Mao (stanford university), Xiaodong Yang (NVIDIA Research), and Bill Dally (stanford university)</i>	
IL2M: Class Incremental Learning With Dual Memory	583
<i>Eden Belouadah (CEA LIST) and Adrian Popescu (CEA LIST)</i>	

Segmentation, Grouping, & Shape

Asymmetric Non-Local Neural Networks for Semantic Segmentation	593
<i>Zhen Zhu (Huazhong University of Science and Technology), Mengdu Xu (Huazhong University of Science and Technology), Song Bai (University of Oxford), Tengpeng Huang (Huazhong University of Science and Technology), and Xiang Bai (Huazhong University of Science and Technology)</i>	
CCNet: Criss-Cross Attention for Semantic Segmentation	603
<i>Zilong Huang (Huazhong Univ. of Science and Technology), Xinggang Wang (Huazhong University of Science and Technology), Lichao Huang (Horizon Robotics), Chang Huang (Horizon Robotics), Yunchao Wei (UIUC), and Wenyu Liu (Huazhong University of Science and Technology)</i>	
Convex Shape Prior for Multi-Object Segmentation Using a Single Level Set Function	613
<i>Shousheng Luo (Beijing Computational Science Research Center), Xue-Cheng Tai (Hong Kong Baptist University), Limei Huo (Henan University), Yang Wang (HKUST), and Roland Glowinski (University of Houston)</i>	
Feature Weighting and Boosting for Few-Shot Segmentation	622
<i>Khoi Nguyen (Oregon State University) and Sinisa Todorovic (Oregon State U)</i>	
Surface Networks via General Covers	632
<i>Niv Haim (Weizmann Institute of Science), Nimrod Segol (Weizmann Institute of Science), Heli Ben-Hamu (Weizmann Institute of Science), Haggai Maron (Weizmann Institute of Science), and Yaron Lipman (Weizmann Institute of Science)</i>	
SSAP: Single-Shot Instance Segmentation With Affinity Pyramid	642
<i>Naiyu Gao (Institute of Automation. Chinese Academy of Sciences; University of Chinese Academy of Sciences), Yanhu Shan (Horizon Robotics), Yupei Wang (Institute of Automation. Chinese Academy of Sciences; University of Chinese Academy of Sciences), Xin Zhao (Institute of Automation. Chinese Academy of Sciences), Yinan Yu (Horizon Robotics), Ming Yang (Horizon Robotics), and Kaiqi Huang (Institute of Automation. Chinese Academy of Sciences)</i>	
Learning Propagation for Arbitrarily-Structured Data	652
<i>Sifei Liu (NVIDIA), Xueting Li (University of California Merced), Varun Jampani (Nvidia Research), Shalini Mello (NVIDIA Research), and Jan Kautz (NVIDIA)</i>	
MultiSeg: Semantically Meaningful, Scale-Diverse Segmentations From Minimal User Input	662
<i>Jun Hao Liew (NUS), Scott Cohen (Adobe Research), Brian Price (Adobe), Long Mai (Adobe Research), Sim-Heng Ong (NUS), and Jiashi Feng (NUS)</i>	
Robust Motion Segmentation From Pairwise Matches	671
<i>Federica Arrigoni (Czech Technical University in Prague) and Tomas Pajdla (Czech Technical University in Prague)</i>	
InstaBoost: Boosting Instance Segmentation via Probability Map Guided Copy-Pasting	682
<i>Hao-Shu Fang (SJTU), Jianhua Sun (Shanghai Jiaotong Univ), Runzhong Wang (Shanghai Jiao Tong University), Minghao Gou (Shanghai Jiao Tong University), Yong-Lu Li (Shanghai Jiao Tong University), and Cewu Lu (Shanghai Jiao Tong University)</i>	

Face & Body

Racial Faces in the Wild: Reducing Racial Bias by Information Maximization Adaptation Network	692
<i>Mei Wang (Beijing University of Posts and Telecommunications), Weihong Deng (Beijing University of Posts and Telecommunications), Jiani Hu (Beijing University of Posts and Telecommunications), Xunqiang Tao (Canon Information Technology. Beijing Co.. Ltd), and Yaohai Huang (Canon Information Technology. Beijing Co.. Ltd)</i>	
Uncertainty Modeling of Contextual-Connections Between Tracklets for Unconstrained Video-Based Face Recognition	703
<i>Jingxiao Zheng (University of Maryland. College Park), Ruichi Yu (Waymo LLC.), Jun-Cheng Chen (University of Maryland), Boyu Lu (University of Maryland), Carlos Castillo (University of Maryland), and Rama Chellappa (University of Maryland)</i>	
Spatio-Temporal Fusion Based Convolutional Sequence Learning for Lip Reading	713
<i>Xingxuan Zhang (Shanghai Jiao Tong University), Feng Cheng (Shanghai Jiao Tong University), and Wang Shilin (sjtu)</i>	
Occlusion-Aware Networks for 3D Human Pose Estimation in Video	723
<i>Yu Cheng (National University of Singapore), Bo Yang (Tencent America), Bo Wang (Tencent America), Yan Wending (National University of Singapore), and Robby Tan (Yale-NUS College. Singapore)</i>	
Context-Aware Feature and Label Fusion for Facial Action Unit Intensity Estimation With Partially Labeled Data	733
<i>Yong Zhang (Tencent AI Lab), Haiyong Jiang (NTU), Baoyuan Wu (Tencent AI Lab), Yanbo Fan (Tencent AI Lab), and Qiang Ji (Rensselaer Polytechnic Institute)</i>	
Distill Knowledge From NRSfM for Weakly Supervised 3D Pose Learning	743
<i>Chaoyang Wang (Carnegie Mellon University), Chen Kong (CMU), and Simon Lucey (CMU)</i>	
MONET: Multiview Semi-Supervised Keypoint Detection via Epipolar Divergence	753
<i>Yuan Yao (University of Minnesota), Yasamin Jafarian (University of Minnesota), and Hyun Soo Park (The University of Minnesota)</i>	
Talking With Hands 16.2M: A Large-Scale Dataset of Synchronized Body-Finger Motion and Audio for Conversational Motion Analysis and Synthesis	763
<i>Gilwoo Lee (University of Washington), Zhiwei Deng (Simon Fraser University), Shugao Ma (Facebook), Takaaki Shiratori (Facebook Reality Labs), Siddhartha Srinivasa (University of Washington), and Yaser Sheikh (Facebook Reality Labs)</i>	
Occlusion Robust Face Recognition Based on Mask Learning With Pairwise Differential Siamese Network	773
<i>Lingxue Song (Tsinghua University), Dihong Gong (Tencent AI Lab), Zhifeng Li (Tencent AI Lab), Changsong Liu (Tsinghua University), and Wei Liu (Tencent)</i>	
Teacher Supervises Students How to Learn From Partially Labeled Images for Facial Landmark Detection.....	783
<i>Xuanyi Dong (UTS) and Yi Yang (UTS)</i>	

A2J: Anchor-to-Joint Regression Network for 3D Articulated Pose Estimation From a Single Depth Image.....	793
<i>Fu Xiong (Huazhong University of Science and Technology), Boshen Zhang (School of Artificial Intelligence and Automation. HUST), Yang Xiao (Huazhong Univ. of Sci.&Tech.), Zhiguo Cao (Huazhong Univ. of Sci.&Tech.), Taidong Yu (Huazhong University of Science and Technology), Joey Tianyi Zhou (IHPC. ASTAR), and Junsong Yuan (State University of New York at Buffalo. USA)</i>	
TexturePose: Supervising Human Mesh Estimation With Texture Consistency	803
<i>Georgios Pavlakos (University of Pennsylvania), Nikos Kolotouros (University of Pennsylvania), and Kostas Daniilidis (University of Pennsylvania)</i>	
FreiHAND: A Dataset for Markerless Capture of Hand Pose and Shape From Single RGB Images	813
<i>Christian Zimmermann (Computer Vision Group - University Of Freiburg), Duygu Ceylan (Adobe Research), Jimei Yang (Adobe), Bryan Russell (Adobe Research), Max J. Argus (University Of Freiburg), and Thomas Brox (University of Freiburg)</i>	
Markerless Outdoor Human Motion Capture Using Multiple Autonomous Micro Aerial Vehicles	823
<i>Nitin Saini (Max Planck Institute for Intelligent Systems), Eric Price (Max Planck Institute for Intelligent Systems), Rahul Tallamraju (Max Planck Institute for Intelligent Systems), Raffi Enficiaud (Max Planck Institute for Intelligent Systems), Roman Ludwig (Max Planck Institute for Intelligent Systems), Igor Martinovic (Max Planck Institute for Intelligent Systems), Aamir Ahmad (Max Planck Institute for Intelligent Systems), and Michael Black (Max Planck Institute for Intelligent Systems)</i>	

Action & Video

Toyota Smarthome: Real-World Activities of Daily Living	833
<i>Srijan Das (INRIA), Rui Dai (INRIA), Michal Koperski (INRIA), Luca Minciullo (Toyota-Europe), Lorenzo Garattoni (Toyota-Europe), Francois Bremond (Inria Sophia Antipolis. France), and Gianpiero Francesca (Toyota-Europe)</i>	
Relation Parsing Neural Network for Human-Object Interaction Detection	843
<i>Penghao Zhou (Fudan university) and Mingmin Chi (Fudan university)</i>	
DistInit: Learning Video Representations Without a Single Labeled Video	852
<i>Rohit Girdhar (Carnegie Mellon University), Du Tran (Facebook Research), Lorenzo Torresani (Facebook AI Research), and Deva Ramanan (Carnegie Mellon University)</i>	
Zero-Shot Anticipation for Instructional Activities	862
<i>Fadime Sener (University of Bonn) and Angela Yao (National University of Singapore)</i>	
Making the Invisible Visible: Action Recognition Through Walls and Occlusions	872
<i>Tianhong Li (MIT), Lijie Fan (MIT), Mingmin Zhao (MIT), Yingcheng Liu (MIT), and Dina Katabi (Massachusetts Institute of Technology)</i>	

Recursive Visual Sound Separation Using Minus-Plus Net	882
<i>Xudong Xu (the Chinese University of Hong Kong), Bo Dai (the Chinese University of Hong Kong), and Dahua Lin (The Chinese University of Hong Kong)</i>	

Motion & Tracking

Unsupervised Video Interpolation Using Cycle Consistency	892
<i>Fitsum Reda (NVIDIA), Deqing Sun (Google), Aysegul Dundar (NVIDIA), Mohammad Shoeybi (NVIDIA), Guilin Liu (NVIDIA), Kevin Shih (NVIDIA), Andrew Tao (NVIDIA), Jan Kautz (NVIDIA), and Bryan Catanzaro (NVIDIA)</i>	
Deformable Surface Tracking by Graph Matching	901
<i>Tao Wang (Beijing Jiaotong University), Haibin Ling (Temple University), Congyan Lang (Beijing Jiaotong University), Songhe Feng (School of Computer and Information Technology, Beijing Jiaotong University), and Xiaohui Hou (HiScene)</i>	
Deep Meta Learning for Real-Time Target-Aware Visual Tracking	911
<i>Janghoon Choi (Seoul National University), Junseok Kwon (Chung-Ang Univ., Korea), and Kyoung Mu Lee (Seoul National University)</i>	
Looking to Relations for Future Trajectory Forecast	921
<i>Chiho Choi (Honda Research Institute US) and Behzad Dariush (Honda Research Institute US)</i>	
Anchor Diffusion for Unsupervised Video Object Segmentation	931
<i>Zhao Yang (University of Oxford), Qiang Wang (CASIA), Luca Bertinetto (University of Oxford), Song Bai (University of Oxford), Weiming Hu (Institute of Automation Chinese Academy of Sciences), and Philip Torr (University of Oxford)</i>	
Tracking Without Bells and Whistles	941
<i>Philipp Bergmann (Technical University of Munich), Tim Meinhardt (TUM), and Laura Leal-Taixé (TUM)</i>	

Scene Understanding

Perspective-Guided Convolution Networks for Crowd Counting	952
<i>Zhaoyi Yan (Harbin Institute of Technology), Yuchen Yuan (Baidu Inc.), Wangmeng Zuo (Harbin Institute of Technology, China), Xiao Tan (Baidu Inc.), Yezhen Wang (Harbin Institute of Technology), Shilei Wen (Baidu Research), and Errui Ding (Baidu Inc.)</i>	
End-to-End Wireframe Parsing	962
<i>Yichao Zhou (UC Berkeley), Haozhi Qi (UC Berkeley), and Yi Ma (UC Berkeley)</i>	
Incremental Class Discovery for Semantic Segmentation With RGBD Sensing	972
<i>Yoshikatsu Nakajima (Keio University), Byeongkeun Kang (Seoul National University of Science and Technology), Hideo Saito (Keio University), and Kris Kitani (CMU)</i>	

SSF-DAN: Separated Semantic Feature Based Domain Adaptation Network for Semantic Segmentation	982
<i>Liang Du (Fudan University), Jingang Tan (SIMIT), Hongye Yang (SIMIT), Jianfeng Feng (Fudan University), Xiangyang Xue (Fudan University), Qibao Zheng (Fudan University), Xiaoqing Ye (SIMIT), and Xiaolin Zhang (SIMIT)</i>	
SpaceNet MVOI: A Multi-View Overhead Imagery Dataset	992
<i>Nicholas Weir (CosmiQ Works. In-Q-Tel), David Lindenbaum (CosmiQ Works. In-Q-Tel), Alexei Bastidas (Intel AI Lab), Adam Van Etten (CosmiQ Works. In-Q-Tel), Varun Kumar (Intel AI Lab), Sean Mcpherson (Intel AI Lab), Jacob Shermeyer (CosmiQ Works. In-Q-Tel), and Hanlin Tang (Intel Corporation)</i>	
Multi-Level Bottom-Top and Top-Bottom Feature Fusion for Crowd Counting	1002
<i>Vishwanath Sindagi (Johns Hopkins University) and Vishal Patel (Johns Hopkins University)</i>	
Learning Lightweight Lane Detection CNNs by Self Attention Distillation	1013
<i>Yuenan Hou (The Chinese University of Hong Kong), Zheng Ma (SenseTime Research), Chunxiao Liu (SenseTime Research), and Chen Change Loy (Nanyang Technological University)</i>	
SplitNet: Sim2Sim and Task2Task Transfer for Embodied Visual Navigation	1022
<i>Daniel Gordon (University of Washington), Abhishek Kadian (Facebook AI Research), Devi Parikh (Georgia Tech & Facebook AI Research), Judy Hoffman (Facebook AI Research and Georgia Tech), and Dhruv Batra (Georgia Tech & Facebook AI Research)</i>	

3D From Multiview & Sensors

Cascaded Parallel Filtering for Memory-Efficient Image-Based Localization	1032
<i>Wentao Cheng (Nanyang Technological University), Weisi Lin (Nanyang Technological University. Singapore), Kan Chen (Fraunhofer Singapore), and Xinfeng Zhang (City University of Hong Kong)</i>	
Pixel2Mesh++: Multi-View 3D Mesh Generation via Deformation	1042
<i>Chao Wen (Fudan University), Yinda Zhang (Google LLC), Zhuwen Li (Nuro), and Yanwei Fu (Fudan University)</i>	
A Differential Volumetric Approach to Multi-View Photometric Stereo	1052
<i>Fotios Logothetis (cambridge university), Roberto Mecca (Toshiba Research Europe Ltd), and Roberto Cipolla (University of Cambridge)</i>	
Revisiting Radial Distortion Absolute Pose	1062
<i>Viktor Larsson (ETH Zurich), Torsten Sattler (Chalmers University of Technology), Zuzana Kukelova (Czech Technical University in Prague), and Marc Pollefeys (ETH Zurich / Microsoft)</i>	
Estimating the Fundamental Matrix Without Point Correspondences With Application to Transmission Imaging	1072
<i>Tobias Wuerfl (Friedrich-Alexander-Universität Erlangen-Nürnberg), Andre Aichert (Siemens Healthineers), Nicole Maass (Siemens Healthineers), Frank Dennerlein (Siemens Healthineers), and Andreas Maier (Pattern Recognition Lab. FAU Erlangen-Nuremberg)</i>	

QUARCH: A New Quasi-Affine Reconstruction Stratum From Vague Relative Camera Orientation Knowledge 1082

Devesh Adlakha (University of Strasbourg), Adlane Habed (University of Strasbourg), Fabio Morbidi (University of Picardie Jules vernes), Cedric Demonceaux (Univ. Bourgogne Franche-Comte. France), and Michel De Mathelin (University of Strasbourg)

Homography From Two Orientation- and Scale-Covariant Features 1091
Dániel Baráth (MTA SZTAKI. CMP Prague) and Zuzana Kukelova (Czech Technical University in Prague)

Applications. Medical, & Robotics

Hiding Video in Audio via Reversible Generative Models 1100
Hyukryul Yang (HKUST), Hao Ouyang (HKUST), Vladlen Koltun (Intel Labs), and Qifeng Chen (HKUST)

GSLAM: A General SLAM Framework and Benchmark 1110
Yong Zhao (Northwestern Polytechnical University), Shibiao Xu (Institute of Automation. Chinese Academy of Sciences), Shuhui Bu (Northwestern Polytechnical University), Hongkai Jiang (Northwestern Polytechnical University), and Pengcheng Han (Northwestern Polytechnical University)

Elaborate Monocular Point and Line SLAM With Robust Initialization 1121
Sang Jun Lee (Handong Global University) and Sung Soo Hwang (Handong Global University)

Adaptive Density Map Generation for Crowd Counting 1130
Jia Wan (City University of Hong Kong) and Antoni Chan (City University of Hong Kong. Hong. Kong)

Attention-Aware Polarity Sensitive Embedding for Affective Image Retrieval 1140
Xingxu Yao (Nankai University), Dongyu She (Nankai University), Sicheng Zhao (University of California Berkeley), Jie Liang (Nankai University), Yu-Kun Lai (Cardiff University), and Jufeng Yang (Nankai University)

Zero-Shot Emotion Recognition via Affective Structural Embedding 1151
Chi Zhan (Nankai University), Dongyu She (Nankai University), Sicheng Zhao (University of California Berkeley), Ming-Ming Cheng (Nankai University), and Jufeng Yang (Nankai University)

FW-GAN: Flow-Navigated Warping GAN for Video Virtual Try-On 1161
Haoye Dong (Sun Yat-sen University), Xiaodan Liang (Sun Yat-sen University), Xiaohui Shen (ByteDance AI Lab), Bowen Wu (Sun Yat-sen University), Bing-Cheng Chen (Sun Yat-sen University), and Jian Yin (Sun Yat-Sen University)

Interactive Sketch & Fill: Multiclass Sketch-to-Image Translation 1171
Arnab Ghosh (University of Oxford), Richard Zhang (Adobe), Puneet Dokania (University of Oxford), Oliver Wang (Adobe Systems Inc), Alexei Efros (UC Berkeley), Philip Torr (University of Oxford), and Eli Shechtman (Adobe Research. US)

Attention-Based Autism Spectrum Disorder Screening With Privileged Modality	1181
<i>Shi Chen (University of Minnesota) and Qi Zhao (University of Minnesota)</i>	
Image Aesthetic Assessment Based on Pairwise Comparison - A Unified Approach to Score Regression, Binary Classification, and Personalization	1191
<i>Jun-Tae Lee (Korea University) and Chang-Su Kim (Korea university)</i>	
Delving Into Robust Object Detection From Unmanned Aerial Vehicles: A Deep Nuisance Disentanglement Approach	1201
<i>Zhenyu Wu (Texas A&M University), Karthik Suresh (Texas A&M University), Priya Narayanan (U.S. Army Research Laboratory), Hongyu Xu (University of Maryland), Heesung Kwon (U.S. Army Research Laboratory), and Zhangyang Wang (TAMU)</i>	
Bit-Flip Attack: Crushing Neural Network With Progressive Bit Search	1211
<i>Adnan Siraj Rakin (University of Central Florida), Zhezhi He (University of Central Florida), and Deliang Fan (Arizona State University)</i>	
Pushing the Frontiers of Unconstrained Crowd Counting: New Dataset and Benchmark Method	1221
<i>Vishwanath Sindagi (Johns Hopkins University), Rajeep Yasarla (Johns Hopkins University. Whiting School of Engineering), and Vishal Patel (Johns Hopkins University)</i>	
Employing Deep Part-Object Relationships for Salient Object Detection	1232
<i>Yi Liu (Xidian University), Qiang Zhang (Xidian University), Dingwen Zhang (Xidian University), and Jungong Han (Lancaster University)</i>	
Self-Supervised Deep Depth Denoising	1242
<i>Vladimiro Sterzentsenko (CERTH), Leonidas Saroglou (CERTH), Anargyros Chatzitofis (CERTH), Spiros Thermos (CERTH), Nikolaos Zioulis (CERTH / CENTRE FOR RESEARCH AND TECHNOLOGY HELLAS), Alexandros Doumanoglou (CERTH), Dimitrios Zarpalas (CERTH / CENTRE FOR RESEARCH AND TECHNOLOGY HELLAS), and Petros Daras (ITI-CERTH. Greece)</i>	
Cost-Aware Fine-Grained Recognition for IoTs Based on Sequential Fixations	1252
<i>Hanxiao Wang (Boston University), Venkatesh Saligrama (Boston University), Stan Sclaroff (Boston University), and Vitaly Ablavsky (Boston University)</i>	
Layout-Induced Video Representation for Recognizing Agent-in-Place Actions	1262
<i>Ruichi Yu (Waymo LLC.), Hongcheng Wang (Comcast), Ang Li (DeepMind. Mountain View), Jingxiao Zheng (University of Maryland. College Park), Vlad Morariu (Adobe Research), and Larry Davis (University of Maryland)</i>	
Anomaly Detection in Video Sequence With Appearance-Motion Correspondence	1273
<i>Trong Nguyen Nguyen (University of Montreal) and Jean Meunier (Canada)</i>	

Oral 1.2A

Architectures, Multi-Task Learning, Domain Adaptation

Exploring Randomly Wired Neural Networks for Image Recognition	1284
<i>Saining Xie (Facebook AI Research), Alexander Kirillov (Facebook AI Reserach), Ross Girshick (FAIR), and Kaiming He (Facebook AI Research)</i>	
Progressive Differentiable Architecture Search: Bridging the Depth Gap Between Search and Evaluation.....	1294
<i>Xin Chen (Tongji University), Lingxi Xie (Huawei Noah's Ark Lab), Jun Wu (Tongji Univeristy), and Qi Tian (Huawei Noah's Ark Lab)</i>	
Multinomial Distribution Learning for Effective Neural Architecture Search	1304
<i>Xiawu Zheng (Xiamen University), Rongrong Ji (Xiamen University. China), Lang Tang (Xiamen University. China), Baochang Zhang (Beihang University), Jianzhuang Liu (Huawei Noah's Ark Lab), and Qi Tian (Huawei Noah's Ark Lab)</i>	
Searching for MobileNetV3	1314
<i>Andrew Howard (Google), Mark Sandler (Google), Bo Chen (Google), Weijun Wang (Google), Liang-Chieh Chen (Google Inc.), Mingxing Tan (Google Brain), Grace Chu (Google), Vijay Vasudevan (Google Brain), Yukun Zhu (Google Inc.), Ruoming Pang (Google Brain), Hartwig Adam (Google), and Quoc Le (Google Brain)</i>	
Data-Free Quantization Through Weight Equalization and Bias Correction	1325
<i>Markus Nagel (Qualcomm), Mart Van Baalen (Qualcomm), Tijmen Blankevoort (Qualcomm), and Max Welling (University of Amsterdam)</i>	
A Camera That CNNs: Towards Embedded Neural Networks on Pixel Processor Arrays	1335
<i>Laurie Bose (University of Bristol), Piotr Dudek (School of Electrical and Electronic Engineering. The University of Manchester. UK), Jianing Chen (The University of Manchester), Stephen Carey (The University of Manchester), and Walterio Mayol-Cuevas (Bristol University)</i>	
Knowledge Distillation via Route Constrained Optimization	1345
<i>Xiao Jin (Sensetime Group Limited), Baoyun Peng (National University of Defense Technology), Yichao Wu (Sensetime Group Limited), Yu Liu (The Chinese University of Hong Kong), Jiaheng Liu (Beihang University), Ding Liang (Sensetime Group Limited), Junjie Yan (Sensetime Group Limited), and Xiaolin Hu (Tsinghua University)</i>	
Distillation-Based Training for Multi-Exit Architectures	1355
<i>Mary Phuong (IST Austria) and Christoph Lampert (IST Austria)</i>	
Similarity-Preserving Knowledge Distillation	1365
<i>Fred Tung (Simon Fraser University) and Greg Mori (Simon Fraser University)</i>	
Many Task Learning With Task Routing	1375
<i>Gjorgji Strezoski (University of Amsterdam), Nanne Noord (University of Amsterdam), and Marcel Worring (University of Amsterdam)</i>	
Stochastic Filter Groups for Multi-Task CNNs: Learning Specialist and Generalist Convolution Kernels.....	1385
<i>Felix Bragman (University College London), Ryutaro Tanno (University College London), Sebastien Ourselin (King's College London), Daniel Alexander (University College London), and Jorge Cardoso (Kings College London)</i>	

Transferability and Hardness of Supervised Classification Tasks	1395
<i>Anh Tran (Amazon Web Services), Cuong Nguyen (Amazon), and Tal Hassner (Open University of Israel)</i>	
Moment Matching for Multi-Source Domain Adaptation	1406
<i>Xingchao Peng (Boston University), Qinxun Bai (Boston University), Xide Xia (Boston University), Zijun Huang (Columbia University), Kate Saenko (Boston University), and Bo Wang (Vector Institute)</i>	
Unsupervised Domain Adaptation via Regularized Conditional Alignment	1416
<i>Safa Cicek (UCLA) and Stefano Soatto (UCLA)</i>	
Larger Norm More Transferable: An Adaptive Feature Norm Approach for Unsupervised Domain Adaptation	1426
<i>Ruijia Xu (Sun Yat-sen University), Guanbin Li (Sun Yat-sen University), Jihan Yang (Sun Yat-sen University), and Liang Lin (Sun Yat-sen University)</i>	
UM-Adapt: Unsupervised Multi-Task Adaptation Using Adversarial Cross-Task Distillation	1436
<i>Jogendra Nath Kundu (Indian Institute of Science), Nishank Lakkakula (Indian Institute of Science), and Venkatesh Babu Radhakrishnan (Indian Institute of Science)</i>	
Episodic Training for Domain Generalization	1446
<i>Da Li (QMUL), Jianshu Zhang (University of Science and Technology of China), Yongxin Yang (University of Edinburgh), Cong Liu (iFLYTEK Research), Yi-Zhe Song (University of Surrey), and Timothy Hospedales (Edinburgh University)</i>	
Domain Adaptation for Structured Output via Discriminative Patch Representations	1456
<i>Yi-Hsuan Tsai (NEC Labs America), Kihyuk Sohn (Google Cloud AI), Samuel Schuster (NEC Labs), and Manmohan Chandraker (NEC Labs America)</i>	
Semi-Supervised Learning by Augmented Distribution Alignment	1466
<i>Qin Wang (ETH Zurich), Wen Li (ETHZ), and Luc Van Gool (ETH Zurich)</i>	
S4L: Self-Supervised Semi-Supervised Learning	1476
<i>Lucas Beyer (Google Brain), Xiaohua Zhai (Google Brain), Avital Oliver (Google Brain), and Alexander Kolesnikov (Google Brain)</i>	

Oral 1.2B

Multi-View Geometry, 3D Scene Understanding

Privacy Preserving Image Queries for Camera Localization	1486
<i>Pablo Speciale (ETH Zurich), Johannes Schönberger (Microsoft), Sudipta Sinha (Microsoft Research), and Marc Pollefeys (ETH Zurich / Microsoft)</i>	
Calibration Wizard: A Guidance System for Camera Calibration Based on Modelling Geometric and Corner Uncertainty	1497
<i>Songyou Peng (ETH Zurich) and Peter Sturm (INRIA Grenoble Rhone-Alpes)</i>	
Gated2Depth: Real-Time Dense Lidar From Gated Images	1506
<i>Tobias Gruber (Daimler AG), Frank Julca-Aguilar (Algolux), Mario Bijelic (Daimler AG), and Felix Heide (Princeton University)</i>	

X-Section: Cross-Section Prediction for Enhanced RGB-D Fusion	1517
<i>Andrea Nicastro (Imperial College London), Ronald Clark (Imperial College London), and Stefan Leutenegger (Imperial College London)</i>	
Learning an Event Sequence Embedding for Dense Event-Based Deep Stereo	1527
<i>Stepan Tulyakov (École polytechnique fédérale de Lausanne), Francois Fleuret (Idiap Research Institute), Martin Kiefel (Amazon), Peter Gehler (Amazon), and Michael Hirsch (Amazon)</i>	
Point-Based Multi-View Stereo Network	1538
<i>Rui Chen (Tsinghua University), Songfang Han (HKUST), Jing Xu (Tsinghua University), and Hao Su (UCSD)</i>	
Discrete Laplace Operator Estimation for Dynamic 3D Reconstruction	1548
<i>Xiangyu Xu (STEVENS INSTITUTE OF TECHNOLOGY) and Enrique Dunn (Stevens Institute of Technology)</i>	
Deep Non-Rigid Structure From Motion	1558
<i>Chen Kong (CMU) and Simon Lucey (CMU)</i>	
Equivariant Multi-View Networks	1568
<i>Carlos Esteves (University of Pennsylvania), Yinshuang Xu (University of Pennsylvania), Christine Allec-Blanchette (University of Pennsylvania), and Kostas Daniilidis (University of Pennsylvania)</i>	
Interpolated Convolutional Networks for 3D Point Cloud Understanding	1578
<i>Jiageng Mao (The Chinese University of Hong Kong), Xiaogang Wang (Chinese University of Hong Kong, Hong Kong), and Hongsheng Li (Chinese University of Hong Kong)</i>	
Revisiting Point Cloud Classification: A New Benchmark Dataset and Classification Model on Real-World Data	1588
<i>Mikaela Angelina Uy (Hong Kong University of Science and Technology), Quang-Hieu Pham (Singapore University of Technology and Design), Binh-Son Hua (The University of Tokyo), Thanh Nguyen (Deakin University, Australia), and Sai-Kit Yeung (Hong Kong University of Science and Technology)</i>	
PointCloud Saliency Maps	1598
<i>Tianhang Zheng (University of Toronto), Changyou Chen (University at Buffalo), Junsong Yuan (State University of New York at Buffalo, USA), Bo Li (Zhejiang University), and Kui Ren (Zhejiang University)</i>	
ShellNet: Efficient Point Cloud Convolutional Neural Networks Using Concentric Shells Statistics	1607
<i>Zhiyuan Zhang (Singapore University of Technology and Design), Binh-Son Hua (The University of Tokyo), and Sai-Kit Yeung (Hong Kong University of Science and Technology)</i>	
Unsupervised Deep Learning for Structured Shape Matching	1617
<i>Jean-Michel Roufousse (École Polytechnique), Abhishek Sharma (Ecole Polytechnique), and Maks Ovsjanikov (Ecole polytechnique)</i>	
Linearly Converging Quasi Branch and Bound Algorithms for Global Rigid Registration	1628
<i>Nadav Dym (Duke University) and Shahar Kovalsky (Duke University)</i>	
Consensus Maximization Tree Search Revisited	1637
<i>Zhipeng Cai (The University of Adelaide), Tat-Jun Chin (University of Adelaide), and Vladlen Koltun (Intel Labs)</i>	

Quasi-Globally Optimal and Efficient Vanishing Point Estimation in Manhattan World	1646
<i>Haoang Li (The Chinese University of Hong Kong), Ji Zhao (TuSimple), Jean-Charles Bazin (KAIST), Wen Chen (The Chinese University of Hong Kong), Zhe Liu (The Chinese University of Hong Kong), and Yunhui Liu (CUHK)</i>	
An Efficient Solution to the Homography-Based Relative Pose Problem With a Common Reference Direction	1655
<i>Yaqing Ding (Nanjing University of Science and Technology), Jian Yang (Nanjing University of Science and Technology), Jean Ponce (Inria), and Hui Kong (Nanjing University of Science and Technology)</i>	
A Quaternion-Based Certifiably Optimal Solution to the Wahba Problem With Outliers	1665
<i>Heng Yang (MIT) and Luca Carlone (Massachusetts Institute of Technology)</i>	
PLMP - Point-Line Minimal Problems in Complete Multi-View Visibility	1675
<i>Timothy Duff (Georgia Institute of Technology), Kathlen Kohn (Universitet i Oslo), Anton Leykin (Georgia Institute of Technology), and Tomas Pajdla (Czech Technical University in Prague)</i>	

Poster 1.2

Deep Learning

Variational Few-Shot Learning	1685
<i>Jian Zhang (Shanghai Jiaotong University), Chenglong Zhao (Shanghai Jiao Tong University), Bingbing Ni (Shanghai Jiao Tong University), Minghao Xu (Shanghai Jiaotong University), and Xiaokang Yang (Shanghai Jiao Tong University of China)</i>	
Generative Adversarial Minority Oversampling	1695
<i>Sankha Subhra Mullick (Indian Statistical Institute), Shounak Datta (Indian Statistical Institute), and Swagatam Das (Indian Statistical Institute)</i>	
Memorizing Normality to Detect Anomaly: Memory-Augmented Deep Autoencoder for Unsupervised Anomaly Detection	1705
<i>Dong Gong (The University of Adelaide), Lingqiao Liu (University of Adelaide), Vuong Le (Deakin University), Budhaditya Saha (Deakin University), Moussa Reda Mansour (iCetana), Svetha Venkatesh (Deakin University), and Anton Van Den Hengel (University of Adelaide)</i>	
Topological Map Extraction From Overhead Images	1715
<i>Zuoyue Li (ETH Zurich), Jan Dirk Wegner (ETH Zurich), and Aurelien Lucchi (ETH Zurich)</i>	
Exploiting Temporal Consistency for Real-Time Video Depth Estimation	1725
<i>Haokui Zhang (Northwestern Polytechnical University), Ying Li (Northwestern Polytechnical University), Yuanzhouhan Cao (The University of Adelaide), Yu Liu (The University of Adelaide), Chunhua Shen (University of Adelaide), and Youliang Yan (Huawei)</i>	

The Sound of Motions	1735
<i>Hang Zhao (MIT), Chuang Gan (MIT-Watson AI Lab), Wei-Chiu Ma (MIT), and Antonio Torralba (MIT)</i>	
SC-FEGAN: Face Editing Generative Adversarial Network With User’s Sketch and Color	1745
<i>Youngjoo Jo (ETRI) and Jongyoul Park (ETRI)</i>	
Exploring Overall Contextual Information for Image Captioning in Human-Like Cognitive Style	1754
<i>Hongwei Ge (Dalian University of Technology), Zehang Yan (Dalian University of Technology), Kai Zhang (Dalian University of Technology), Mingde Zhao (Mila. McGill University), and Liang Sun (Dalian University of Technology)</i>	
Order-Aware Generative Modeling Using the 3D-Craft Dataset	1764
<i>Zhuoyuan Chen (Facebook AI Research), Demi Guo (Facebook AI Research), Tong Xiao (Facebook), Saining Xie (Facebook AI Research), Xinlei Chen (Facebook AI Research), Arthur Szlam (Facebook), Shubham Tulsiani (Facebook AI Research), Haonan Yu (Facebook AI Research), Jonathan Gray (Facebook AI Research), Kavya Srinet (Facebook AI Research), Charles R. Qi (Facebook AI Research), Haoqi Fan (Facebook AI Research), Jerry Ma (Facebook AI Research), and Larry Zitnick (Facebook AI Research)</i>	
Crowd Counting With Deep Structured Scale Integration Network	1774
<i>Lingbo Liu (Sun Yat-sen University), Zhilin Qiu (Sun Yat-sen University), Guanbin Li (Sun Yat-sen University), Shufan Liu (The University of Sydney), Wanli Ouyang (The University of Sydney), and Liang Lin (Sun Yat-sen University)</i>	
Bidirectional One-Shot Unsupervised Domain Mapping	1784
<i>Tomer Cohen (Tel Aviv University) and Lior Wolf (Tel Aviv University. Israel)</i>	
Evolving Space-Time Neural Architectures for Videos	1793
<i>Aj Piergiovanni (Indiana University), Anelia Angelova (Google), Alexander Toshev (Google), and Michael Ryoo (Google Brain; Indiana University)</i>	
Universally Slimmable Networks and Improved Training Techniques	1803
<i>Jiahui Yu (UIUC) and Thomas Huang (UIUC)</i>	
AutoDispNet: Improving Disparity Estimation With AutoML	1812
<i>Tonmoy Saikia (University of Freiburg), Yassine Marrakchi (University of Freiburg), Arber Zela (University of Freiburg), Frank Hutter (University of Freiburg), and Thomas Brox (University of Freiburg)</i>	
Deep Meta Functionals for Shape Representation	1824
<i>Gidi Littwin (Tel Aviv University) and Lior Wolf (Tel Aviv University. Israel)</i>	
Differentiable Kernel Evolution	1834
<i>Yu Liu (The Chinese University of Hong Kong), Jihao Liu (Sensetime), Xiaogang Wang (Chinese University of Hong Kong. Hong Kong), and Ailing Zeng (CUHK)</i>	
Batch Weight for Domain Adaptation With Mass Shift	1844
<i>Mikolaj Binkowski (Imperial College London), Devon Hjelm (Microsoft Research), and Aaron Courville (MILA. Université de Montréal)</i>	

SRM: A Style-Based Recalibration Module for Convolutional Neural Networks	1854
<i>Hyunjae Lee (Lunit Inc.), Hyo-Eun Kim (Lunit Inc.), and Hyeonseob Nam (Lunit Inc.)</i>	
Switchable Whitening for Deep Representation Learning	1863
<i>Xingang Pan (The Chinese University of Hong Kong), Xiaohang Zhan (The Chinese University of Hong Kong), Jianping Shi (Sensetime Group Limited), Xiaoou Tang (The Chinese University of Hong Kong), and Ping Luo (The Chinese University of Hong Kong)</i>	
Adaptative Inference Cost With Convolutional Neural Mixture Models	1872
<i>Adria Ruiz (Inria. LJK) and Jakob Verbeek (INRIA)</i>	
On Network Design Spaces for Visual Recognition	1882
<i>Ilija Radosavovic (Facebook AI Research), Justin Johnson (Facebook AI Research), Saining Xie (Facebook AI Research), Wan-Yen Lo (Facebook AI Research), and Piotr Dollar (FAIR)</i>	
Improved Techniques for Training Adaptive Deep Networks	1891
<i>Hao Li (Tsinghua University), Hong Zhang (Baidu Inc.), Xiaojuan Qi (University of Oxford), Yang Ruigang (Baidu), and Gao Huang (Tsinghua)</i>	
Resource Constrained Neural Network Architecture Search: Will a Submodularity Assumption Help?	1901
<i>Yunyang Xiong (University of Wisconsin-Madison), Ronak Mehta (University of Wisconsin-Madison), and Vikas Singh (University of Wisconsin-Madison USA)</i>	
ACNet: Strengthening the Kernel Skeletons for Powerful CNN via Asymmetric Convolution Blocks	1911
<i>Xiaohan Ding (Tsinghua University), Yuchen Guo (Tsinghua University), Guiguang Ding (Tsinghua University. China), and Jungong Han (Lancaster University)</i>	
A Comprehensive Overhaul of Feature Distillation	1921
<i>Byeongho Heo (Seoul National University), Jeessoo Kim (Seoul National University), Sangdoo Yun (Clova AI Research. NAVER Corp.), Hyojin Park (Seoul National University), Nojun Kwak (Seoul National University), and Jin Young Choi (Seoul National University)</i>	

Recognition

Transferable Semi-Supervised 3D Object Detection From RGB-D Data	1931
<i>Yew Siang Tang (National University of Singapore) and Gim Hee Lee (National University of Singapore)</i>	
DPOD: 6D Pose Object Detector and Refiner	1941
<i>Sergey Zakharov (TU Munich / Siemens Corporate Technology), Ivan Shugurov (TU Munich / Siemens Corporate Technology), and Slobodan Ilic (TUM)</i>	
STD: Sparse-to-Dense 3D Object Detector for Point Cloud	1951
<i>Zetong Yang (Tencent), Yanan Sun (Tencent), Shu Liu (Tencent), Xiaoyong Shen (Tencent), and Jiaya Jia (Chinese University of Hong Kong)</i>	

DUP-Net: Denoiser and Upsampler Network for 3D Adversarial Point Clouds Defense	1961
<i>Hang Zhou (University of Science and Technology of China), Kejiang Chen (University of Science and Technology of China), Weiming Zhang (University of Science and Technology of China), Han Fang (University of Science and Technology of China), Wenbo Zhou (University of Science and Technology of China), and Nenghai Yu (University of Science and Technology of China)</i>	
Learning Rich Features at High-Speed for Single-Shot Object Detection	1971
<i>Tiancai Wang (Tianjin University), Rao Muhammad Anwer (Inception Institute of Artificial Intelligence), Hisham Cholakkal (Inception Institute of Artificial Intelligence), Fahad Shahbaz Khan (Inception Institute of Artificial Intelligence), Yanwei Pang (Tianjin University), and Ling Shao (Inception Institute of Artificial Intelligence)</i>	
Detecting Unseen Visual Relations Using Analogies	1981
<i>Julia Peyre (INRIA), Josef Sivic (Inria and Czech Technical University), Ivan Laptev (INRIA Paris), and Cordelia Schmid (INRIA)</i>	
Disentangling Monocular 3D Object Detection	1991
<i>Andrea Simonelli (Mapillary Research), Samuel Rota Bulò (Mapillary Research), Lorenzo Porzi (Mapillary Research), Manuel Lopez-Antequera (Mapillary), and Peter Kotschieder (Mapillary Research)</i>	
STM: SpatioTemporal and Motion Encoding for Action Recognition	2000
<i>Boyuan Jiang (Zhejiang University), Mengmeng Wang (SenseTime Group Limited), Weihao Gan (SenseTime Group Limited), Wei Wu (SenseTime Group Limited), and Junjie Yan (Sensetime Group Limited)</i>	
Dynamic Context Correspondence Network for Semantic Alignment	2010
<i>Shuaiyi Huang (ShanghaiTech University), Qiuyue Wang (Northwestern Polytechnical University), Songyang Zhang (ShanghaiTech University), Shipeng Yan (ShanghaiTech University), and Xuming He (ShanghaiTech University)</i>	
Fooling Network Interpretation in Image Classification	2020
<i>Akshayvarun Subramanya (UMBC), Vipin Pillai (UMBC), and Hamed Pirsiavash (UMBC)</i>	
Unconstrained Foreground Object Search	2030
<i>Yinan Zhao (University of Texas at Austin), Brian Price (Adobe), Scott Cohen (Adobe Research), and Danna Gurari (University of Texas at Austin)</i>	
Embodied Amodal Recognition: Learning to Move to Perceive Objects	2040
<i>Jianwei Yang (Georgia Institute of Technology), Zhile Ren (Georgia Institute of Technology), Mingze Xu (Indiana University), Xinlei Chen (Facebook AI Research), David Crandall (Indiana University), Devi Parikh (Georgia Tech & Facebook AI Research), and Dhruv Batra (Georgia Tech & Facebook AI Research)</i>	
SpatialSense: An Adversarially Crowdsourced Benchmark for Spatial Relation Recognition	2051
<i>Kaiyu Yang (Princeton University), Olga Russakovsky (Princeton University), and Jia Deng (Princeton University)</i>	

TensorMask: A Foundation for Dense Object Segmentation	2061
<i>Xinlei Chen (Facebook AI Research), Ross Girshick (FAIR), Kaiming He (Facebook AI Research), and Piotr Dollar (FAIR)</i>	
Integral Object Mining via Online Attention Accumulation	2070
<i>Peng-Tao Jiang (Nankai University), Qibin Hou (Nankai University), Yang Cao (Nankai University), Ming-Ming Cheng (Nankai University), Yunchao Wei (UIUC), and Hongkai Xiong (Shanghai Jiao Tong University)</i>	

Segmentation, Grouping, & Shape

Accelerated Gravitational Point Set Alignment With Altered Physical Laws	2080
<i>Vladislav Golyanik (MPI), Christian Theobalt (MPI Informatik), and Didier Stricker (DFKI)</i>	
Domain Adaptation for Semantic Segmentation With Maximum Squares Loss	2090
<i>Minghao Chen (Zhejiang University), Hongyang Xue (Zhejiang University), and Deng Cai (Zhejiang University)</i>	
Domain Randomization and Pyramid Consistency: Simulation-to-Real Generalization Without Accessing Target Domain Data	2100
<i>Xiangyu Yue (UC Berkeley), Yang Zhang (University of Central Florida), Sicheng Zhao (University of California Berkeley), Alberto Sangiovanni-Vincentelli (University of California, Berkeley), Kurt Keutzer (EECS, UC Berkeley), and Boqing Gong (Google / ICSI Berkeley)</i>	
Semi-Supervised Skin Detection by Network With Mutual Guidance	2111
<i>Yi He (DUT-RU International School of Information Science & Engineering, Dalian University of Technology), Jiayuan Shi (Face++ Megvii), Chuan Wang (Face++ Megvii), Haibin Huang (Megvii Technology), Jiaming Liu (Megvii), Guanbin Li (Sun Yat-sen University), Risheng Liu (Dalian University of Technology), and Jue Wang (Megvii Technology)</i>	
ACE: Adapting to Changing Environments for Semantic Segmentation	2121
<i>Zuxuan Wu (UMD), Xin Wang (UC Berkeley), Joseph Gonzalez (UC Berkeley), Tom Goldstein (University of Maryland, College Park), and Larry Davis (University of Maryland)</i>	
Efficient Segmentation: Learning Downsampling Near Semantic Boundaries	2131
<i>Dmitrii Marin (University of Waterloo), Zijian He (Facebook), Peter Vajda (Facebook), Priyam Chatterjee (Facebook), Sam Tsai (Facebook), Fei Yang (Facebook), and Yuri Boykov (University of Waterloo)</i>	
Recurrent U-Net for Resource-Constrained Segmentation	2142
<i>Wei Wang (EPFL), Kaicheng Yu (EPFL), Joachim Hugonot (EPFL), Pascal Fua (EPFL, Switzerland), and Mathieu Salzmann (EPFL)</i>	
Detecting the Unexpected via Image Resynthesis	2152
<i>Krzysztof Lis (EPFL), Krishna Kanth Nakka (EPFL), Pascal Fua (EPFL, Switzerland), and Mathieu Salzmann (EPFL)</i>	

3D From Single View & RGBD

Self-Supervised Monocular Depth Hints	2162
<i>Jamie Watson (Niantic), Michael Firman (Niantic), Gabriel Brostow (University College London), and Daniyar Turmukhambetov (Niantic)</i>	
3D Scene Reconstruction With Multi-Layer Depth and Epipolar Transformers	2172
<i>Daeyun Shin (University of California, Irvine), Zhile Ren (Georgia Institute of Technology), Erik Sudderth (University of California, Irvine), and Charless Fowlkes (UC Irvine)</i>	
How Do Neural Networks See Depth in Single Images?	2183
<i>Tom Van Dijk (Technische Universiteit Delft) and Guido De Croon (TU Delft)</i>	
On Boosting Single-Frame 3D Human Pose Estimation via Monocular Videos	2192
<i>Zhi Li (Xi'an Jiaotong University), Xuan Wang (Xi'an Jiaotong University), Fei Wang (Xi'an Jiaotong University), and Peilin Jiang (Xi'an Jiaotong University)</i>	
Canonical Surface Mapping via Geometric Cycle Consistency	2202
<i>Nilesh Kulkarni (Carnegie Mellon University), Shubham Tulsiani (Facebook AI Research), and Abhinav Gupta (CMU/FAIR)</i>	
3D-RelNet: Joint Object and Relational Network for 3D Prediction	2212
<i>Nilesh Kulkarni (Carnegie Mellon University), Ishan Misra (Facebook AI Research), Shubham Tulsiani (Facebook AI Research), and Abhinav Gupta (CMU/FAIR)</i>	
GP2C: Geometric Projection Parameter Consensus for Joint 3D Pose and Focal Length Estimation in the Wild	2222
<i>Alexander Grabner (Graz University of Technology), Peter Roth (Graz University of Technology, Austria), and Vincent Lepetit (TU Graz)</i>	

Face & Body

Moulding Humans: Non-Parametric 3D Human Shape Estimation From Single Images	2232
<i>Valentin Gabeur (Inria), Jean-Sebastien Franco (INRIA), Xavier Martin (Inria), Cordelia Schmid (INRIA), and Gregory Rogez (NAVER LABS Europe)</i>	
3DPeople: Modeling the Geometry of Dressed Humans	2242
<i>Albert Pumarola (IRI), Jordi Sanchez (IRI), Gary P. T. Choi (Harvard University), Alberto Sanfeliu (Industrial Robotics Institute), and Francesc Moreno (IRI)</i>	
Learning to Reconstruct 3D Human Pose and Shape via Model-Fitting in the Loop	2252
<i>Nikos Kolotouros (University of Pennsylvania), Georgios Pavlakos (University of Pennsylvania), Michael Black (Max Planck Institute for Intelligent Systems), and Kostas Daniilidis (University of Pennsylvania)</i>	
Optimizing Network Structure for 3D Human Pose Estimation	2262
<i>Hai Ci (Peking University), Chunyu Wang (Microsoft Research asia), Xiaoxuan Ma (Peking University), and Yizhou Wang (PKU)</i>	

Exploiting Spatial-Temporal Relationships for 3D Pose Estimation via Graph Convolutional Networks	2272
<i>Yujun Cai (Nanyang Technological University), Lihao Ge (Nanyang Technological University), Jun Liu (Nanyang Technological University), Jianfei Cai (Nanyang Technological University), Tat-Jen Cham (Nanyang Technological University), Junsong Yuan (State University of New York at Buffalo, USA), and Nadia Magnenat Thalmann (Nanyang Technological University)</i>	
Resolving 3D Human Pose Ambiguities With 3D Scene Constraints	2282
<i>Mohamed Hassan (Max Planck Institute for Intelligent Systems), Vasileios Choutas (Max Planck Institute for Intelligent Systems), Dimitrios Tzionas (Max Planck Institute for Intelligent Systems), and Michael Black (Max Planck Institute for Intelligent Systems)</i>	
Tex2Shape: Detailed Full Human Body Geometry From a Single Image	2293
<i>Thiemo Alldieck (TU Braunschweig), Gerard Pons-Moll (MPII, Germany), Christian Theobalt (MPI Informatik), and Marcus Magnor (TU Braunschweig)</i>	
PIFu: Pixel-Aligned Implicit Function for High-Resolution Clothed Human Digitization	2304
<i>Shunsuke Saito (University of Southern California), Zeng Huang (University of Southern California), Ryota Natsume (Waseda University), Shigeo Morishima (Waseda Research Institute for Science and Engineering), Hao Li (Pinscreen/University of Southern California/USC ICT), and Angjoo Kanazawa (University of California Berkeley)</i>	
DF2Net: A Dense-Fine-Finer Network for Detailed 3D Face Reconstruction	2315
<i>Xiaoxing Zeng (Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences), Xiaojiang Peng (SIAT), and Yu Qiao (Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences)</i>	
Monocular 3D Human Pose Estimation by Generation and Ordinal Ranking	2325
<i>Saurabh Sharma (Max Planck Institute for Informatics), Pavan Teja Varigonda (IIT Bombay), Prashast Bindal (IIT Bombay), Abhishek Sharma (Facebook Inc), and Arjun Jain (Indian Institute Of Technology Bombay)</i>	
Aligning Latent Spaces for 3D Hand Pose Estimation	2335
<i>Linlin Yang (University of Bonn), Shile Li (Technical University of Munich), Dongheui Lee (Technical University of Munich), and Angela Yao (National University of Singapore)</i>	
HEMlets Pose: Learning Part-Centric Heatmap Triplets for Accurate 3D Human Pose Estimation	2344
<i>Kun Zhou (Cloudream Technology Co., Ltd), Xiaoguang Han (Shenzhen Research Institute of Big Data, the Chinese University of Hong Kong, Shenzhen), Nianjuan Jiang (Cloudream Technology Co., Ltd), Kui Jia (South China University of Technology), and Jiangbo Lu (Shenzhen Cloudream Technology Co., Ltd)</i>	
End-to-End Hand Mesh Recovery From a Monocular RGB Image	2354
<i>Xiong Zhang (Y-Lab, Kwai), Qiang Li (Y-Lab, Kwai), Hong Mo (State Key Laboratory of Virtual Reality Technology and Systems, School of Computer Science and Engineering, Beihang University), Wenbo Zhang (Y-Lab, Kwai), and Wen Zheng (Y-Lab, Kwai)</i>	

Motion & Tracking

Robust Multi-Modality Multi-Object Tracking	2365
<i>Wenwei Zhang (SenseTime Group Limited), Hui Zhou (SenseTime Group Limited.), Shuyang Sun (The University of Sydney), Zhe Wang (SenseTime Group Limited), Jianping Shi (SenseTime Group Limited), and Chen Change Loy (Nanyang Technological University)</i>	
The Trajectron: Probabilistic Multi-Agent Trajectory Modeling With Dynamic Spatiotemporal Graphs	2375
<i>Boris Ivanovic (Stanford University) and Marco Pavone (Stanford University)</i>	
‘Skimming-Perusal’ Tracking: A Framework for Real-Time and Robust Long-Term Tracking	2385
<i>Bin Yan (Dalian University of Technology), Haojie Zhao (Dalian University of Technology), Dong Wang (Dalian University of Technology), Huchuan Lu (Dalian University of Technology), and Xiaoyun Yang (China Science IntelliCloud Technology Co.. Ltd)</i>	
TASSED-Net: Temporally-Aggregating Spatial Encoder-Decoder Network for Video Saliency Detection	2394
<i>Kyle Min (University of Michigan) and Jason Corso (University of Michigan)</i>	
Attacking Optical Flow	2404
<i>Anurag Ranjan (MPI for Intelligent Systems), Joel Janai (Max Planck Institute for Intelligent Systems), Andreas Geiger (MPI-IS and University of Tuebingen), and Michael Black (Max Planck Institute for Intelligent Systems)</i>	

Computational Photography & Graphics

Pro-Cam SSfM: Projector-Camera System for Structure and Spectral Reflectance From Motion	2414
<i>Chunyu Li (Tokyo Institute of Technology), Yusuke Monno (Tokyo Institute of Technology), Hironori Hidaka (Tokyo Institute of Technology), and Masatoshi Okutomi (Tokyo Institute of Technology)</i>	
Mop Moiré Patterns Using MopNet	2424
<i>Bin He (Peking University), Ce Wang (Peking University), Boxin Shi (Peking University), and Lingyu Duan (Peking University)</i>	
Kernel Modeling Super-Resolution on Real Low-Resolution Images	2433
<i>Ruofan Zhou (EPFL) and Sabine Süsstrunk (EPFL)</i>	
Learning to Jointly Generate and Separate Reflections	2444
<i>Daiqian Ma (Peking University), Renjie Wan (Nanyang Technological University), Boxin Shi (Peking University), Alex Kot (Nanyang Technological University), and Lingyu Duan (Peking University)</i>	
Deep Multi-Model Fusion for Single-Image Dehazing	2453
<i>Zijun Deng (South China University of Technology), Lei Zhu (The Chinese University of Hong Kong), Xiaowei Hu (The Chinese University of Hong Kong), Chi-Wing Fu (The Chinese University of Hong Kong), Xuemiao Xu (South China University of Technology), Qing Zhang (Sun Yat-sen University), Jing Qin (The Hong Kong Polytechnic University), and Pheng-Ann Heng (The Chinese University of Hong Kong)</i>	

Deep Learning for Seeing Through Window With Raindrops	2463
<i>Yuhui Quan (South China Univ Tech), Shijie Deng (South China Univ Tech), Yixin Chen (South China Univ Tech), and Hui Ji (National University of Singapore)</i>	
Mask-ShadowGAN: Learning to Remove Shadows From Unpaired Data	2472
<i>Xiaowei Hu (The Chinese University of Hong Kong), Yitong Jiang (The Chinese University of Hong Kong), Chi-Wing Fu (The Chinese University of Hong Kong), and Pheng-Ann Heng (The Chinese University of Hong Kong)</i>	

Low-Level Vision & Optimization

Spatio-Temporal Filter Adaptive Network for Video Deblurring	2482
<i>Shangchen Zhou (Nanyang Technological University), Jiawei Zhang (SenseTime Research), Jinshan Pan (Nanjing University of Science and Technology), Wangmeng Zuo (Harbin Institute of Technology, China), Haozhe Xie (Harbin Institute of Technology), and Jimmy Ren (SenseTime Research)</i>	
Learning Deep Priors for Image Dehazing	2492
<i>Yang Liu (Dalian University of Technology), Jinshan Pan (Nanjing University of Science and Technology), Jimmy Ren (SenseTime Research), and Zhixun Su (Dalian University of Technology)</i>	
JPEG Artifacts Reduction via Deep Convolutional Sparse Coding	2501
<i>Xueyang Fu (University of Science and Technology of China), Zheng-Jun Zha (University of Science and Technology of China), Feng Wu (University of Science and Technology of China), Xinghao Ding (Xiamen University), and John Paisley (Columbia University)</i>	
Self-Guided Network for Fast Image Denoising	2511
<i>Shuhang Gu (ETH Zurich), Yawei Li (ETH Zurich), Luc Van Gool (ETH Zurich), and Radu Timofte (ETH Zurich)</i>	
Non-Local Intrinsic Decomposition With Near-Infrared Priors	2521
<i>Ziang Cheng (Australian National University), Yinqiang Zheng (National Institute of Informatics), Shaodi You (Data61-CSIRO), and Imari Sato (National Institute of Informatics)</i>	

Scene Understanding

VideoMem: Constructing, Analyzing, Predicting Short-Term and Long-Term Video Memorability	2531
<i>Romain Cohendet (Technicolor), Claire-Helene Demarty (Technicolor), Ngoc Duong (Technicolor), and Martin Engilberge (Technicolor)</i>	
Rescan: Inductive Instance Segmentation for Indoor RGBD Scans	2541
<i>Maciej Halber (Princeton University), Yifei Shi (National University of Defense Technology), Kai Xu (National University of Defense Technology), and Thomas Funkhouser (Princeton University)</i>	

End-to-End CAD Model Retrieval and 9DoF Alignment in 3D Scans	2551
<i>Armen Avetisyan (Technical University of Munich), Angela Dai (Technical University of Munich), and Matthias Niessner (Technical University of Munich)</i>	
Making History Matter: History-Advantage Sequence Training for Visual Dialog	2561
<i>Tianhao Yang (University of Science and Technology of China), Zheng-Jun Zha (University of Science and Technology of China), and Hanwang Zhang (Nanyang Technological University)</i>	
Stochastic Attraction-Repulsion Embedding for Large Scale Image Localization	2570
<i>Liu Liu (ANU. Australian National University), Hongdong Li (Australian National University. Australia), and Yuchao Dai (Northwestern Polytechnical University)</i>	
Scene Graph Prediction With Limited Labels	2580
<i>Ranjay Krishna (Stanford University), Vincent Chen (Stanford University), Paroma Varma (Stanford University), Michael Bernstein (Stanford University), Christopher Re (Stanford University), and Li Fei-Fei (Stanford University)</i>	

Language & Reasoning

Taking a HINT: Leveraging Explanations to Make Vision and Language Models More Grounded	2591
<i>Ramprasaath Ramasamy Selvaraju (Georgia Institute of Technology), Stefan Lee (Georgia Institute of Technology), Yilin Shen (Samsung Research America), Hongxia Jin (Samsung Research America), Shalini Ghosh (Samsung Research America), Larry Heck (Samsung Research America), Dhruv Batra (Georgia Tech & Facebook AI Research), and Devi Parikh (Georgia Tech & Facebook AI Research)</i>	
Align2Ground: Weakly Supervised Phrase Grounding Guided by Image-Caption Alignment	2601
<i>Samyak Datta (Georgia Tech), Karan Sikka (SRI International), Anirban Roy (SRI International), Karuna Ahuja (SRI International), Devi Parikh (Georgia Tech & Facebook AI Research), and Ajay Divakaran (SRI. USA)</i>	
Adaptive Reconstruction Network for Weakly Supervised Referring Expression Grounding	2611
<i>Xuejing Liu (CAS), Liang Li (Chinese Academy of Sciences), Shuhui Wang (VIPL), Zheng-Jun Zha (University of Science and Technology of China), Dechao Meng (vipl), and Qingming Huang (University of Chinese Academy of Sciences)</i>	
Hierarchy Parsing for Image Captioning	2621
<i>Ting Yao (JD AI Research), Yingwei Pan (JD AI Research), Yehao Li (Sun Yat-Sen University), and Tao Mei (AI Research of JD.com)</i>	
HowTo100M: Learning a Text-Video Embedding by Watching Hundred Million Narrated Video Clips	2630
<i>Antoine Miech (Inria), Dimitri Zhukov (Inria), Jean-Baptiste Alayrac (DeepMind), Makarand Tapaswi (INRIA), Ivan Laptev (INRIA Paris), and Josef Sivic (Inria and Czech Technical University)</i>	
Controllable Video Captioning With POS Sequence Guidance Based on Gated Fusion Network	2641
<i>Bairui Wang (Shandong University), Lin Ma (Tencent AI Lab), Wei Zhang (Shandong University), Wenhao Jiang (Tencent AI Lab), Jingwen Wang (Tencent AI Lab), and Wei Liu (Tencent)</i>	

3D From Multiview & Sensors

Multi-View Stereo by Temporal Nonparametric Fusion	2651
<i>Yuxin Hou (Aalto University), Juho Kannala (Aalto University, Finland), and Arno Solin (Aalto University)</i>	
Floor-SP: Inverse CAD for Floorplans by Sequential Room-Wise Shortest Path	2661
<i>Jiacheng Chen (Simon Fraser University), Chen Liu (Washington University in St. Louis), Jiaye Wu (Washington University in St. Louis), and Yasutaka Furukawa (Simon Fraser University)</i>	
Polarimetric Relative Pose Estimation	2671
<i>Zhaopeng Cui (ETH Zurich), Viktor Larsson (ETH Zurich), and Marc Pollefeys (ETH Zurich / Microsoft)</i>	
Closed-Form Optimal Two-View Triangulation Based on Angular Errors	2681
<i>Seong Hun Lee (University of Zaragoza) and Javier Civera (Universidad de Zaragoza)</i>	
Pix2Vox: Context-Aware 3D Reconstruction From Single and Multi-View Images	2690
<i>Haozhe Xie (Harbin Institute of Technology), Hongxun Yao (Harbin Institute of Technology), Xiaoshuai Sun (Harbin Institute of Technology), Shangchen Zhou (Nanyang Technological University), and Shengping Zhang (Harbin Institute of Technology)</i>	

Image & Video Synthesis

Unsupervised Robust Disentangling of Latent Characteristics for Image Synthesis	2699
<i>Patrick Esser (Heidelberg University), Johannes Haux (Heidelberg University), and Bjorn Ommer (Heidelberg University)</i>	
SROBB: Targeted Perceptual Loss for Single Image Super-Resolution	2710
<i>Mohammad Saeed Rad (EPFL), Behzad Bozorgtabar (EPFL), Urs-Viktor Marti (Swisscom), Max Basler (Swisscom), Hazim Kemal Ekenel (EPFL), and Jean-Philippe Thiran (École Polytechnique Fédérale de Lausanne)</i>	
An Internal Learning Approach to Video Inpainting	2720
<i>Haotian Zhang (Stanford University), Long Mai (Adobe Research), Hailin Jin (Adobe Research), Zhaowen Wang (Adobe Research), Ning Xu (Adobe Research), and John Collomosse (University of Surrey)</i>	
Deep CG2Real: Synthetic-to-Real Translation via Image Disentanglement	2730
<i>Sai Bi (University of California, San Diego), Kalyan Sunkavalli (Adobe Research), Federico Perazzi (Adobe Research), Eli Shechtman (Adobe Research, US), Vladimir Kim (Adobe), and Ravi Ramamoorthi (University of California San Diego)</i>	
Adversarial Defense via Learning to Generate Diverse Attacks	2740
<i>Yunseok Jang (University of Michigan), Tianchen Zhao (University of Michigan), Seunghoon Hong (University of Michigan), and Honglak Lee (University of Michigan Ann Arbor, USA)</i>	

Image Generation From Small Datasets via Batch Statistics Adaptation	2750
<i>Atsuhiko Noguchi (The University of Tokyo) and Tatsuya Harada (The University of Tokyo / RIKEN)</i>	
Lifelong GAN: Continual Learning for Conditional Image Generation	2759
<i>Mengyao Zhai (Simon Fraser University), Lei Chen (Simon Fraser University), Fred Tung (Simon Fraser University), Jiawei He (Simon Fraser University), Megha Nawhal (Simon Fraser University), and Greg Mori (Simon Fraser University)</i>	

Applications. Medical, & Robotics

Bayesian Relational Memory for Semantic Visual Navigation	2769
<i>Yi Wu (UC Berkeley), Yuxin Wu (Facebook AI Research), Aviv Tamar (UC Berkeley), Stuart Russell (UC Berkeley), Georgia Gkioxari (Facebook), and Yuandong Tian (Facebook)</i>	
Mono-SF: Multi-View Geometry Meets Single-View Depth for Monocular Scene Flow Estimation of Dynamic Traffic Scenes	2780
<i>Fabian Brickwedde (Robert Bosch GmbH), Steffen Abraham (Robert Bosch GmbH), and Rudolf Mester (NTNU Trondheim)</i>	
Prior Guided Dropout for Robust Visual Localization in Dynamic Environments	2791
<i>Zhaoyang Huang (Zhejiang University), Yan Xu (Sensetime Group Limited), Jianping Shi (Sensetime Group Limited), Xiaowei Zhou (Zhejiang Univ.. China), Hujun Bao (Zhejiang University), and Guofeng Zhang (Zhejiang University)</i>	
Drive&Act: A Multi-Modal Dataset for Fine-Grained Driver Behavior Recognition in Autonomous Vehicles.	2801
<i>Manuel Martin (Fraunhofer IOSB), Alina Roitberg (KIT), Monica Haurilet (KIT), Matthias Horne (Herr), Simon Reiß (Karlsruhe Institute of Technology), Michael Voit (Fraunhofer IOSB), and Rainer Stiefelwagen (Karlsruhe Institute of Technology)</i>	
Depth Completion From Sparse LiDAR Data With Depth-Normal Constraints	2811
<i>Yan Xu (Sensetime Group Limited), Xinge Zhu (The Chinese University of Hong Kong), Jianping Shi (Sensetime Group Limited), Guofeng Zhang (Zhejiang University), Hujun Bao (Zhejiang University), and Hongsheng Li (Chinese University of Hong Kong)</i>	
PRECOG: PREDiction Conditioned on Goals in Visual Multi-Agent Settings	2821
<i>Nicholas Rhinehart (CMU), Rowan Mcallister (UC Berkeley), Kris Kitani (CMU), and Sergey Levine (UC Berkeley)</i>	
LPD-Net: 3D Point Cloud Learning for Large-Scale Place Recognition and Environment Analysis	2831
<i>Zhe Liu (The Chinese University of Hong Kong), Shunbo Zhou (The Chinese University of Hong Kong), Chuanzhe Suo (The Chinese University of Hong Kong), Peng Yin (CMU), Wen Chen (The Chinese University of Hong Kong), Hesheng Wang (Shanghai Jiao Tong University), Haoang Li (The Chinese University of Hong Kong), and Yunhui Liu (CUHK)</i>	
Local Supports Global: Deep Camera Relocalization With Sequence Enhancement	2841
<i>Fei Xue (Peking University), Xin Wang (Peking University), Zike Yan (Peking University), Qiuyuan Wang (Peking University), Junqiu Wang (Beijing Changcheng Aviation Measurement and Control Institute), and Hongbin Zha (Peking University. China)</i>	

Sequential Adversarial Learning for Self-Supervised Deep Visual Odometry	2851
<i>Shunkai Li (Peking University), Fei Xue (Peking University), Xin Wang (Peking University), Zike Yan (Peking University), and Hongbin Zha (Peking University. China)</i>	
TextPlace: Visual Place Recognition and Topological Localization Through Reading Scene Texts	2861
<i>Ziyang Hong (Heriot-Watt University), Yvan Petillot (Heriot-Watt University), David Lane (Heriot-Watt University), Yishu Miao (University of Oxford), and Sen Wang (Heriot-Watt University)</i>	
CamNet: Coarse-to-Fine Retrieval for Camera Re-Localization	2871
<i>Mingyu Ding (Renmin University of China), Zhe Wang (SenseTime Group Limited), Jiankai Sun (SenseTime Group Limited), Jianping Shi (SenseTime Group Limited), and Ping Luo (The Chinese University of Hong Kong)</i>	
Situational Fusion of Visual Representation for Visual Navigation	2881
<i>William Shen (Stanford), Danfei Xu (Stanford), Yuke Zhu (Stanford University), Li Fei-Fei (Stanford University), Leonidas Guibas (Stanford University), and Silvio Savarese (Stanford University)</i>	
Learning Aberrance Repressed Correlation Filters for Real-Time UAV Tracking	2891
<i>Ziyuan Huang (National University of Singapore), Changhong Fu (Tongji University), Yiming Li (Tongji University), Fuling Lin (Tongji University), and Peng Lu (Hong Kong Polytechnic University)</i>	
6-DOF GraspNet: Variational Grasp Generation for Object Manipulation	2901
<i>Arsalan Mousavian (NVIDIA), Clemens Eppner (NVIDIA), and Dieter Fox (NVIDIA)</i>	
DAGMapper: Learning to Map by Discovering Lane Topology	2911
<i>Namdar Homayounfar (University of Toronto), Justin Liang (Uber ATG), Wei-Chiu Ma (MIT), Jack Fan (Uber), Xinyu Wu (UberATG), and Raquel Urtasun (Uber ATG)</i>	
3D-LaneNet: End-to-End 3D Multiple Lane Detection	2921
<i>Noa Garnett (GM), Rafi Cohen (GM), Tomer Pe'er (GM), Roei Lahav (GM), and Dan Levi (General Motors)</i>	

Oral 2.1A

Feature Representations, Similarity Learning

Sampling-Free Epistemic Uncertainty Estimation Using Approximated Variance Propagation	2931
<i>Janis Postels (Autonomous Intelligent Driving GmbH), Francesco Ferroni (Autonomous Intelligent Driving GmbH), Huseyin Coskun (Technical University of Munich), Nassir Navab (TU Munich. Germany), and Federico Tombari (TUM. Google)</i>	
Universal Adversarial Perturbation via Prior Driven Uncertainty Approximation	2941
<i>Hong Liu (Xiamen University), Rongrong Ji (Xiamen University. China), Jie Li (Xiamen University), Baochang Zhang (Beihang University), Yue Gao (Tsinghua University), Yongjian Wu (Tencent Technology. Shanghai Co.), and Feiyue Huang (Tencent)</i>	

Understanding Deep Networks via Extremal Perturbations and Smooth Masks	2950
<i>Ruth Fong (University of Oxford), Mandela Patrick (University of Oxford), and Andrea Vedaldi (Oxford University)</i>	
Unsupervised Pre-Training of Image Features on Non-Curated Data	2959
<i>Mathilde Caron (Facebook Artificial Intelligence Research), Piotr Bojanowski (Facebook), Julien Mairal (INRIA), and Armand Joulin (Facebook AI Research)</i>	
Learning Local Descriptors With a CDF-Based Dynamic Soft Margin	2969
<i>Linguang Zhang (Princeton University) and Szymon Rusinkiewicz (Princeton University)</i>	
Bayes-Factor-VAE: Hierarchical Bayesian Deep Auto-Encoder Models for Factor Disentanglement	2979
<i>Minyoung Kim (Samsung AI Center), Yuting Wang (Rutgers University), Pritish Sahu (Rutgers University), and Vladimir Pavlovic (Rutgers University)</i>	
Linearized Multi-Sampling for Differentiable Image Transformation	2988
<i>Wei Jiang (University of Victoria), Weiwei Sun (University of Victoria), Andrea Tagliasacchi (Google Inc.), Eduard Trulls (Google), and Kwang Moo Yi (University of Victoria)</i>	
AdaTransform: Adaptive Data Transformation	2998
<i>Zhiqiang Tang (Rutgers), Xi Peng (University of Delaware), Tingfeng Li (Rutgers University), Yizhe Zhu (Rutgers University), and Dimitris Metaxas (Rutgers)</i>	
CARAFE: Content-Aware ReAssembly of FEatures	3007
<i>Jiaqi Wang (The Chinese University of Hong Kong), Kai Chen (The Chinese University of Hong Kong), Rui Xu (CUHK), Ziwei Liu (The Chinese University of Hong Kong), Chen Change Loy (Nanyang Technological University), and Dahua Lin (The Chinese University of Hong Kong)</i>	
AFD-Net: Aggregated Feature Difference Learning for Cross-Spectral Image Patch Matching	3017
<i>Dou Quan (School of Artificial Intelligence, Xidian University, China), Xuefeng Liang (Xidian University), Shuang Wang (Xidian University), Shaowei Wei (Xidian University), Yanfeng Li (Xidian University), Ning Huyan (Xidian University), and Licheng Jiao (Xidian University)</i>	
Deep Joint-Semantics Reconstructing Hashing for Large-Scale Unsupervised Cross-Modal Retrieval	3027
<i>Shupeng Su (Peking University), Zhisheng Zhong (Peking University), and Chao Zhang (Peking University)</i>	
Unsupervised Neural Quantization for Compressed-Domain Similarity Search	3036
<i>Stanislav Morozov (Yandex) and Artem Babenko (MIPT/Yandex)</i>	
Siamese Networks: The Tale of Two Manifolds	3046
<i>Soumava Roy (AUSTRALIAN NATIONAL UNIVERSITY), Mehrtash Harandi (Monash University), Richard Nock (Data61, CSIRO), and Richard Hartley (Australian National University, Australia)</i>	
Learning Combinatorial Embedding Networks for Deep Graph Matching	3056
<i>Runzhong Wang (Shanghai Jiao Tong University), Junchi Yan (Shanghai Jiao Tong University), and Xiaokang Yang (Shanghai Jiao Tong University of China)</i>	

Fashion Retrieval via Graph Reasoning Networks on a Similarity Pyramid	3066
<i>Zhanghui Kuang (SenseTime Ltd.), Yiming Gao (Sun Yat-sen University), Guanbin Li (Sun Yat-sen University), Ping Luo (The Chinese University of Hong Kong), Yimin Chen (senseTime), Liang Lin (Sun Yat-sen University), and Wayne Zhang (SenseTime Research)</i>	

Oral 2.1B

Low Level Vision

Wavelet Domain Style Transfer for an Effective Perception-Distortion Tradeoff in Single Image Super-Resolution	3076
<i>Xin Deng (Imperial College London), Ren Yang (ETH Zurich), Mai Xu (BUAA), and Pier Luigi Dragotti (Imperial College London)</i>	
Toward Real-World Single Image Super-Resolution: A New Benchmark and a New Model	3086
<i>Jianrui Cai (The Hong Kong Polytechnic University. Hong Kong. China), Hui Zeng (The Hong Kong Polytechnic University), Hongwei Yong (The Hong Kong Polytechnic University), Zisheng Cao (Da-Jiang Innovations), and Lei Zhang (Hong Kong Polytechnic University. Hong Kong. China)</i>	
RankSRGAN: Generative Adversarial Networks With Ranker for Image Super-Resolution	3096
<i>Wenlong Zhang (SIAT), Yihao Liu (University of Chinese Academy of Sciences), Chao Dong (SIAT), and Yu Qiao (Shenzhen Institutes of Advanced Technology. Chinese Academy of Sciences)</i>	
Progressive Fusion Video Super-Resolution Network via Exploiting Non-Local Spatio-Temporal Correlations	3106
<i>Peng Yi (National Engineering Research Center for Multimedia Software. Wuhan University. China), Zhongyuan Wang (National Engineering Research Center for Multimedia Software. Wuhan University. China), Kui Jiang (Wuhan University), Junjun Jiang (Harbin Institute of Technology), and Jiayi Ma (Wuhan University)</i>	
Deep SR-ITM: Joint Learning of Super-Resolution and Inverse Tone-Mapping for 4K UHD HDR Applications.....	3116
<i>Soo Ye Kim (KAIST), Jihyong Oh (KAIST), and Munchurl Kim (KAIST-VICLab)</i>	
Dynamic PET Image Reconstruction Using Nonnegative Matrix Factorization Incorporated With Deep Image Prior	3126
<i>Tatsuya Yokota (Nagoya Institute of Technology), Kazuya Kawai (Nagoya Institute of Technology), Muneyuki Sakata (Tokyo Metropolitan Institute of Gerontology), Yuichi Kimura (Kindai University), and Hidekata Hontani (Nagoya Institute of Technology)</i>	
DSIC: Deep Stereo Image Compression	3136
<i>Jerry Liu (Uber ATG), Shenlong Wang (Uber ATG. University of Toronto), and Raquel Urtasun (Uber ATG)</i>	
Variable Rate Deep Image Compression With a Conditional Autoencoder	3146
<i>Yoojin Choi (Samsung Research USA), Mostafa El-Khamy (Samsung Research USA), and Jungwon Lee (Samsung Research USA)</i>	

Real Image Denoising With Feature Attention	3155
<i>Saeed Anwar (ANU) and Nick Barnes (CSIRO(Data61))</i>	
Noise Flow: Noise Modeling With Conditional Normalizing Flows	3165
<i>Abdelrahman Abdelhamed (York University), Marcus Brubaker (Borealis AI), and Michael Brown (York University)</i>	
Bottleneck Potentials in Markov Random Fields	3174
<i>Ahmed Abbas (MPI-INF) and Paul Swoboda (MPI fuer Informatik. Saarbruecken)</i>	
Seeing Motion in the Dark	3184
<i>Chen Chen (UIUC), Qifeng Chen (HKUST), Minh Do (UIUC), and Vladlen Koltun (Intel Labs)</i>	
SENSE: A Shared Encoder Network for Scene-Flow Estimation	3194
<i>Huaizu Jiang (UMass Amherst), Deqing Sun (Google), Varun Jampani (Nvidia Research), Zhaoyang Lv (GEORGIA TECH), Erik Learned-Miller (University of Massachusetts. Amherst), and Jan Kautz (NVIDIA)</i>	

Poster 2.1

Deep Learning

Adversarial Feedback Loop	3204
<i>Firas Shama (Technion), Roey Mechrez (Technion), Alon Shoshan (Technion), and Lihi Zelnik-Manor (Technion)</i>	
Dynamic-Net: Tuning the Objective Without Re-Training for Synthesis Tasks	3214
<i>Alon Shoshan (Technion), Roey Mechrez (Technion), and Lihi Zelnik-Manor (Technion)</i>	
AutoGAN: Neural Architecture Search for Generative Adversarial Networks	3223
<i>Xinyu Gong (Texas A&M University), Shiyu Chang (IBM Research), Yifan Jiang (Huazhong University of Science and Technology), and Zhangyang Wang (TAMU)</i>	
Co-Evolutionary Compression for Unpaired Image Translation	3234
<i>Han Shu (Huawei Noah's Ark Lab), Yunhe Wang (Huawei Noah's Ark Lab), Xu Jia (Huawei Noah's Ark Lab), Kai Han (Huawei Noah's Ark Lab), Hanting Chen (Peking University), Chunjing Xu (Huawei Noah's Ark Lab), Qi Tian (Huawei Noah's Ark Lab), and Chang Xu (University of Sydney)</i>	
Self-Supervised Representation Learning From Multi-Domain Data	3244
<i>Zeyu Feng (University of Sydney), Chang Xu (University of Sydney), and Dacheng Tao (University of Sydney)</i>	
Controlling Neural Networks via Energy Dissipation	3255
<i>Michael Moeller (University of Siegen), Thomas Moellenhoff (Technical University of Munich), and Daniel Cremers (TU Munich)</i>	
Indices Matter: Learning to Index for Deep Image Matting	3265
<i>Hao Lu (The University of Adelaide), Yutong Dai (The University of Adelaide), Chunhua Shen (University of Adelaide), and Songcen Xu (Noah's Ark Lab. Huawei Technologies Co.. Ltd.)</i>	

LAP-Net: Level-Aware Progressive Network for Image Dehazing	3275
<i>Yunan Li (Xidian University), Qiguang Miao (Xidian University), Wanli Ouyang (The University of Sydney), Zhenxin Ma (Xidian University), Huijuan Fang (Xidian University), Chao Dong (SIAT), and Yining Quan (Xidian University)</i>	
Attention Augmented Convolutional Networks	3285
<i>Irwan Bello (Google), Barret Zoph (Google), Quoc Le (Google Brain), Ashish Vaswani (Google Brain), and Jonathon Shlens (Google)</i>	
MetaPruning: Meta Learning for Automatic Neural Network Channel Pruning	3295
<i>Zechun Liu (HKUST), Haoyuan Mu (Tsinghua University), Xiangyu Zhang (Megvii Technology), Zichao Guo (Megvii Inc), Xin Yang (Huazhong University of Science and Technology), Kwang-Ting Cheng (Hong Kong University of Science and Technology), and Jian Sun (Megvii Technology)</i>	
Accelerate CNN via Recursive Bayesian Pruning	3305
<i>Yuefu Zhou (Cooperative Medianet Innovation Center. Shanghai Jiao Tong University), Ya Zhang (Cooperative Medianet Innovation Center. Shanghai Jiao Tong University), Yan-Feng Wang (Cooperative medianet innovation center of Shanghai Jiao Tong University), and Qi Tian (Huawei Noah's Ark Lab)</i>	
HBONet: Harmonious Bottleneck on Two Orthogonal Dimensions	3315
<i>Duo Li (Intel Labs China), Aojun Zhou (Intel Labs China), and Anbang Yao (Intel Labs China)</i>	
O2U-Net: A Simple Noisy Label Detection Approach for Deep Neural Networks	3325
<i>Jinchi Huang (Alibaba Group), Lie Qu (Alibaba Group), Rongfei Jia (Alibaba Group), and Binqiang Zhao (Alibaba)</i>	
Continual Learning by Asymmetric Loss Approximation With Single-Side Overestimation	3334
<i>Dongmin Park (Seoul National University), Seokil Hong (Seoul National University), Bohyung Han (Seoul National University), and Kyoung Mu Lee (Seoul National University)</i>	
Label-PENet: Sequential Label Propagation and Enhancement Networks for Weakly Supervised Instance Segmentation	3344
<i>Weifeng Ge (The University of Hong Kong), Weilin Huang (Malong Technologies), Sheng Guo (Malong Technologies), and Matthew Scott (Malong Technologies)</i>	
LIP: Local Importance-Based Pooling	3354
<i>Ziteng Gao (Nanjing University), Limin Wang (Nanjing University), and Gangshan Wu (Nanjing University)</i>	
Global Feature Guided Local Pooling	3364
<i>Takumi Kobayashi (National Institute of Advanced Industrial Science and Technology)</i>	
Conditional Coupled Generative Adversarial Networks for Zero-Shot Domain Adaptation	3374
<i>Jinghua Wang (Shenzhen University) and Jianmin Jiang (Shenzhen University)</i>	

Adversarial Defense by Restricting the Hidden Space of Deep Neural Networks	3384
<i>Aamir Mustafa (Inception Institute of Artificial Intelligence), Salman Khan (Australian National University. ANU), Munawar Hayat (University of Canberra), Roland Goecke (University of Canberra), Jianbing Shen (Inception Institute of Artificial Intelligence), and Ling Shao (Inception Institute of Artificial Intelligence)</i>	
Hyperpixel Flow: Semantic Correspondence With Multi-Layer Neural Features	3394
<i>Juhong Min (POSTECH), Jongmin Lee (POSTECH), Jean Ponce (Inria), and Minsu Cho (POSTECH)</i>	
Information Entropy Based Feature Pooling for Convolutional Neural Networks	3404
<i>Weitao Wan (Tsinghua University), Jiansheng Chen (Tsinghua University), Tianpeng Li (Tsinghua University), Yiqing Huang (Tsinghua University), Jingqi Tian (Tsinghua University), Cheng Yu (Tsinghua University), and Youze Xue (Tsinghua University)</i>	
Patchwork: A Patch-Wise Attention Network for Efficient Object Detection and Segmentation in Video Streams	3414
<i>Yuning Chai (Alphabet)</i>	
AttentionRNN: A Structured Spatial Attention Mechanism	3424
<i>Siddhesh Khandelwal (University of British Columbia) and Leonid Sigal (University of British Columbia)</i>	
Drop an Octave: Reducing Spatial Redundancy in Convolutional Neural Networks With Octave Convolution.	3434
<i>Yunpeng Chen (National University of Singapore), Haoqi Fan (Facebook AI Research), Bing Xu (Facebook AI), Zhicheng Yan (Facebook AI), Yannis Kalantidis (Facebook Research), Marcus Rohrbach (Facebook AI Research), Yan Shuicheng (National University of Singapore), and Jiashi Feng (NUS)</i>	
Domain Intersection and Domain Difference	3444
<i>Sagie Benaim (Tel Aviv University), Michael Khaitov (Tel Aviv University), Tomer Galanti (Tel Aviv University), and Lior Wolf (Tel Aviv University. Israel)</i>	
Learned Video Compression	3453
<i>Oren Rippel (WaveOne. Inc.), Sanjay Nair (WaveOne. Inc.), Carissa Lew (WaveOne. Inc.), Steve Branson (WaveOne. Inc.), Alexander Anderson (WaveOne. Inc.), and Lubomir Bourdev (WaveOne. Inc.)</i>	
Local Relation Networks for Image Recognition	3463
<i>Han Hu (Microsoft Research Asia), Zheng Zhang (MSRA), Zhenda Xie (Tsinghua University), and Stephen Lin (Microsoft Research)</i>	
DiscoNet: Shapes Learning on Disconnected Manifolds for 3D Editing	3473
<i>Eloi Mehr (Dassault Systèmes), Ariane Jourdan (Dassault Systèmes), Nicolas Thome (CNAM. Paris), Matthieu Cord (Sorbonne University), and Vincent Guittény (Dassault Systèmes)</i>	
Deep Residual Learning in the JPEG Transform Domain	3483
<i>Max Ehrlich (University of Maryland) and Larry Davis (University of Maryland)</i>	
Approximated Bilinear Modules for Temporal Modeling	3493
<i>Xinqi Zhu (University of Sydney), Chang Xu (University of Sydney), Langwen Hui (Shanghai Jiaotong University), Cewu Lu (Shanghai Jiao Tong University), and Dacheng Tao (University of Sydney)</i>	

Customizing Student Networks From Heterogeneous Teachers via Adaptive Knowledge Amalgamation	3503
<i>Chengchao Shen (Zhejiang University), Mengqi Xue (Zhejiang University), Xinchao Wang (Stevens Institute of Technology), Jie Song (College of Computer Science and Technology, Zhejiang University), Li Sun (Zhejiang University), and Mingli Song (Zhejiang University)</i>	
Data-Free Learning of Student Networks	3513
<i>Hanting Chen (Peking University), Yunhe Wang (Huawei Noah's Ark Lab), Chang Xu (University of Sydney), Zhaohui Yang (Peking University), Chuanjian Liu (Huawei Noah's Ark Lab), Boxin Shi (Peking University), Chunjing Xu (Huawei Noah's Ark Lab), Chao Xu (Peking University), and Qi Tian (Huawei Noah's Ark Lab)</i>	
Deep Closest Point: Learning Representations for Point Cloud Registration	3522
<i>Yue Wang (MIT) and Justin Solomon (MIT)</i>	
Orientation-Aware Semantic Segmentation on Icosahedron Spheres	3532
<i>Chao Zhang (Toshiba Research Europe Ltd), Stephan Liwicki (Toshiba Research Europe Ltd), William Smith (University of York), and Roberto Cipolla (University of Cambridge)</i>	
Differentiable Learning-to-Group Channels via Groupable Convolutional Neural Networks	3541
<i>Zhaoyang Zhang (The Chinese University of Hong Kong), Jingyu Li (SenseTime Research), Wenqi Shao (The Chinese University of Hong Kong), Zhanglin Peng (SenseTime), Ruimao Zhang (SenseTime Research), Xiaogang Wang (Chinese University of Hong Kong, Hong Kong), and Ping Luo (The Chinese University of Hong Kong)</i>	
HarDNet: A Low Memory Traffic Network	3551
<i>Ping Chao (university of michigan), Chao-Yang Kao (National Tsing Hua University), Yushan Ruan (National Tsing Hua University), Chien-Hsiang Huang (NTHU), and Youn-Long Lin (National Tsing Hua University)</i>	
Dynamic Multi-Scale Filters for Semantic Segmentation	3561
<i>Junjun He (SJTU), Zhongying Deng (SIAT), and Yu Qiao (Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences)</i>	
Online Model Distillation for Efficient Video Inference	3572
<i>Ravi Teja Mullanpudi (CMU), Steven Chen (Stanford University), Keyi Zhang (Stanford University), Deva Ramanan (Carnegie Mellon University), and Kayvon Fatahalian (Stanford)</i>	

Recognition

Rethinking Zero-Shot Learning: A Conditional Visual Classification Perspective	3582
<i>Kai Li (Northeastern University), Martin Renqiang Min (NEC Labs America-Princeton), and Yun Fu (Northeastern University)</i>	
Task-Driven Modular Networks for Zero-Shot Compositional Learning	3592
<i>Senthil Purushwalkam (Carnegie Mellon University), Maximillian Nickel (Facebook AI Research), Abhinav Gupta (CMU/FAIR), and Marc'aurelio Ranzato (Facebook)</i>	

Transductive Episodic-Wise Adaptive Metric for Few-Shot Learning	3602
<i>Limeng Qiao (Peking University), Yemin Shi (Peking University), Jia Li (Beihang University), Yonghong Tian (PKU), Tiejun Huang (Peking University), and Yaowei Wang (Beijing Institute of Technology)</i>	
Deep Multiple-Attribute-Perceived Network for Real-World Texture Recognition	3612
<i>Wei Zhai (University of Science and Technology of China), Yang Cao (University of Science and Technology of China), Jing Zhang (The University of Sydney), and Zheng-Jun Zha (University of Science and Technology of China)</i>	
RGB-Infrared Cross-Modality Person Re-Identification via Joint Pixel and Feature Alignment	3622
<i>Guan'an Wang (Chinese Academy of Sciences), Tianzhu Zhang (University of Science and Technology of China), Jian Cheng (Chinese Academy of Sciences. China), Si Liu (Beihang University), Yang Yang (Chinese Academy of Sciences), and Zengguang Hou (Chinese Academy of Sciences)</i>	
EvalNorm: Estimating Batch Normalization Statistics for Evaluation	3632
<i>Saurabh Singh (Google) and Abhinav Shrivastava (University of Maryland)</i>	
Beyond Human Parts: Dual Part-Aligned Representations for Person Re-Identification	3641
<i>Jianyuan Guo (Peking University), Yuhui Yuan (Microsoft Research), Lang Huang (Peking University), Chao Zhang (Peking University), Jin-Ge Yao (Microsoft), and Kai Han (Huawei Noah's Ark Lab)</i>	
Person Search by Text Attribute Query As Zero-Shot Learning	3651
<i>Qi Dong (Queen Mary University of London), Xiatian Zhu (Vision Semantics Limited), and Shaogang Gong (Queen Mary University of London)</i>	
Semantic-Aware Knowledge Preservation for Zero-Shot Sketch-Based Image Retrieval	3661
<i>Qing Liu (Johns Hopkins University), Lingxi Xie (Huawei Noah's Ark Lab), Huiyu Wang (Johns Hopkins University), and Alan Yuille (Johns Hopkins University)</i>	
Active Learning for Deep Detection Neural Networks	3671
<i>Hamed H. Aghdam (Computer Vision Center. UAB), Abel Gonzalez-Garcia (Computer Vision Center), Antonio Lopez (CVC & UAB), and Joost Weijer (Computer Vision Center)</i>	
One-Shot Neural Architecture Search via Self-Evaluated Template Network	3680
<i>Xuanyi Dong (UTS) and Yi Yang (UTS)</i>	
Batch DropBlock Network for Person Re-Identification and Beyond	3690
<i>Zuozhuo Dai (Alibaba A.I. Labs), Mingqiang Chen (Alibaba A.I. Labs), Xiaodong Gu (Alibaba A.I. Labs), Siyu Zhu (Alibaba A.I. Labs), and Ping Tan (Simon Fraser University)</i>	
Omni-Scale Feature Learning for Person Re-Identification	3701
<i>Kaiyang Zhou (University of Surrey), Yongxin Yang (University of Edinburgh), Andrea Cavallaro (Queen Mary University of London. UK), and Tao Xiang (University of Surrey)</i>	

Be Your Own Teacher: Improve the Performance of Convolutional Neural Networks via Self Distillation	3712
<i>Linfeng Zhang (Tsinghua University), Jiebo Song (Institute for Interdisciplinary Information Core Technology), Anni Gao (Institute for Interdisciplinary Information Core Technology), Jingwei Chen (Tsinghua University), Chenglong Bao (Tsinghua University), and Kaisheng Ma (Tsinghua University)</i>	
Diversity With Cooperation: Ensemble Methods for Few-Shot Classification	3722
<i>Nikita Dvornik (INRIA), Julien Mairal (INRIA), and Cordelia Schmid (INRIA)</i>	
Enhancing 2D Representation via Adjacent Views for 3D Shape Retrieval	3731
<i>Cheng Xu (Beihang University), Zhaoqun Li (Beihang University), Qiang Qiu (Duke University), Biao Leng (Beihang University), and Jingfei Jiang (Science and Technology on Parallel and Distributed Processing Laboratory. National University of Defense Technology)</i>	
Adversarial Fine-Grained Composition Learning for Unseen Attribute-Object Recognition	3740
<i>Kun Wei (Xidian University), Muli Yang (Xidian University), Hao Wang (Xidian University), Cheng Deng (Xidian University), and Xianglong Liu (Beihang University)</i>	
Auto-ReID: Searching for a Part-Aware ConvNet for Person Re-Identification	3749
<i>Ruijie Quan (University of Technology Sydney), Xuanyi Dong (UTS), Yu Wu (University of Technology Sydney), Linchao Zhu (University of Technology. Sydney), and Yi Yang (UTS)</i>	
Second-Order Non-Local Attention Networks for Person Re-Identification	3759
<i>Bryan Bryan (Ning Xia. University of Notre Dame), Yuan Gong (University of Notre Dame), Yizhe Zhang (University of Notre Dame), and Christian Poellabauer (University of Notre Dame)</i>	

Segmentation, Grouping, & Shape

Fast Computation of Content-Sensitive Superpixels and Supervoxels Using Q-Distances	3769
<i>Zipeng Ye (Tsinghua University), Ran Yi (Tsinghua University), Minjing Yu (Tsinghua University), Yong-Jin Liu (Tsinghua University), and Ying He (Nanyang Technological University)</i>	
Progressive-X: Efficient, Anytime, Multi-Model Fitting Algorithm	3779
<i>Dániel Baráth (MTA SZTAKI. CMP Prague) and Jiri Matas (CMP CTU FEE)</i>	
Structured Modeling of Joint Deep Feature and Prediction Refinement for Salient Object Detection	3788
<i>Yingyue Xu (University of Oulu), Dan Xu (University of Oxford), Xiaopeng Hong (Xi'an Jiaotong University), Wanli Ouyang (The University of Sydney), Rongrong Ji (Xiamen University. China), Min Xu (University of Technology Sydney), and Guoying Zhao (University of Oulu)</i>	
Selectivity or Invariance: Boundary-Aware Salient Object Detection	3798
<i>Jinming Su (Beihang University), Jia Li (Beihang University), Yu Zhang (Beihang University), Changqun Xia (Beihang University), and Yonghong Tian (PKU)</i>	

Online Unsupervised Learning of the 3D Kinematic Structure of Arbitrary Rigid Bodies	3808
<i>Urbano Miguel Nunes (Imperial College London) and Yiannis Demiris (Imperial College London)</i>	

3D From Single View & RGBD

Few-Shot Generalization for Single-Image 3D Reconstruction via Priors	3817
<i>Bram Wallace (Cornell University) and Bharath Hariharan (Cornell University)</i>	
Digging Into Self-Supervised Monocular Depth Estimation	3827
<i>Clement Godard (University College London), Oisín Mac Aodha (Caltech), Michael Firman (Niantic), and Gabriel Brostow (University College London)</i>	
Learning Object-Specific Distance From a Monocular Image	3838
<i>Jing Zhu (New York University) and Yi Fang (New York University)</i>	
Unsupervised 3D Reconstruction Networks	3848
<i>Geonho Cha (Seoul National University), Minsik Lee (Hanyang University), and Songhwai Oh (Seoul National University)</i>	
3D Point Cloud Generative Adversarial Network Based on Tree Structured Graph Convolutions	3858
<i>Dongwook Shu (Chung-Ang Univ.. Korea), Sung Woo Park (Chung-Ang Univ.. Korea), and Junseok Kwon (Chung-Ang Univ.. Korea)</i>	
Visualization of Convolutional Neural Networks for Monocular Depth Estimation	3868
<i>Junjie Hu (Tohoku University), Yan Zhang (RIKEN AIP), and Takayuki Okatani (Tohoku University/RIKEN AIP)</i>	

Action & Video

Co-Separating Sounds of Visual Objects	3878
<i>Ruohan Gao (University of Texas at Austin) and Kristen Grauman (Facebook AI Research & UT Austin)</i>	
BMN: Boundary-Matching Network for Temporal Action Proposal Generation	3888
<i>Tianwei Lin (Baidu), Xiao Liu (Baidu), Xin Li (Baidu), Errui Ding (Baidu Inc.), and Shilei Wen (Baidu Research)</i>	
Weakly Supervised Temporal Action Localization Through Contrast Based Evaluation Networks	3898
<i>Ziyi Liu (Xi'an Jiaotong University), Le Wang (Xi'an Jiaotong University), Qilin Zhang (HERE Technologies), Zhanning Gao (Alibaba Group), Zhenxing Niu (Alibaba Group), Nanning Zheng (Xi'an Jiaotong University), and Gang Hua (Wormpex AI Research)</i>	
Progressive Sparse Local Attention for Video Object Detection	3908
<i>Chaoxu Guo (Institute of Automation. Chinese Academy of Science), Bin Fan (Institute of Automation. Chinese Academy of Sciences. China), Jie Gu (National Laboratory of Pattern Recognition. Institute of Automation. Chinese Academy of Sciences), Qian Zhang (Horizon Robotics), Shiming Xiang (Chinese Academy of Sciences. China), Véronique Prinnet (The Hebrew University of Jerusalem. Israel), and Chunhong Pan (Institute of Automation. Chinese Academy of Sciences)</i>	

Reasoning About Human-Object Interactions Through Dual Attention Networks	3918
<i>Tete Xiao (Peking University), Quanfu Fan (IBM Research), Danny Gutfreund (IBM), Mathew Monfort (MIT), Aude Oliva (MIT), and Bolei Zhou (CUHK)</i>	
DMM-Net: Differentiable Mask-Matching Network for Video Object Segmentation	3928
<i>Xiaohui Zeng (University of Toronto), Renjie Liao (University of Toronto), Li Gu (University of Toronto), Yuwen Xiong (Uber ATG), Sanja Fidler (University of Toronto. NVIDIA), and Raquel Urtasun (Uber ATG)</i>	
Asymmetric Cross-Guided Attention Network for Actor and Action Video Segmentation From Natural Language Query	3938
<i>Hao Wang (Xidian University), Cheng Deng (Xidian University), Junchi Yan (Shanghai Jiao Tong University), and Dacheng Tao (University of Sydney)</i>	
AGSS-VOS: Attention Guided Single-Shot Video Object Segmentation	3948
<i>Huaijia Lin (The Chinese University of Hong Kong), Xiaojuan Qi (University of Oxford), and Jiaya Jia (Chinese University of Hong Kong)</i>	
Global-Local Temporal Representations for Video Person Re-Identification	3957
<i>Jianing Li (Peking University), Shiliang Zhang (Peking University), Jingdong Wang ((), Wen Gao (PKU), and Qi Tian (Huawei Noah's Ark Lab)</i>	
AdvIT: Adversarial Frames Identifier Based on Temporal Consistency in Videos	3967
<i>Chaowei Xiao (University of Michigan. Ann Arbor), Ruizhi Deng (Simon Fraser University), Bo Li (University of Illinois at Urbana-Champaign), Taesung Lee (IBM Research), Benjamin Edwards (IBM Research), Jinfeng Yi (JD AI Research), Dawn Song (UC Berkeley), Mingyan Liu (university of Michigan. Ann Arbor), and Ian Molloy (IBM Research AI)</i>	

Motion & Tracking

RANet: Ranking Attention Network for Fast Video Object Segmentation	3977
<i>Ziqin Wang (Xi'an Jiaotong University), Jun Xu (Inception Institute of Artificial Intelligence), Li Liu (the inception institute of artificial intelligence), Fan Zhu (Inception Institute of Artificial Intelligence), and Ling Shao (Inception Institute of Artificial Intelligence)</i>	
Spatial-Temporal Relation Networks for Multi-Object Tracking	3987
<i>Jiarui Xu (Hong Kong University of Science and Technology. HKUST), Yue Cao (Tsinghua University), Zheng Zhang (MSRA), and Han Hu (Microsoft Research Asia)</i>	
Bridging the Gap Between Detection and Tracking: A Unified Approach	3998
<i>Lianghua Huang (CASIA), Xin Zhao (Institute of Automation. Chinese Academy of Sciences), and Kaiqi Huang (Institute of Automation. Chinese Academy of Sciences)</i>	

Learning the Model Update for Siamese Trackers	4009
<i>Lichao Zhang (Computer Vision Center), Abel Gonzalez-Garcia (Computer Vision Center), Joost Van De Weijer (Computer Vision Center), Martin Danelljan (ETH Zurich), and Fahad Shahbaz Khan (Inception Institute of Artificial Intelligence)</i>	
Fast-deepKCF Without Boundary Effect	4019
<i>Linyu Zheng (Institute of Automation. Chinese Academy of Sciences), Ming Tang (Chinese Academy of Sciences. China), Yingying Chen (CASIA), Jinqiao Wang (Institute of Automation. Chinese Academy of Sciences), and Hanqing Lu (NLPR. Institute of Automation. CAS)</i>	

Computational Photography & Graphics

Program-Guided Image Manipulators	4029
<i>Xiuming Zhang (MIT), Jiayuan Mao (Tsinghua University), Yikai Li (Shanghai Jiao Tong University), William Freeman (MIT), Joshua Tenenbaum (MIT), and Jiajun Wu (MIT)</i>	
Calibration of Axial Fisheye Cameras Through Generic Virtual Central Models	4039
<i>Pierre-André Brousseau (Université de Montréal) and Sebastien Roy (Montreal University)</i>	
Micro-Baseline Structured Light	4048
<i>Vishwanath Saragadam Raja Venkata (Carnegie Mellon University), Jian Wang (Snap), Shree Nayar (Snap), and Mohit Gupta (University of Wisconsin-Madison. USA)</i>	
lambda-Net: Reconstruct Hyperspectral Images From a Snapshot Measurement	4058
<i>Xin Miao (The University of Texas at Arlington), Xin Yuan (Bell Labs), Yunchen Pu (facebook), and Vassilis Athitsos (University of Texas at Arlington)</i>	
Deep Depth From Aberration Map	4069
<i>Masako Kashiwagi (Toshiba), Nao Mishima (Toshiba Research and Development Center), Tatsuo Kozakaya (Toshiba), and Shinsaku Hiura (Hiroshima city univ.)</i>	
A Dataset of Multi-Illumination Images in the Wild	4079
<i>Lukas Murmann (MIT CSAIL), Michael Gharbi (Adobe), Miika Aittala (MIT), and Fredo Durand (MIT)</i>	
Monocular Neural Image Based Rendering With Continuous View Control	4089
<i>Jie Song (ETH Zurich), Xu Chen (ETH Zürich), and Otmar Hilliges (ETH Zurich)</i>	
Multi-View Image Fusion	4100
<i>Marc Comino Trinidad (Universitat Politècnica de Catalunya), Ricardo Martin-Brualla (Google), Florian Kainz (Google), and Janne Kontkanen (Google)</i>	

Low-Level & Optimization

Enhancing Low Light Videos by Exploring High Sensitivity Camera Noise	4110
<i>Wei Wang (Nanjing University), Xin Chen (Nanjing University), Cheng Yang (Nanjing University), Xiang Li (Nanjing University), Xuemei Hu (Nanjing University), and Tao Yue (Nanjing University)</i>	
Deep Restoration of Vintage Photographs From Scanned Halftone Prints	4119
<i>Qifan Gao (Shanghai Jiao Tong University), Xiao Shu (McMaster University; Shanghai Jiao Tong University), and Xiaolin Wu (McMaster University; Shanghai Jiao Tong University)</i>	
Context-Aware Image Matting for Simultaneous Foreground and Alpha Estimation	4129
<i>Qiqi Hou (Portland State University) and Feng Liu (Portland State University)</i>	
CFSNet: Toward a Controllable Feature Space for Image Restoration	4139
<i>Wei Wang (Tsinghua University), Ruiming Guo (The Chinese University of Hong Kong), Yapeng Tian (University of Rochester), and Wenming Yang (Tsinghua University)</i>	
Deep Blind Hyperspectral Image Fusion	4149
<i>Wu Wang (Xiamen University), Weihong Zeng (Xiamen University), Yue Huang (Xiamen University), Xinghao Ding (Xiamen University), and John Paisley (Columbia University)</i>	
Fully Convolutional Pixel Adaptive Image Denoiser	4159
<i>Sungmin Cha (Sungkyunkwan University) and Taesup Moon (Sungkyunkwan University)</i>	
Coherent Semantic Attention for Image Inpainting	4169
<i>Hongyu Liu (Hunan University), Bin Jiang (Hunan University), Yi Xiao (Hunan University), and Chao Yang (Hunan University)</i>	
Embedded Block Residual Network: A Recursive Restoration Model for Single-Image Super-Resolution	4179
<i>Yajun Qiu (Yunnan University), Ruxin Wang (Union Vision Innovation Co Ltd.), Dapeng Tao (Yunnan University), and Jun Cheng (Shenzhen Institutes of Advanced Technology. Chinese Academy of Sciences)</i>	
Fast Image Restoration With Multi-Bin Trainable Linear Units	4189
<i>Shuhang Gu (ETH Zurich), Wen Li (ETH Zurich), Luc Van Gool (ETH Zurich), and Radu Timofte (ETH Zurich)</i>	

Scene Understanding

Counting With Focus for Free	4199
<i>Zenglin Shi (University of Amsterdam), Pascal Mettes (University of Amsterdam), and Cees Snoek (University of Amsterdam)</i>	
SynDeMo: Synergistic Deep Feature Alignment for Joint Learning of Depth and Ego-Motion	4209
<i>Behzad Bozorgtabar (EPFL), Mohammad Saeed Rad (EPFL), Dwarikanath Mahapatra (IBM Research Australia), and Jean-Philippe Thiran (École Polytechnique Fédérale de Lausanne)</i>	
Diverse Image Synthesis From Semantic Layouts via Conditional IMLE	4219
<i>Ke Li (UC Berkeley), Tianhao Zhang (Nanjing University), and Jitendra Malik (University of California at Berkley)</i>	

Towards Bridging Semantic Gap to Improve Semantic Segmentation	4229
<i>Yanwei Pang (Tianjin University), Yazhao Li (Tianjin University), Jianbing Shen (Inception Institute of Artificial Intelligence), and Ling Shao (Inception Institute of Artificial Intelligence)</i>	

Language & Reasoning

Generating Diverse and Descriptive Image Captions Using Visual Paraphrases	4239
<i>Lixin Liu (Peking University), Jiajun Tang (Peking University), Xiaojun Wan (Peking University), and Zongming Guo (Peking University)</i>	
Learning to Collocate Neural Modules for Image Captioning	4249
<i>Xu Yang (Nanyang Technological University), Hanwang Zhang (Nanyang Technological University), and Jianfei Cai (Nanyang Technological University)</i>	
Sequential Latent Spaces for Modeling the Intention During Diverse Image Captioning	4260
<i>Jyoti Aneja (University of Illinois. Urbana-Champaign), Harsh Agrawal (Georgia Institute of Technology), Dhruv Batra (Georgia Tech & Facebook AI Research), and Alexander Schwing (UIUC)</i>	
Why Does a Visual Question Have Different Answers?	4270
<i>Nilava Bhattacharya (University of Texas at Austin), Qing Li (University of California. Los Angeles), and Danna Gurari (University of Texas at Austin)</i>	
G3graphGround: Graph-Based Language Grounding	4280
<i>Mohit Bajaj (University of British Columbia), Lanjun Wang (Huawei Technologies Canada Co. Ltd.), and Leonid Sigal (University of British Columbia)</i>	
Scene Text Visual Question Answering	4290
<i>Ali Furkan Biten (Computer Vision Center), Rubèn Tito (Centre de Visió per Computador. CVC), Andrés Mafla (Computer Vision Centre), Lluís Gomez (Universitat Autònoma de Barcelona), Marçal Rusiñol (Computer Vision Center. UAB), C.V. Jawahar (IIIT-Hyderabad), Ernest Valveny (Universitat Autònoma de Barcelona), and Dimosthenis Karatzas (Computer Vision Centre)</i>	

3D From Multiview & Sensors

Unsupervised Collaborative Learning of Keyframe Detection and Visual Odometry Towards Monocular Deep SLAM	4301
<i>Lu Sheng (Beihang University), Dan Xu (University of Oxford), Wanli Ouyang (The University of Sydney), and Xiaogang Wang (Chinese University of Hong Kong. Hong Kong)</i>	
MVSCRf: Learning Multi-View Stereo With Conditional Random Fields	4311
<i>Youze Xue (Tsinghua University), Jiansheng Chen (Tsinghua University), Weitao Wan (Tsinghua University), Yiqing Huang (Tsinghua University), Cheng Yu (Tsinghua University), Tianpeng Li (Tsinghua University), and Jiayu Bao (Tsinghua University)</i>	

Neural-Guided RANSAC: Learning Where to Sample Model Hypotheses	4321
<i>Eric Brachmann (Heidelberg University) and Carsten Rother (University of Heidelberg)</i>	
Efficient Learning on Point Clouds With Basis Point Sets	4331
<i>Sergey Prokudin (Max Planck Institute for Intelligent Systems), Christoph Lassner (Facebook Reality Labs), and Javier Romero (Amazon)</i>	
Cross View Fusion for 3D Human Pose Estimation	4341
<i>Haibo Qiu (University of Science and Technology of China), Chunyu Wang (Microsoft Research asia), Jingdong Wang ((), Naiyan Wang (TuSimple), and Wenjun Zeng (Microsoft Research)</i>	
Shape-Aware Human Pose and Shape Reconstruction Using Multi-View Images	4351
<i>Junbang Liang (University of Maryland. College Park) and Ming Lin (UMD-CP & UNC-CH)</i>	
Monocular Piecewise Depth Estimation in Dynamic Scenes by Exploiting Superpixel Relations	4362
<i>Di Yan (Tsinghua university), Henrique Morimitsu (Tsinghua University), Shan Gao (nwpu), and Xiangyang Ji (Tsinghua University)</i>	
Is This the Right Place? Geometric-Semantic Pose Verification for Indoor Visual Localization	4372
<i>Hajime Taira (Tokyo Institute of Technology), Ignacio Rocco (Inria), Jiri Sedlar (CVUT), Masatoshi Okutomi (Tokyo Institute of Technology), Josef Sivic (Inria and Czech Technical University), Tomas Pajdla (Czech Technical University in Prague), Torsten Sattler (Chalmers University of Technology), and Akihiko Torii (Tokyo Institute of Technology. Japan)</i>	
DeepPruner: Learning Efficient Stereo Matching via Differentiable PatchMatch	4383
<i>Shivam Duggal (Delhi Technological University), Shenlong Wang (Uber ATG. University of Toronto), Wei-Chiu Ma (MIT), Rui Hu (Uber), and Raquel Urtasun (Uber ATG)</i>	

Image & Video Synthesis

Convolutional Sequence Generation for Skeleton-Based Action Synthesis	4393
<i>Sijie Yan (Chinese University of Hong Kong), Zhizhong Li (CUHK), Yuanjun Xiong (Amazon), Huahan Yan (Trootopia), and Dahua Lin (The Chinese University of Hong Kong)</i>	
Onion-Peel Networks for Deep Video Completion	4402
<i>Seoung Wug Oh (Yonsei Univeristy), Sungho Lee (Yonsei University), Joon-Young Lee (Adobe Research), and Seon Joo Kim (Yonsei Univ.)</i>	
Copy-and-Paste Networks for Deep Video Inpainting	4412
<i>Sungho Lee (Yonsei University), Seoung Wug Oh (Yonsei Univeristy), Daeyeun Won (Hyundai-mnsoft), and Seon Joo Kim (Yonsei Univ.)</i>	
Content and Style Disentanglement for Artistic Style Transfer	4421
<i>Dmytro Kotovenko (Heidelberg University), Artsiom Sanakoyeu (Heidelberg University), Sabine Lang (Heidelberg University), and Bjorn Ommer (Heidelberg University)</i>	

Oral 3.1A

Generative Modeling & Synthesis

Image2StyleGAN: How to Embed Images Into the StyleGAN Latent Space?	4431
<i>Rameen Abdal (KAUST), Yipeng Qin (KAUST), and Peter Wonka (KAUST)</i>	
Controllable Artistic Text Style Transfer via Shape-Matching GAN	4441
<i>Shuai Yang (Peking University), Zhangyang Wang (TAMU), Zhaowen Wang (Adobe Research), Ning Xu (Adobe Research), Jiaying Liu (Peking University), and Zongming Guo (Peking University)</i>	
Understanding Generalized Whitening and Coloring Transform for Universal Style Transfer	4451
<i>Tai-Yin Chiu (University of Texas)</i>	
Learning Implicit Generative Models by Matching Perceptual Features	4460
<i>Cicero Nogueira Dos Santos (IBM Research), Youssef Mroueh (IBM Research), Inkit Padhi (IBM Research), and Pierre Dognin (IBM Research)</i>	
Free-Form Image Inpainting With Gated Convolution	4470
<i>Jiahui Yu (UIUC), Zhe Lin (Adobe Research), Jimei Yang (Adobe), Xiaohui Shen (ByteDance AI Lab), Xin Lu (Adobe), and Thomas Huang (UIUC)</i>	
FiNet: Compatible and Diverse Fashion Image Inpainting	4480
<i>Xintong Han (Malong Technologies), Zuxuan Wu (UMD), Weilin Huang (Malong Technologies), Matthew Scott (Malong Technologies), and Larry Davis (University of Maryland)</i>	
InGAN: Capturing and Retargeting the “DNA” of a Natural Image	4491
<i>Assaf Shocher (Weizmann Institute of Science), Shai Bagon (Weizmann Institute of Science), Phillip Isola (MIT), and Michal Irani (Weizmann Institute. Israel)</i>	
Seeing What a GAN Cannot Generate	4501
<i>David Bau (MIT), Jun-Yan Zhu (MIT), Jonas Wulff (MIT), William Peebles (MIT), Bolei Zhou (CUHK), Hendrik Strobelt (IBM Research), and Antonio Torralba (MIT)</i>	
COCO-GAN: Generation by Parts via Conditional Coordinating	4511
<i>Chieh Hubert Lin (National Tsing Hua University), Chia-Che Chang (MediaTek), Yu-Sheng Chen (National Taiwan University), Da-Cheng Juan (Google), Wei Wei (Google), and Hwann-Tzong Chen (National Tsing Hua University)</i>	
Neural Turtle Graphics for Modeling City Road Layouts	4521
<i>Hang Chu (University of Toronto), Daiqing Li (NVIDIA), David Acuna (University of Toronto), Amlan Kar (University of Toronto), Maria Shugrina (University of Toronto), Xinkai Wei (University of Waterloo), Ming-Yu Liu (NVIDIA), Antonio Torralba (MIT), and Sanja Fidler (University of Toronto. NVIDIA)</i>	

Texture Fields: Learning Texture Representations in Function Space	4530
<i>Michael Oechsle (MPI-IS, University of Tuebingen and ETAS GmbH), Lars Mescheder (MPI-IS and University of Tuebingen), Michael Niemeyer (MPI-IS and University of Tuebingen), Thilo Strauss (ETAS GmbH), and Andreas Geiger (MPI-IS and University of Tuebingen)</i>	
PointFlow: 3D Point Cloud Generation With Continuous Normalizing Flows	4540
<i>Guandao Yang (Cornell University), Xun Huang (Cornell University), Zekun Hao (Cornell University), Ming-Yu Liu (NVIDIA), Serge Belongie (Cornell University), and Bharath Hariharan (Cornell University)</i>	
Meta-Sim: Learning to Generate Synthetic Datasets	4550
<i>Amlan Kar (University of Toronto), Aayush Prakash (Nvidia), Ming-Yu Liu (NVIDIA), Eric Cameracci (Nvidia), Justin Yuan (Nvidia), Matt Rusiniak (NVIDIA), David Acuna (University of Toronto), Antonio Torralba (MIT), and Sanja Fidler (University of Toronto, NVIDIA)</i>	
Specifying Object Attributes and Relations in Interactive Scene Generation	4560
<i>Oron Ashual (Tel-Aviv University) and Lior Wolf (Tel Aviv University, Israel)</i>	
SinGAN: Learning a Generative Model From a Single Natural Image	4569
<i>Tamar Rott Shaham (Technion), Tali Dekel (Google), and Tomer Michaeli (Technion)</i>	

Oral 3.1B

Vision, Language, & Text

VaTeX: A Large-Scale, High-Quality Multilingual Dataset for Video-and-Language Research	4580
<i>Xin Wang (University of California, Santa Barbara), Jiawei Wu (University of California, Santa Barbara), Junkun Chen (Fudan University), Lei Li (ByteDance AI Lab), Yuan-Fang Wang (UC Santa Barbara), and William Yang Wang (UC Santa Barbara)</i>	
A Graph-Based Framework to Bridge Movies and Synopses	4591
<i>Yu Xiong (The Chinese University of HK), Qingqiu Huang (The Chinese University of Hong Kong), Lingfeng Guo (UCSD), Hang Zhou (The Chinese University of Hong Kong), Bolei Zhou (CUHK), and Dahua Lin (The Chinese University of Hong Kong)</i>	
From Strings to Things: Knowledge-Enabled VQA Model That Can Read and Reason	4601
<i>Ajeet Kumar Singh (TCS Research), Anand Mishra (IISc Bangalore), Shashank Shekhar (Indian Institute Of Science), and Anirban Chakraborty (Indian Institute of Science)</i>	
Counterfactual Critic Multi-Agent Training for Scene Graph Generation	4612
<i>Long Chen (Zhejiang University), Hanwang Zhang (Nanyang Technological University), Jun Xiao (Zhejiang University), Xiangnan He (University of Science and Technology of China), Shiliang Pu (Hikvision Research Institute), and Shih-Fu Chang (Columbia University)</i>	
Robust Change Captioning	4623
<i>Dong Huk Park (UC Berkeley), Trevor Darrell (UC Berkeley), and Anna Rohrbach (UC Berkeley)</i>	

Attention on Attention for Image Captioning	4633
<i>Lun Huang (Peking University), Wenmin Wang (Peking University), Jie Chen (Peng Cheng Laboratory), and Xiao-Yong Wei (Peng Cheng Laboratory)</i>	
Dynamic Graph Attention for Referring Expression Comprehension	4643
<i>Sibe Yang (The University of Hong Kong), Guanbin Li (Sun Yat-sen University), and Yizhou Yu (Deepwise AI Lab)</i>	
Visual Semantic Reasoning for Image-Text Matching	4653
<i>Kunpeng Li (Northeastern University), Yulun Zhang (Northeastern University), Kai Li (Northeastern University), Yuanyuan Li (Northeastern University), and Yun Fu (Northeastern University)</i>	
Phrase Localization Without Paired Training Examples	4662
<i>Josiah Wang (Imperial College London) and Lucia Specia (University of Sheffield)</i>	
Learning to Assemble Neural Module Tree Networks for Visual Grounding	4672
<i>Daqing Liu (University of Science and Technology of China), Hanwang Zhang (Nanyang Technological University), Zheng-Jun Zha (University of Science and Technology of China), and Feng Wu (University of Science and Technology of China)</i>	
A Fast and Accurate One-Stage Approach to Visual Grounding	4682
<i>Zhengyuan Yang (University of Rochester), Boqing Gong (Google / ICSI Berkeley), Liwei Wang (Tencent AI Lab), Wenbing Huang (Tencent AI Lab), Dong Yu (Tencent), and Jiebo Luo (U. Rochester)</i>	
Zero-Shot Grounding of Objects From Natural Language Queries	4693
<i>Arka Sadhu (University of Southern California), Kan Chen (Facebook Research), and Ram Nevatia (U of Southern California)</i>	
Towards Unconstrained End-to-End Text Spotting	4703
<i>Siyang Qin (Google Inc.), Alessandro Bissaco (Google Inc.), Michalis Raptis (Google Inc.), Yasuhisa Fujii (Google Inc.), and Ying Xiao (Google)</i>	
What Is Wrong With Scene Text Recognition Model Comparisons? Dataset and Model Analysis	4714
<i>Jeonghun Baek (Clova AI Research. NAVER Corp.), Geewook Kim (Kyoto University), Junyeop Lee (Clova AI Research. NAVER Corp.), Sungrae Park (Clova AI Research. NAVER Corp.), Dongyoon Han (Clova AI Research. NAVER Corp.), Sangdoon Yun (Clova AI Research. NAVER Corp.), Seong Joon Oh (Clova AI Research. LINE Plus), and Hwalsuk Lee (Clova AI Research. NAVER Corp.)</i>	

Poster 3.1

Deep Learning

Sparse and Imperceivable Adversarial Attacks	4723
<i>Francesco Croce (University of Tübingen) and Matthias Hein (University of Tübingen)</i>	

Enhancing Adversarial Example Transferability With an Intermediate Level Attack	4732
<i>Qian Huang (Cornell University), Isay Katsman (Cornell University), Zeqi Gu (Cornell University), Horace He (Cornell University), Serge Belongie (Cornell University), and Ser-Nam Lim (Facebook AI)</i>	
Implicit Surface Representations As Layers in Neural Networks	4742
<i>Mateusz Michalkiewicz (Queensland University of Technology), Jhony Kaesemodel Pontes (Queensland University of Technology), Dominic Jack (Queensland University of Technology), Mahsa Baktashmottlagh (University of Queensland), and Anders Eriksson (University of Queensland)</i>	
A Tour of Convolutional Networks Guided by Linear Interpreters	4752
<i>Pablo Navarrete Michelini (BOE Technology Group Co.. Ltd.), Hanwen Liu (BOE Technology Group Co.. Ltd.), Yunhua Lu (BOE Technology Group Co.. Ltd.), and Xingqun Jiang (BOE Technology Group Co.. Ltd.)</i>	
Small Steps and Giant Leaps: Minimal Newton Solvers for Deep Learning	4762
<i>Joao Henriques (University of Oxford), Sebastien Ehrhardt (University of Oxford), Samuel Albanie (University of Oxford), and Andrea Vedaldi (Oxford University)</i>	
Semantic Adversarial Attacks: Parametric Transformations That Fool Deep Classifiers	4772
<i>Ameya Joshi (Iowa State University), Amitangshu Mukherjee (Iowa State University), Soumik Sarkar (Iowa State University), and Chinmay Hegde (Iowa State University)</i>	
Hilbert-Based Generative Defense for Adversarial Examples	4783
<i>Yang Bai (TBSI(Tsinghua Berkley Shenzhen Institute), Yan Feng (Tsinghua University), Yisen Wang (JD AI Research), Tao Dai (Tsinghua University), Shutao Xia (Tsinghua University), and Yong Jiang (Tsinghua University)</i>	
On the Efficacy of Knowledge Distillation	4793
<i>Jang Hyun Cho (Cornell University) and Bharath Hariharan (Cornell University)</i>	
Sym-Parameterized Dynamic Inference for Mixed-Domain Image Translation	4802
<i>Simyung Chang (Seoul National University), Seonguk Park (Seoul National University), John Yang (Seoul National University), and Nojun Kwak (Seoul National University)</i>	
Better and Faster: Exponential Loss for Image Patch Matching	4811
<i>Shuang Wang (Xidian University), Yanfeng Li (Xidian University), Xuefeng Liang (Xidian University), Dou Quan (School of Artificial Intelligence, Xidian University, China), Bowu Yang (Xidian University), Shaowei Wei (Xidian University), and Licheng Jiao (Xidian University)</i>	
Physical Adversarial Textures That Fool Visual Object Tracking	4821
<i>Rey Wiyatno (Element AI) and Anqi Xu (Element AI)</i>	
Wasserstein GAN With Quadratic Transport Cost	4831
<i>Huidong Liu (Stony Brook University), Xianfeng Gu (Stony Brook University), and Dimitris Samaras (Stony Brook University)</i>	

Scalable Verified Training for Provably Robust Image Classification	4841
<i>Sven Gowal (DeepMind), Krishnamurthy Dvijotham (DeepMind), Robert Stanforth (Deepmind), Rudy Bunel (University of Oxford), Chongli Qin (DeepMind), Jonathan Uesato (DeepMind), Relja Arandjelovic (DeepMind), Timothy Arthur Mann, and Pushmeet Kohli (DeepMind)</i>	
Differentiable Soft Quantization: Bridging Full-Precision and Low-Bit Neural Networks	4851
<i>Ruihao Gong (Beihang University), Xianglong Liu (Beihang University), Shenghu Jiang (Beihang University), Tianxiang Li (Beijing Institute of Technology), Peng Hu (Beihang University), Jiazhen Lin (Beihang University), Fengwei Yu (Sensetime Research), and Junjie Yan (Sensetime Group Limited)</i>	
The LogBarrier Adversarial Attack: Making Effective Use of Decision Boundary Information	4861
<i>Chris Finlay (McGill University), Aram-Alexandre Pooladian (McGill University), and Adam Oberman (McGill University)</i>	
Proximal Mean-Field for Neural Network Quantization	4870
<i>Thalaiyasingam Ajanthan (ANU), Puneet Dokania (University of Oxford), Richard Hartley (Australian National University. Australia), and Philip Torr (University of Oxford)</i>	
Improving Adversarial Robustness via Guided Complement Entropy	4880
<i>Hao-Yun Chen (National Tsing Hua University), Jhao-Hong Liang (National Tsing Hua University), Shih-Chieh Chang (National Tsing Hua University), Jia-Yu Pan (Google), Yu-Ting Chen (Google), Wei Wei (Google), and Da-Cheng Juan (Google)</i>	
A Geometry-Inspired Decision-Based Attack	4889
<i>Yujia Liu (University of Science and Technology of China), Seyed-Mohsen Moosavi-Dezfooli (EPFL), and Pascal Frossard (EPFL)</i>	
Universal Perturbation Attack Against Image Retrieval	4898
<i>Jie Li (Xiamen University), Rongrong Ji (Xiamen University. China), Hong Liu (Xiamen University), Xiaopeng Hong (Xi'an Jiaotong University), Yue Gao (Tsinghua University), and Qi Tian (Huawei Noah's Ark Lab)</i>	
Bayesian Optimized 1-Bit CNNs	4908
<i>Jiaxin Gu (Beihang University), Junhe Zhao (Beihang University), Xiaolong Jiang (Beihang University), Baochang Zhang (Beihang University), Jianzhuang Liu (Huawei Noah's Ark Lab), Guodong Guo (Baidu), and Rongrong Ji (Xiamen University. China)</i>	
Rethinking ImageNet Pre-Training	4917
<i>Kaiming He (Facebook AI Research), Ross Girshick (FAIR), and Piotr Dollar (FAIR)</i>	
Defending Against Universal Perturbations With Shared Adversarial Training	4927
<i>Chaithanya Kumar Mummadi (Robert Bosch GmbH), Thomas Brox (University of Freiburg), and Jan Hendrik Metzen (Robert Bosch GmbH)</i>	
Adaptive Activation Thresholding: Dynamic Routing Type Behavior for Interpretability in Convolutional Neural Networks	4937
<i>Yiyou Sun (University of Wisconsin Madison), Sathya Ravi (University of Wisconsin-Madison), and Vikas Singh (University of Wisconsin-Madison USA)</i>	

XRAI: Better Attributions Through Regions	4947
<i>Andrei Kapishnikov (Google), Tolga Bolukbasi (Google), Fernanda Viegas (Google), and Michael Terry (Google)</i>	
Guessing Smart: Biased Sampling for Efficient Black-Box Adversarial Attacks	4957
<i>Thomas Brunner (fortiss GmbH), Frederik Diehl (Fortiss GmbH), Michael Truong Le (fortiss GmbH), and Alois Knoll (Robotics and Embedded Systems)</i>	

Recognition

Mask-Guided Attention Network for Occluded Pedestrian Detection	4966
<i>Yanwei Pang (Tianjin University), Jin Xie (Tianjin University), Muhammad Haris Khan (Inception Institute of Artificial Intelligence), Rao Muhammad Anwer (Inception Institute of Artificial Intelligence), Fahad Shahbaz Khan (Inception Institute of Artificial Intelligence), and Ling Shao (Inception Institute of Artificial Intelligence)</i>	
Spectral Feature Transformation for Person Re-Identification	4975
<i>Chuanchen Luo (CASIA), Yuntao Chen (CASIA), Naiyan Wang (TuSimple), and Zhao-Xiang Zhang (Chinese Academy of Sciences. China)</i>	
Permutation-Invariant Feature Restructuring for Correlation-Aware Image Set-Based Recognition	4985
<i>Xiaofeng Liu (CMU), Zhenhua Guo (Graduate School at Shenzhen. Tsinghua University. China), Site Li (Carnegie Mellon University), Ping Jia (CIOMP), Lingsheng Kong (Changchun Institute of Optics. Fine Mechanics and Physics), Jane You (HK Poly U), and B. V. K. Vijaya Kumar (CMU. USA)</i>	
Improving Pedestrian Attribute Recognition With Weakly-Supervised Multi-Scale Attribute-Specific Localization	4996
<i>Chufeng Tang (Tsinghua University), Lu Sheng (Beihang University), Zhao-Xiang Zhang (Chinese Academy of Sciences. China), and Xiaolin Hu (Tsinghua University)</i>	
Correlation Congruence for Knowledge Distillation	5006
<i>Baoyun Peng (National University of Defense Technology), Xiao Jin (SenseTime Group Limited), Dongsheng Li (School of Computer Science. National University of Defense Technology), Shunfeng Zhou (SenseTime), Yichao Wu (SenseTime Group Limited), Jiaheng Liu (Beihang University), Zhaoning Zhang (School of Computer Science. National University of Defense Technology), and Yu Liu (The Chinese University of Hong Kong)</i>	
Dynamic Curriculum Learning for Imbalanced Data Classification	5016
<i>Yiru Wang (SenseTime Group Limited), Weihao Gan (SenseTime Group Limited), Jie Yang (SenseTime Group Limited), Wei Wu (SenseTime Group Limited), and Junjie Yan (SenseTime Group Limited)</i>	
Video Face Clustering With Unknown Number of Clusters	5026
<i>Makarand Tapaswi (INRIA), Marc Law (NVIDIA), and Sanja Fidler (University of Toronto. NVIDIA)</i>	

Targeted Mismatch Adversarial Attack: Query With a Flower to Retrieve the Tower	5036
<i>Giorgos Tolias (Czech Technical University in Prague. Faculty of Electrical Engineering. Visual Recognition Group), Filip Radenovic (Visual Recognition Group. CTU Prague), and Ondrej Chum (Vision Recognition Group. Czech Technical University in Prague)</i>	
Fashion++: Minimal Edits for Outfit Improvement	5046
<i>Wei-Lin Hsiao (UT-Austin), Isay Katsman (Cornell University), Chao-Yuan Wu (UT Austin), Devi Parikh (Georgia Tech & Facebook AI Research), and Kristen Grauman (Facebook AI Research & UT Austin)</i>	
Semi-Supervised Pedestrian Instance Synthesis and Detection With Mutual Reinforcement	5056
<i>Si Wu (South China University of Technology), Sihao Lin (South China University of Technology), Wenhao Wu (South China University of Technology), Mohamed Azzam (City University of Hong Kong), and Hau San Wong (City University of Hong Kong)</i>	
SILCO: Show a Few Images, Localize the Common Object	5066
<i>Tao Hu (University of Amsterdam), Pascal Mettes (University of Amsterdam), Jia-Hong Huang (University of Amsterdam), and Cees Snoek (University of Amsterdam)</i>	
A Deep Step Pattern Representation for Multimodal Retinal Image Registration	5076
<i>Jimmy Lee (Cixi Institute of Biomedical Engineering. Chinese Academy of Sciences), Peng Liu (University of Electronic Science and Technology of China), Jun Cheng (Cixi Institute of Biomedical Engineering. Chinese Academy of Sciences), and Huazhu Fu (Inception Institute of Artificial Intelligence)</i>	
Deep Graphical Feature Learning for the Feature Matching Problem	5086
<i>Zhen Zhang (NUS) and Wee Sun Lee (NUS)</i>	
Minimum Delay Object Detection From Video	5096
<i>Dong Lao (KAUST) and Ganesh Sundaramoorthi (Kaust)</i>	
Learning With Average Precision: Training Image Retrieval With a Listwise Loss	5106
<i>Jerome Revaud (Naver Labs Europe), Jon Almazan (V7 Labs), Rafael Rezende (Naver Labs), and Cesar De Souza (NAVER LABS Europe)</i>	
Learning to Find Common Objects Across Few Image Collections	5116
<i>Amirreza Shaban (Georgia Institute of Technology), Amir Rahimi (Australian National University), Shray Bansal (Georgia Institute of Technology), Stephen Gould (Australian National University. Australia), Byron Boots (Georgia Institute of Technology), and Richard Hartley (Australian National University. Australia)</i>	
Weakly Aligned Cross-Modal Learning for Multispectral Pedestrian Detection	5126
<i>Lu Zhang (Institute of Automation. Chinese Academy of Sciences), Xiangyu Zhu (Chinese Academy of Science), Xiangyu Chen (Renmin University of China), Xu Yang (Chinese Academy of Sciences), Zhen Lei (NLPR. CASIA. China), and Zhiyong Liu (State Key Lab of Management and Control for Complex Systems. Institute of Automation. Chinese Academy of Sciences)</i>	
Deep Self-Learning From Noisy Labels	5137
<i>Jiangfan Han (The Chinese University of Hong Kong), Ping Luo (The Chinese University of Hong Kong), and Xiaogang Wang (Chinese University of Hong Kong. Hong Kong)</i>	

DSCConv: Efficient Convolution Operator	5147
<i>Marcelo Gennari Do Nascimento (University of Oxford), Victor Prisacariu (University of Oxford), and Roger Fawcett (Intel)</i>	
Once a MAN: Towards Multi-Target Attack via Learning Multi-Target Adversarial Network Once	5157
<i>Jiangfan Han (The Chinese University of Hong Kong), Xiaoyi Dong (University of Science and Technology of China), Ruimao Zhang (The Chinese University of Hong Kong), Dongdong Chen (university of science and technology of china), Weiming Zhang (University of Science and Technology of China), Nenghai Yu (University of Science and Technology of China), Ping Luo (The Chinese University of Hong Kong), and Xiaogang Wang (Chinese University of Hong Kong. Hong Kong)</i>	

Segmentation, Grouping, & Shape

Explicit Shape Encoding for Real-Time Instance Segmentation	5167
<i>Wenqiang Xu (Shanghai Jiaotong University), Haiyang Wang (SHANG HAI JIAO TONG UNIVERSITY), Fubo Qi (Shanghai JiaoTong University), and Cewu Lu (Shanghai Jiao Tong University)</i>	
IMP: Instance Mask Projection for High Accuracy Semantic Segmentation of Things	5177
<i>Cheng-Yang Fu (UNC-Chapel Hill), Tamara Berg (University on North carolina), and Alexander Berg (University of North Carolina. USA)</i>	
Video Instance Segmentation	5187
<i>Linjie Yang (ByteDance AI Lab), Yuchen Fan (University of Illinois at Urbana-Champaign), and Ning Xu (Adobe Research)</i>	
Attention Bridging Network for Knowledge Transfer	5197
<i>Kunpeng Li (Northeastern University), Yulun Zhang (Northeastern University), Kai Li (Northeastern University), Yuanyuan Li (Northeastern University), and Yun Fu (Northeastern University)</i>	
Self-Supervised Difference Detection for Weakly-Supervised Semantic Segmentation	5207
<i>Wataru Shimoda (The University of Electro-Communications. Tokyo) and Keiji Yanai (Univ. Electro-Comm.. Tokyo)</i>	
SPGNet: Semantic Prediction Guidance for Scene Parsing	5217
<i>Bowen Cheng (UIUC), Liang-Chieh Chen (Google Inc.), Yunchao Wei (UIUC), Yukun Zhu (Google Inc.), Zilong Huang (Huazhong Univ. of Science and Technology), Jinjun Xiong (IBM Thomas J. Watson Research Center), Thomas Huang (UIUC), Wen-Mei Hwu (University of Illinois at Urbana-Champaign), Honghui Shi (IBM), and Uiu Uiu</i>	
Gated-SCNN: Gated Shape CNNs for Semantic Segmentation	5228
<i>Towaki Takikawa (University of Waterloo), David Acuna (University of Toronto), Varun Jampani (Nvidia Research), and Sanja Fidler (University of Toronto. NVIDIA)</i>	
DensePoint: Learning Densely Contextual Representation for Efficient Point Cloud Processing	5238
<i>Yongcheng Liu (Institute of Automation. Chinese Academy of Sciences), Bin Fan (Institute of Automation. Chinese Academy of Sciences. China), Gaofeng Meng (Chinese Academy of Sciences), Jiwen Lu (Tsinghua University), Shiming Xiang (Chinese Academy of Sciences. China), and Chunhong Pan (Institute of Automation. Chinese Academy of Sciences)</i>	

AMP: Adaptive Masked Proxies for Few-Shot Segmentation	5248
<i>Mennatullah Siam (University of Alberta), Boris Oreshkin (Element AI), and Martin Jagersand (University of Alberta)</i>	
Universal Semi-Supervised Semantic Segmentation	5258
<i>Tarun Kalluri (IIIT Hyderabad), Girish Varma (IIIT Hyderabad), Manmohan Chandraker (UC San Diego), and C.V. Jawahar (IIIT-Hyderabad)</i>	

Statistics, Physics, Theory & Datasets

Accelerate Learning of Deep Hashing With Gradient Attention	5270
<i>Long-Kai Huang (NTU, Singapore), Jianda Chen (Nanyang Technological University), and Sinno Pan (NTU, Singapore)</i>	
SVD: A Large-Scale Short Video Dataset for Near-Duplicate Video Retrieval	5280
<i>Qing-Yuan Jiang (Nanjing university), Yi He (Bytedance.com), Gen Li (Bytedance.Inc), Jian Lin (Nanjing university), Lei Li (ByteDance AI Lab), and Wu-Jun Li (Nanjing University)</i>	
Block Annotation: Better Image Annotation With Sub-Image Decomposition	5289
<i>Hubert Lin (Cornell University), Paul Upchurch (Cornell), and Kavita Bala (Cornell)</i>	
Probabilistic Deep Ordinal Regression Based on Gaussian Processes	5300
<i>Yanzhu Liu (Nanyang Technological University), Fan Wang (Nanyang Technological University), and Wai-Kin Adams Kong (Nanyang Technological University)</i>	
Balanced Datasets Are Not Enough: Estimating and Mitigating Gender Bias in Deep Image Representations	5309
<i>Tianlu Wang (University of Virginia), Jieyu Zhao (University of California Los Angeles), Mark Yatskar (University of Washington), Kai-Wei Chang (UCLA), and Vicente Ordonez (University of Virginia)</i>	
Teacher Guided Architecture Search	5319
<i>Pouya Bashivan (Massachusetts Institute of Technology), Mark Tensen (University of Amsterdam), and James Dicarlo (Massachusetts Institute of Technology)</i>	

3D From Single View & RGBD

FACSIMILE: Fast and Accurate Scans From an Image in Less Than a Second	5329
<i>David Smith (Amazon), Matthew Loper (Amazon), Xiaochen Hu (Amazon), Paris Mavroidis (Amazon), and Javier Romero (Amazon)</i>	
Delving Deep Into Hybrid Annotations for 3D Human Recovery in the Wild	5339
<i>Yu Rong (The Chinese University of Hong Kong), Ziwei Liu (The Chinese University of Hong Kong), Cheng Li (SenseTime Research), Kaidi Cao (Stanford University), and Chen Change Loy (Nanyang Technological University)</i>	

Human Mesh Recovery From Monocular Images via a Skeleton-Disentangled Representation	5348
<i>Yu Sun (Harbin Institute of technology), Yun Ye (Intel), Wu Liu (AI Research of JD.com), Wenpeng Gao (Harbin Institute of Technology), Yili Fu (HIT), and Tao Mei (AI Research of JD.com)</i>	
Three-D Safari: Learning to Estimate Zebra Pose, Shape, and Texture From Images “In the Wild”	5358
<i>Silvia Zuffi (IMATI-CNR), Angjoo Kanazawa (University of California Berkeley), Tanya Berger-Wolf (University of Illinois at Chicago), and Michael Black (Max Planck Institute for Intelligent Systems)</i>	
Object-Driven Multi-Layer Scene Decomposition From a Single Image	5368
<i>Helisa Dhama (Technical University of Munich), Nassir Navab (TU Munich, Germany), and Federico Tombari (TUM, Google)</i>	
Occupancy Flow: 4D Reconstruction by Learning Particle Dynamics	5378
<i>Michael Niemeyer (MPI-IS and University of Tuebingen), Lars Mescheder (MPI-IS and University of Tuebingen), Michael Oechsle (MPI-IS, University of Tuebingen and ETAS GmbH), and Andreas Geiger (MPI-IS and University of Tuebingen)</i>	
Joint Monocular 3D Vehicle Detection and Tracking	5389
<i>Hou-Ning Hu (National Tsing Hua University), Qi-Zhi Cai (Sinovation Ventures AI Institute), Dequan Wang (UC Berkeley), Ji Lin (MIT), Min Sun (NTHU), Philipp Kraehenbuehl (UT Austin), Trevor Darrell (UC Berkeley), and Fisher Yu (UC Berkeley)</i>	

Face & Body

Fingerspelling Recognition in the Wild With Iterative Visual Attention	5399
<i>Bowen Shi (Toyota Technological Institute at Chicago), Aurora Martinez Del Rio (University of Chicago), Jonathan Keane (University of Chicago), Diane Brentari (University of Chicago), Greg Shakhnarovich (TTI-Chicago), and Karen Livescu (TTI-Chicago)</i>	
PointAE: Point Auto-Encoder for 3D Statistical Shape and Texture Modelling	5409
<i>Hang Dai (Inception Institute of Artificial Intelligence) and Ling Shao (Inception Institute of Artificial Intelligence)</i>	
Multi-Garment Net: Learning to Dress 3D People From Images	5419
<i>Bharat Bhatnagar (MPI-INF), Garvita Tiwari (MPI-INF), Christian Theobalt (MPI Informatik), and Gerard Pons-Moll (MPII, Germany)</i>	
Skeleton-Aware 3D Human Shape Reconstruction From Point Clouds	5430
<i>Haiyong Jiang (NTU), Jianfei Cai (Nanyang Technological University), and Jianmin Zheng (Nanyang Technological University)</i>	
AMASS: Archive of Motion Capture As Surface Shapes	5441
<i>Naureen Mahmood (Meshcapade GmbH), Nima Ghorbani (Max Planck Institute Tübingen), Nikolaus F. Troje (York University), Gerard Pons-Moll (MPII, Germany), and Michael Black (Max Planck Institute for Intelligent Systems)</i>	
Person-in-WiFi: Fine-Grained Person Perception Using WiFi	5451
<i>Fei Wang (Xi'an Jiaotong University), Sanping Zhou (Xi'an Jiaotong University), Stanislav Panev (Carnegie Mellon University), Jinsong Han (Zhejiang University), and Dong Huang (Carnegie Mellon University)</i>	

FAB: A Robust Facial Landmark Detection Framework for Motion-Blurred Videos	5461
<i>Keqiang Sun (Tsinghua University), Wayne Wu (Tsinghua University. SenseTime Research), Tinghao Liu (SenseTime), Shuo Yang (Amazon), Quan Wang (SenseTime), Qiang Zhou (Tsinghua University), Zuochang Ye (Tsinghua University), and Chen Qian (SenseTime)</i>	
Attentional Feature-Pair Relation Networks for Accurate Face Recognition	5471
<i>Bong-Nam Kang (POSTECH), Yonghyun Kim (POSTECH), Bongjin Jun (StradVision. Inc.), and Daijin Kim (Pohang University of Science and Technology)</i>	

Action & Video

Action Recognition With Spatial-Temporal Discriminative Filter Banks	5481
<i>Brais Martínez Martínez (Amazon. USA), Davide Modolo (Amazon), Yuanjun Xiong (Amazon), and Joseph Tighe (Amazon)</i>	
EPIC-Fusion: Audio-Visual Temporal Binding for Egocentric Action Recognition	5491
<i>Evangelos Kazakos (University of Bristol), Arsha Nagrani (Oxford University), Andrew Zisserman (University of Oxford), and Dima Damen (University of Bristol)</i>	
Weakly-Supervised Action Localization With Background Modeling	5501
<i>Phuc Nguyen (UC Irvine), Deva Ramanan (Carnegie Mellon University), and Charless Fowlkes (UC Irvine)</i>	
Grouped Spatial-Temporal Aggregation for Efficient Action Recognition	5511
<i>Chenxu Luo (Johns Hopkins University) and Alan Yuille (Johns Hopkins University)</i>	
Temporal Structure Mining for Weakly Supervised Action Detection	5521
<i>Tan Yu (Nanyang Technological University), Zhou Ren (Wormpex AI Research), Yuncheng Li (Snap), Enxu Yan (Snap Inc.), Ning Xu (-), and Junsong Yuan (State University of New York at Buffalo. USA)</i>	
Temporal Recurrent Networks for Online Action Detection	5531
<i>Mingze Xu (Indiana University), Mingfei Gao (University of Maryland), Yi-Ting Chen (Honda Research Institute USA), Larry Davis (University of Maryland), and David Crandall (Indiana University)</i>	
StartNet: Online Detection of Action Start in Untrimmed Videos	5541
<i>Mingfei Gao (University of Maryland), Mingze Xu (Indiana University), Larry Davis (University of Maryland), Richard Socher (Salesforce), and Caiming Xiong (Salesforce Research)</i>	
Video Classification With Channel-Separated Convolutional Networks	5551
<i>Du Tran (Facebook Research), Heng Wang (Facebook Research), Matt Feiszli (Facebook Research), and Lorenzo Torresani (Facebook AI Research)</i>	
Predicting the Future: A Jointly Learnt Model for Action Anticipation	5561
<i>Harshala Gammulle (Queensland University of Technology), Simon Denman (Queensland University of Technology. Australia), Sridha Sridharan (QUT), and Clinton Fookes (Queensland University of Technology)</i>	

Low-Level & Optimization

Human-Aware Motion Deblurring	5571
<i>Ziyi Shen (Inception Institute of Artificial Intelligence), Wenguan Wang (Inception Institute of Artificial Intelligence), Xiankai Lu (Inception Institute of Artificial Intelligence), Jianbing Shen (Inception Institute of Artificial Intelligence), Haibin Ling (Temple University), Tingfa Xu (Beijing Institute of Technology), and Ling Shao (Inception Institute of Artificial Intelligence)</i>	
Fast Video Object Segmentation via Dynamic Targeting Network	5581
<i>Lu Zhang (Dalian University of Technology), Zhe Lin (Adobe Research), Jianming Zhang (Adobe Research), Huchuan Lu (Dalian University of Technology), and You He (Naval Aviation University)</i>	
Solving Vision Problems via Filtering	5591
<i>Sean Young (Stanford University), Aous Naman (University of New South Wales), Bernd Girod (Stanford University), and David Taubman (University of New South Wales)</i>	
GAN-Based Projector for Faster Recovery With Convergence Guarantees in Linear Inverse Problems	5601
<i>Ankit Raj (University of Illinois at Urbana-Champaign), Yuqi Li (University of Illinois Urbana-Champaign), and Yoram Bresler (UIUC)</i>	
Scoot: A Perceptual Metric for Facial Sketches	5611
<i>Deng-Ping Fan (Inception Institute of Artificial Intelligence), Shengchuan Zhang (Xiamen University), Yu-Huan Wu (Nankai University), Yun Liu (Nankai University), Ming-Ming Cheng (Nankai University), Bo Ren (Nankai University), Paul Rosin (Cardiff University), and Rongrong Ji (Xiamen University. China)</i>	
Learning Filter Basis for Convolutional Neural Network Compression	5622
<i>Yawei Li (ETH Zurich), Shuhang Gu (ETH Zurich), Luc Van Gool (ETH Zurich), and Radu Timofte (ETH Zurich)</i>	
End-to-End Learning of Representations for Asynchronous Event-Based Data	5632
<i>Daniel Gehrig (University of Zurich & ETH Zurich), Antonio Loquercio (ETH / University of Zurich), Konstantinos Derpanis (Ryerson University), and Davide Scaramuzza (University of Zurich & ETH Zurich. Switzerland)</i>	
ERL-Net: Entangled Representation Learning for Single Image De-Raining	5643
<i>Guoqing Wang (UNSW; Csiro Data61), Changming Sun (CSIRO Data61), and Arcot Sowmya (UNSW)</i>	
Perceptual Deep Depth Super-Resolution	5652
<i>Oleg Voynov (Skoltech), Alexey Artemov (Skoltech), Vage Egiazarian (Skoltech), Alexandr Notchenko (Skoltech), Gleb Bobrovskikh (Higher School of Economics), Evgeny Burnaev (Skoltech), and Denis Zorin (New York University)</i>	

Scene Understanding

3D Scene Graph: A Structure for Unified Semantics, 3D Space, and Camera	5663
<i>Iro Armeni (Stanford University), Zhi-Yang He (Stanford University), Amir Zamir (Stanford. UC Berkeley), Junyoung Gwak (Stanford University), Jitendra Malik (University of California at Berkley), Martin Fischer (Stanford University), and Silvio Savarese (Stanford University)</i>	
Floorplan-Jigsaw: Jointly Estimating Scene Layout and Aligning Partial Scans	5673
<i>Cheng Lin (The University of Hong Kong), Changjian Li (The University of Hong Kong), and Wenping Wang (The University of Hong Kong)</i>	
Enforcing Geometric Constraints of Virtual Normal for Depth Prediction	5683
<i>Wei Yin (University of Adelaide), Yifan Liu (University of Adelaide), Chunhua Shen (University of Adelaide), and Youliang Yan (Huawei)</i>	
Deep Contextual Attention for Human-Object Interaction Detection	5693
<i>Tiancai Wang (Tianjin University), Rao Muhammad Anwer (Inception Institute of Artificial Intelligence), Muhammad Haris Khan (Inception Institute of Artificial Intelligence), Fahad Shahbaz Khan (Inception Institute of Artificial Intelligence), Yanwei Pang (Tianjin University), Ling Shao (Inception Institute of Artificial Intelligence), and Jorma Laaksonen (Aalto University)</i>	
Learning Compositional Neural Information Fusion for Human Parsing	5702
<i>Wenguan Wang (Inception Institute of Artificial Intelligence), Zhijie Zhang (Tianjin University), Siyuan Qi (UCLA), Jianbing Shen (Inception Institute of Artificial Intelligence), Yanwei Pang (Tianjin University), and Ling Shao (Inception Institute of Artificial Intelligence)</i>	
Attentional Neural Fields for Crowd Counting	5713
<i>Anran Zhang (Beihang University), Lei Yue (Beihang University), Jiayi Shen (Beihang University), Fan Zhu (Inception Institute of Artificial Intelligence), Xiantong Zhen (Inception Institute of Artificial Intelligence), Xianbin Cao (Beihang University. China), and Ling Shao (Inception Institute of Artificial Intelligence)</i>	
Understanding Human Gaze Communication by Spatio-Temporal Graph Reasoning	5723
<i>Lifeng Fan (University of California. Los Angeles), Wenguan Wang (Inception Institute of Artificial Intelligence), Song-Chun Zhu (UCLA), Xinyu Tang (USTC), and Siyuan Huang (UCLA)</i>	
Controllable Attention for Structured Layered Video Decomposition	5733
<i>Jean-Baptiste Alayrac (DeepMind), Joao Carreira (DeepMind), Relja Arandjelovic (DeepMind), and Andrew Zisserman (University of Oxford)</i>	
GANalyze: Toward Visual Definitions of Cognitive Image Properties	5743
<i>Lore Goetschalckx (MIT), Alex Andonian (MIT), Aude Oliva (MIT), and Phillip Isola (MIT)</i>	

Language & Reasoning

Saliency-Guided Attention Network for Image-Sentence Matching	5753
<i>Zhong Ji (Tianjin University), Haoran Wang (Tianjin University), Jungong Han (Lancaster University), and Yanwei Pang (Tianjin University)</i>	
CAMP: Cross-Modal Adaptive Message Passing for Text-Image Retrieval	5763
<i>Zihao Wang (SenseTime), Xihui Liu (The Chinese University of Hong Kong), Hongsheng Li (Chinese University of Hong Kong), Lu Sheng (Beihang University), Junjie Yan (SenseTime Group Limited), Xiaogang Wang (Chinese University of Hong Kong, Hong Kong), and Jing Shao (SenseTime)</i>	
ACMM: Aligned Cross-Modal Memory for Few-Shot Image and Sentence Matching	5773
<i>Yan Huang (Institute of Automation, Chinese Academy of Sciences) and Liang Wang (NLPR, China)</i>	
Creativity Inspired Zero-Shot Learning	5783
<i>Mohamed Elhoseiny (KAUST and Baidu) and Mohamed Elfeki (CRCV)</i>	
Generating Easy-to-Understand Referring Expressions for Target Identifications	5793
<i>Mikihiro Tanaka (The University of Tokyo), Takayuki Itamochi (NVIDIA), Kenichi Narioka (DENSO CORPORATION), Ikuro Sato (Denso IT Laboratory), Yoshitaka Ushiku (The University of Tokyo), and Tatsuya Harada (The University of Tokyo / RIKEN)</i>	
Language-Agnostic Visual-Semantic Embeddings	5803
<i>Jonatas Wehrmann (PUCRS), Maurício Armani Lopes (PUCRS), Douglas Souza (PUCRS), and Rodrigo Barros (PUCRS)</i>	
Adversarial Representation Learning for Text-to-Image Matching	5813
<i>Nikolaos Sarafianos (University of Houston), Xiang Xu (University of Houston), and Ioannis Kakadiaris (University of Houston)</i>	
Multi-Modality Latent Interaction Network for Visual Question Answering	5824
<i>Gao Peng (Chinese university of hong kong), Haoxuan You (Tsinghua University), Zhanpeng Zhang (SenseTime Group Limited), Xiaogang Wang (Chinese University of Hong Kong, Hong Kong), and Hongsheng Li (Chinese University of Hong Kong)</i>	

3D From Multiview & Sensors

Key.Net: Keypoint Detection by Handcrafted and Learned CNN Filters	5835
<i>Axel Barroso Laguna (Imperial College London), Edgar Riba (Computer Vision Center, Computer Science Department Universitat Autònoma de Barcelona Bellaterra, Barcelona), Daniel Ponsa (Centre de Visió per Computador - Universitat Autònoma de Barcelona), and Krystian Mikolajczyk (Imperial College London)</i>	
Learning Two-View Correspondences and Geometry Using Order-Aware Network	5844
<i>Jiahui Zhang (Tsinghua University), Dawei Sun (Tsinghua University), Zixin Luo (HKUST), Anbang Yao (Intel Labs China), Lei Zhou (HKUST), Tianwei Shen (HKUST), Yurong Chen (Intel Labs China), Hongen Liao (Tsinghua University), and Long Quan (Hong Kong University of Science and Technology)</i>	

Learning Meshes for Dense Visual SLAM	5854
<i>Michael Bloesch (Deepmind), Tristan Laidlow (Imperial College London), Ronald Clark (Imperial College London), Stefan Leutenegger (Imperial College London), and Andrew Davison (Imperial College London)</i>	
EM-Fusion: Dynamic Object-Level SLAM With Probabilistic Data Association	5864
<i>Michael Strecke (Max Planck Institute for Intelligent Systems) and Joerg Stueckler (Max-Planck-Institute for Intelligent Systems)</i>	
ClusterSLAM: A SLAM Backend for Simultaneous Rigid Body Clustering and Motion Estimation	5874
<i>Jiahui Huang (Tsinghua University), Sheng Yang (Tsinghua University), Zishuo Zhao (Tsinghua University), Yu-Kun Lai (Cardiff University), and Shimin Hu (Tsinghua University)</i>	
Efficient and Robust Registration on the 3D Special Euclidean Group	5884
<i>Uttaran Bhattacharya (University of Maryland. College Park) and Venu Madhav Govindu (Indian Institute of Science)</i>	
Algebraic Characterization of Essential Matrices and Their Averaging in Multiview Settings	5894
<i>Yoni Kasten (Weizmann Institute), Amnon Geifman (Weizmann Institute), Meirav Galun (Weizmann Institute of Science), and Ronen Basri (Weizmann Institute of Science)</i>	

Image & Video Synthesis

Liquid Warping GAN: A Unified Framework for Human Motion Imitation, Appearance Transfer and Novel View Synthesis	5903
<i>Wen Liu (ShanghaiTech University), Zhixin Piao (ShanghaiTech University), Jie Min (ShanghaiTech University), Wenhan Luo (Tencent AI Lab), Lin Ma (Tencent AI Lab), and Shenghua Gao (ShanghaiTech University)</i>	
RelGAN: Multi-Domain Image-to-Image Translation via Relative Attributes	5913
<i>Yu-Jing Lin (National Taiwan University), Po-Wei Wu (National Taiwan University), Che-Han Chang (HTC Research & Healthcare), Edward Chang (HTC Research & Healthcare), and Shih-Wei Liao (National Taiwan University)</i>	
Attribute-Driven Spontaneous Motion in Unpaired Image Translation	5922
<i>Ruizheng Wu (The Chinese University of Hong KONG), Xin Tao (Tencent), Xiaodong Gu (Harbin Institute of Technology. Shenzhen), Xiaoyong Shen (Tencent), and Jiaya Jia (Chinese University of Hong Kong)</i>	
Everybody Dance Now	5932
<i>Caroline Chan (UC Berkeley), Shiry Ginosar (UC Berkeley), Tinghui Zhou (Berkeley), and Alexei Efros (UC Berkeley)</i>	
Multimodal Style Transfer via Graph Cuts	5942
<i>Yulun Zhang (Northeastern University), Chen Fang (ByteDance), Yilin Wang (Adobe), Zhaowen Wang (Adobe Research), Zhe Lin (Adobe Research), Yun Fu (Northeastern University), and Jimei Yang (Adobe)</i>	
A Closed-Form Solution to Universal Style Transfer	5951
<i>Ming Lu (Tsinghua University), Hao Zhao (Tsinghua University), Anbang Yao (Intel Labs China), Yurong Chen (Intel Labs China), Feng Xu (Tsinghua University), and Li Zhang (Tsinghua University)</i>	

Progressive Reconstruction of Visual Structure for Image Inpainting	5961
<i>Jingyuan Li (Wuhan University), Fengxiang He (The University of Sydney), Lefei Zhang (Wuhan University), Bo Du (School of Computer Science. Wuhan University), and Dacheng Tao (University of Sydney)</i>	

Oral 3.2A

Recognition, Detection, & Re-Identification

Variational Adversarial Active Learning	5971
<i>Samrath Sinha (University of Toronto), Sayna Ebrahimi (UC Berkeley), and Trevor Darrell (UC Berkeley)</i>	
Confidence Regularized Self-Training	5981
<i>Yang Zou (Carnegie Mellon University), Zhiding Yu (NVIDIA), Xiaofeng Liu (CMU), B. V. K. Vijaya Kumar (CMU. USA), and Jinsong Wang (General Motors)</i>	
Anchor Loss: Modulating Loss Scale Based on Prediction Difficulty	5991
<i>Serim Ryou (California Institute of Technology), Seong-Gyun Jeong (code42.ai), and Pietro Perona (California Institute of Technology)</i>	
Local Aggregation for Unsupervised Learning of Visual Embeddings	6001
<i>Chengxu Zhuang (Stanford University), Alex Zhai (Stanford University), and Daniel Yamins (Stanford University)</i>	
PR Product: A Substitute for Inner Product in Neural Networks	6012
<i>Zhennan Wang (Shenzhen University), Wenbin Zou (Shenzhen University), and Chen Xu (Shenzhen University)</i>	
CutMix: Regularization Strategy to Train Strong Classifiers With Localizable Features	6022
<i>Sangdoon Yun (Clova AI Research. NAVER Corp.), Dongyoon Han (Clova AI Research. NAVER Corp.), Sanghyuk Chun (Clova AI Research. NAVER Corp.), Seong Joon Oh (Clova AI Research. LINE Plus), Youngjoon Yoo (Clova AI Research. NAVER Corp.), and Junsuk Choe (Yonsei University)</i>	
Towards Interpretable Object Detection by Unfolding Latent Structures	6032
<i>Tianfu Wu (NC State University) and Xi Song (None)</i>	
Scaling Object Detection by Transferring Classification Weights	6043
<i>Jason Kuen (Nanyang Technological University), Federico Perazzi (Adobe Research), Zhe Lin (Adobe Research), Jianming Zhang (Adobe Research), and Yap-Peng Tan (Nanyang Technological University. Singapore)</i>	
Scale-Aware Trident Networks for Object Detection	6053
<i>Yanghao Li (Peking University), Yuntao Chen (CASIA), Naiyan Wang (TuSimple), and Zhao-Xiang Zhang (Chinese Academy of Sciences. China)</i>	
Object-Aware Instance Labeling for Weakly Supervised Object Detection	6063
<i>Satoshi Kosugi (The University of Tokyo), Toshihiko Yamasaki (The University of Tokyo), and Kiyoharu Aizawa (The University of Tokyo)</i>	
Generative Modeling for Small-Data Object Detection	6072
<i>Lanlan Liu (University of Michigan), Michael Muelly (Google), Jia Deng (Princeton University), Tomas Pfister (Google), and Li-Jia Li (Stanford)</i>	

Transductive Learning for Zero-Shot Object Detection	6081
<i>Shafin Rahman (Australian National University), Salman Khan (Australian National University. ANU), and Nick Barnes (CSIRO(Data61)</i>	
Self-Training and Adversarial Background Regularization for Unsupervised Domain Adaptive One-Stage Object Detection	6091
<i>Seunghyeon Kim (KAIST), Jaehoon Choi (KAIST), Taekyung Kim (KAIST), and Changick Kim (KAIST)</i>	
Memory-Based Neighbourhood Embedding for Visual Recognition	6101
<i>Suichan Li (University of Science and Technology of China), Dapeng Chen (Sensetime Group Limited), Bin Liu (University of Science and Technology of China), Nenghai Yu (University of Science and Technology of China), and Rui Zhao (SenseTime Group Limited)</i>	
Self-Similarity Grouping: A Simple Unsupervised Cross Domain Adaptation Approach for Person Re-Identification	6111
<i>Yang Fu (UIUC), Yunchao Wei (UIUC), Guanshuo Wang (Shanghai Jiao Tong University), Yuqian Zhou (UIUC), Honghui Shi (IBM), Uiuic Uiuc, and Thomas Huang (UIUC)</i>	
Deep Reinforcement Active Learning for Human-in-the-Loop Person Re-Identification	6121
<i>Zimo Liu (Dalian University of Technology), Jingya Wang (University of Sydney), Shaogang Gong (Queen Mary University of London), Dacheng Tao (University of Sydney), and Huchuan Lu (Dalian University of Technology)</i>	
A Dual-Path Model With Adaptive Attention for Vehicle Re-Identification	6131
<i>Pirazh Khorramshahi (University of Maryland), Amit Kumar (University of Maryland), Neehar Peri (University of Maryland), Sai Saketh Rambhatla (University of Maryland), Jun-Cheng Chen (University of Maryland), and Rama Chellappa (University of Maryland)</i>	
Bayesian Loss for Crowd Count Estimation With Point Supervision	6141
<i>Ziheng Ma (Xi'an Jiaotong University), Xing Wei (Xi'an Jiaotong University), Xiaopeng Hong (Xi'an Jiaotong University), and Yihong Gong (Xi'an Jiaotong University)</i>	
Learning Spatial Awareness to Improve Crowd Counting	6151
<i>Zhi-Qi Cheng (Carnegie Mellon University), Jun-Xiu Li (Microsoft Research), Qi Dai (Microsoft Research), Xiao Wu (Southwest Jiaotong University), and Alexander Hauptmann (Carnegie Mellon University)</i>	

Oral 3.2B

Video & Action Underst&ing

GradNet: Gradient-Guided Network for Visual Object Tracking	6161
<i>Peixia Li (Dalian University of Technology), Boyu Chen (Dalian University of Technology), Wanli Ouyang (The University of Sydney), Dong Wang (Dalian University of Technology), Xiaoyun Yang (China Science IntelliCloud Technology Co.. Ltd), and Huchuan Lu (Dalian University of Technology)</i>	

FAMNet: Joint Learning of Feature, Affinity and Multi-Dimensional Assignment for Online Multiple Object Tracking	6171
<i>Peng Chu (Temple University) and Haibin Ling (Temple University)</i>	
Learning Discriminative Model Prediction for Tracking	6181
<i>Goutam Bhat (ETH Zurich), Martin Danelljan (ETH Zurich), Luc Van Gool (ETH Zurich), and Radu Timofte (ETH Zurich)</i>	
DynamoNet: Dynamic Action and Motion Network	6191
<i>Ali Diba (KU Leuven), Vivek Sharma (Karlsruhe Institute of Technology), Luc Van Gool (ETH Zurich), and Rainer Stiefelhagen (Karlsruhe Institute of Technology)</i>	
SlowFast Networks for Video Recognition	6201
<i>Christoph Feichtenhofer (Facebook AI Research), Haoqi Fan (Facebook AI Research), Jitendra Malik (University of California at Berkley), and Kaiming He (Facebook AI Research)</i>	
Generative Multi-View Human Action Recognition	6211
<i>Lichen Wang (Northeastern University), Zhengming Ding (Indiana University-Purdue University Indianapolis), Zhiqiang Tao (Northeastern University), Yunyu Liu (Northeastern University), and Yun Fu (Northeastern University)</i>	
Multi-Agent Reinforcement Learning Based Frame Sampling for Effective Untrimmed Video Recognition ...	6221
<i>Wenhao Wu (Shenzhen Institutes of Advanced Technology. Chinese Academy of Sciences), Dongliang He (Baidu), Xiao Tan (Baidu Research), Shifeng Chen (SIAT), and Shilei Wen (Baidu Research)</i>	
SCSampler: Sampling Salient Clips From Video for Efficient Action Recognition	6231
<i>Bruno Korbar (Facebook), Du Tran (Facebook Research), and Lorenzo Torresani (Facebook AI Research)</i>	
Weakly Supervised Energy-Based Learning for Action Segmentation	6242
<i>Jun Li (Oregon State University), Peng Lei (Amazon.com Services. Inc.), and Sinisa Todorovic (Oregon State U)</i>	
What Would You Expect? Anticipating Egocentric Actions With Rolling-Unrolling LSTMs and Modality Attention	6251
<i>Antonino Furnari (University of Catania) and Giovanni Farinella (University of Catania. Italy)</i>	
PIE: A Large-Scale Dataset and Models for Pedestrian Intention Estimation and Trajectory Prediction	6261
<i>Amir Rasouli (York University), Iuliia Kotseruba (York University), Toni Kunic (York University), and John Tsotsos (York University)</i>	
STGAT: Modeling Spatial-Temporal Interactions for Human Trajectory Prediction	6271
<i>Yingfan Huang (Institute of Computing Technology Chinese Academy of Sciences), Huikun Bi (Institute of Computing Technology Chinese Academy of Sciences), Zhaoxin Li (Institute of Computing Technology Chinese Academy of Sciences), Tianlu Mao (Institute of Computing Technology Chinese Academy of Sciences), and Zhaoqi Wang (Chinese Academy of Sciences)</i>	
Learning Motion in Feature Space: Locally-Consistent Deformable Convolution Networks for Fine-Grained Action Detection	6281
<i>Khoi-Nguyen Mac (UIUC), Dhiraj Joshi (IBM Research AI), Raymond Yeh (UIUC), Jinjun Xiong (IBM Thomas J. Watson Research Center), Rogerio Feris (IBM Research AI. MIT-IBM Watson AI Lab), and Minh Do (UIUC)</i>	

Dual Attention Matching for Audio-Visual Event Localization	6291
<i>Yu Wu (University of Technology Sydney), Linchao Zhu (University of Technology. Sydney), Yan Yan (Texas State University), and Yi Yang (UTS)</i>	
Uncertainty-Aware Audiovisual Activity Recognition Using Deep Bayesian Variational Inference	6300
<i>Mahesh Subedar (Intel), Ranganath Krishnan (Intel), Paulo Lopez Meyer (Intel), Omesh Tickoo (Intel), and Jonathan Huang (Intel)</i>	
Non-Local Recurrent Neural Memory for Supervised Sequence Modeling	6310
<i>Canmiao Fu (Peking University), Wenjie Pei (Tencent), Qiong Cao (Tencent), Chaopeng Zhang (Peking University), Yong Zhao (Peking University Shenzhen Graduate School), Xiaoyong Shen (Tencent), and Yu-Wing Tai (Tencent)</i>	
Temporal Attentive Alignment for Large-Scale Video Domain Adaptation	6320
<i>Min-Hung Chen (Georgia Institute of Technology), Zsolt Kira (Georgia Institute of Technology), Ghassan Alregib (Georgia Institute of Technology), Jaekwon Yoo (Sony Interactive Entertainment LLC), Ruxin Chen (Sie), and Jian Zheng (Binghamton University)</i>	
Action Assessment by Joint Relation Graphs	6330
<i>Jia-Hui Pan (Sun Yat-sen University), Jibin Gao (Sun Yat-sen University), and Wei-Shi Zheng (Sun Yat-sen University. China)</i>	
Unsupervised Procedure Learning via Joint Dynamic Summarization	6340
<i>Ehsan Elhamifar (Northeastern University) and Zwe Naing (Northeastern University)</i>	
ViSiL: Fine-Grained Spatio-Temporal Video Similarity Learning	6350
<i>Giorgos Kordopatis-Zilos (ITI-CERTH. Queen Mary University of London), Symeon Papadopoulos (Information Technologies Institute / Centre for Research & Technology - Hellas. GR), Ioannis Patras (Queen Mary University of London), and Yiannis Kompatsiaris (CERTH-ITI)</i>	

Poster 3.2

Deep Learning

Unsupervised Learning of Landmarks by Descriptor Vector Exchange	6360
<i>James Thewlis (University of Oxford), Samuel Albanie (University of Oxford), Hakan Bilen (University of Edinburgh), and Andrea Vedaldi (Oxford University)</i>	
Learning Compositional Representations for Few-Shot Recognition	6371
<i>Pavel Tokmakov (CMU), Yu-Xiong Wang (Carnegie Mellon University), and Martial Hebert (Carnegie Mellon University)</i>	
Spectral Regularization for Combating Mode Collapse in GANs	6381
<i>Kanglin Liu (Shenzhen University), Guoping Qiu (University of Nottingham), Wenming Tang (Shenzhen University), and Fei Zhou (Shenzhen University)</i>	
Scaling and Benchmarking Self-Supervised Visual Representation Learning	6390
<i>Priya Goyal (Facebook Inc.), Dhruv Mahajan (Facebook), Abhinav Gupta (CMU/FAIR), and Ishan Misra (Facebook AI Research)</i>	

Learning an Effective Equivariant 3D Descriptor Without Supervision	6400
<i>Riccardo Spezialetti (Universita' degli studi di Bologna), Samuele Salti (University of Bologna), and Luigi Di Stefano (University of Bologna)</i>	
KPConv: Flexible and Deformable Convolution for Point Clouds	6410
<i>Hugues Thomas (Mines ParisTech. PSL Research University), Charles R. Qi (Facebook AI Research), Jean-Emmanuel Deschaud (MINES ParisTech), Beatriz Marcotegui (MINES ParisTech), François Goulette (MINES ParisTech), and Leonidas Guibas (Stanford University)</i>	
Neural Inter-Frame Compression for Video Coding	6420
<i>Abdelaziz Djelouah (Disney Research), Joaquim Campos (Disney Research), Simone Schaub-Meyer (ETH Zurich), and Christopher Schroers (Disney Research)</i>	
Task2Vec: Task Embedding for Meta-Learning	6429
<i>Alessandro Achille (UCLA), Michael Lam (Amazon), Rahul Tewari (Amazon), Avinash Ravichandran (Amazon), Subhransu Maji (University of Massachusetts. Amherst), Charless Fowlkes (UC Irvine), Stefano Soatto (UCLA), and Pietro Perona (California Institute of Technology)</i>	
Deep Clustering by Gaussian Mixture Variational Autoencoders With Graph Embedding	6439
<i>Linxiao Yang (University of Electronic Science and Technology of China; Singapore University of Technology and Design), Ngai-Man Cheung (Singapore University of Technology and Design), Jiaying Li (SUTD), and Jun Fang (University of Electronic Science and Technology of China)</i>	
SoftTriple Loss: Deep Metric Learning Without Triplet Sampling	6449
<i>Qi Qian (Alibaba Group), Lei Shang (Alibaba Group), Baigui Sun (Alibaba Group), Juhua Hu (University of Washington), Tacoma Tacoma, Hao Li (Alibaba Group), and Rong Jin (alibaba group)</i>	
A Weakly Supervised Fine Label Classifier Enhanced by Coarse Supervision	6458
<i>Fariborz Taherkhani (West Virginia University), Hadi Kazemi (WVU), Ali Dabouei (West Virginia university), Jeremy Dawson (West Virginia University), and Nasser Nasrabadi (West Virginia University)</i>	
Gaussian Affinity for Max-Margin Class Imbalanced Learning	6468
<i>Munawar Hayat (University of Canberra), Salman Khan (Australian National University. ANU), Syed Waqas Zamir (IIAI), Jianbing Shen (Inception Institute of Artificial Intelligence), and Ling Shao (Inception Institute of Artificial Intelligence)</i>	
AttPool: Towards Hierarchical Feature Representation in Graph Convolutional Networks via Attention Mechanism	6479
<i>Jingjia Huang (Peking University), Zhangheng Li (Peking University), Nannan Li (Peking University Shenzhen Graduate School), Shan Liu (Tencent America), and Ge Li (SECE. Shenzhen Graduate School. Peking University)</i>	
Deep Metric Learning With Tuplet Margin Loss	6489
<i>Baosheng Yu (University of Sydney) and Dacheng Tao (University of Sydney)</i>	

Normalized Wasserstein for Mixture Distributions With Applications in Adversarial Learning and Domain Adaptation	6499
<i>Yogesh Balaji (University of Maryland. College Park), Rama Chellappa (University of Maryland), and Soheil Feizi (University of Maryland)</i>	
Fast and Practical Neural Architecture Search	6508
<i>Jiequan Cui (Chinese University of Hong Kong), Pengguang Chen (Chinese University of Hong Kong), Ruiyu Li (Tencent), Shu Liu (Tencent), Xiaoyong Shen (Tencent), and Jiaya Jia (Chinese University of Hong Kong)</i>	
Symmetric Graph Convolutional Autoencoder for Unsupervised Graph Representation Learning	6518
<i>Jiwoong Park (Seoul National University), Minsik Lee (Hanyang University), Hyung Jin Chang (University of Birmingham), Kyuewang Lee (Seoul National University), and Jin Young Choi (Seoul National University)</i>	
Deep Elastic Networks With Model Selection for Multi-Task Learning	6528
<i>Chanho Ahn (Department of ECE and ASRI. Seoul National University), Eunwoo Kim (University of Oxford), and Songhwai Oh (Seoul National University)</i>	
Metric Learning With HORDE: High-Order Regularizer for Deep Embeddings	6538
<i>Pierre Jacob (ETIS), David Picard (ETIS - CNRS), Aymeric Histace (ETIS), and Edouard Klein (C3N Gendarmerie Nationale)</i>	
Adversarial Learning With Margin-Based Triplet Embedding Regularization	6548
<i>Yaoyao Zhong (Beijing University of Posts and Telecommunications) and Weihong Deng (Beijing University of Posts and Telecommunications)</i>	

Recognition

Simultaneous Multi-View Instance Detection With Learned Geometric Soft-Constraints	6558
<i>Ahmed Nassar (IRISA Institute), Sébastien Lefèvre (Université de Bretagne Sud / IRISA), and Jan Dirk Wegner (ETH Zurich)</i>	
CenterNet: Keypoint Triplets for Object Detection	6568
<i>Kaiwen Duan (University of Chinese Academy of Sciences), Song Bai (University of Oxford), Lingxi Xie (Huawei Noah's Ark Lab), Honggang Qi (University of Chinese Academy of Sciences), Qingming Huang (University of Chinese Academy of Sciences), and Qi Tian (Huawei Noah's Ark Lab)</i>	
Online Hyper-Parameter Learning for Auto-Augmentation Strategy	6578
<i>Chen Lin (Sensetime), Minghao Guo (Sensetime), Chuming Li (sensetime), Xin Yuan (Tsinghua University), Wei Wu (SenseTime Group Limited), Junjie Yan (Sensetime Group Limited), Dahua Lin (The Chinese University of Hong Kong), and Wanli Ouyang (The University of Sydney)</i>	
DANet: Divergent Activation for Weakly Supervised Object Localization	6588
<i>Haolan Xue (University of Chinese Academy of Sciences), Chang Liu (University of Chinese Academy of Sciences), Fang Wan (University of Chinese Academy of Sciences), Jianbin Jiao (University of Chinese Academy of Sciences), Xiangyang Ji (Tsinghua University), and Qixiang Ye (University of Chinese Academy of Sciences. China)</i>	

Selective Sparse Sampling for Fine-Grained Image Recognition	6598
<i>Yao Ding (University of Chinese Academy of Sciences), Yanzhao Zhou (University of Chinese Academy of Sciences), Yi Zhu (University of Chinese Academy of Sciences), Qixiang Ye (University of Chinese Academy of Sciences. China), and Jianbin Jiao (University of Chinese Academy of Sciences)</i>	
Dynamic Anchor Feature Selection for Single-Shot Object Detection	6608
<i>Shuai Li (PolyU), Lingxiao Yang (The Hong Kong Polytechnic University), Jianqiang Huang (Alibaba Group), Xian-Sheng Hua (Alibaba Group), and Lei Zhang (Hong Kong Polytechnic University. Hong Kong. China)</i>	
Incremental Learning Using Conditional Adversarial Networks	6618
<i>Ye Xiang (Beijing Institute of Technology), Ying Fu (Beijing Institute of Technology), Pan Ji (NEC Laboratories America), and Hua Huang (Beijing Institute of Technology)</i>	
Bilateral Adversarial Training: Towards Fast Training of More Robust Models Against Adversarial Attacks	6628
<i>Jianyu Wang (Baidu Research USA) and Haichao Zhang (Baidu Research)</i>	
View Confusion Feature Learning for Person Re-Identification	6638
<i>Fangyi Liu (ChongQing University) and Lei Zhang (Chongqing University)</i>	
Auto-FPN: Automatic Network Architecture Adaptation for Object Detection Beyond Classification	6648
<i>Hang Xu (Huawei Noah's Ark Lab), Lewei Yao (Huawei Noah's Ark Lab), Zhenguo Li (Huawei Noah's Ark Lab), Xiaodan Liang (Sun Yat-sen University), and Wei Zhang (Noah's Ark Lab. Huawei Inc.)</i>	
PARN: Position-Aware Relation Networks for Few-Shot Learning	6658
<i>Ziyang Wu (South China University of Technology), Yuwei Li (South China University of Technology), Lihua Guo (South China University of Technology), and Kui Jia (South China University of Technology)</i>	
Multi-Adversarial Faster-RCNN for Unrestricted Object Detection	6667
<i>Zhenwei He (Chongqing University) and Lei Zhang (Chongqing University)</i>	
Object Guided External Memory Network for Video Object Detection	6677
<i>Hanming Deng (Shanghai Jiao Tong University), Yang Hua (Queen's University Belfast), Tao Song (Shanghai Jiao Tong University), Zongpu Zhang (Shanghai Jiao Tong University), Zhengui Xue (Shanghai Jiao Tong University), Ruhui Ma (Shanghai Jiao Tong University), Neil Robertson (Queen's University Belfast), and Haibing Guan (Shanghai Jiao Tong University)</i>	
An Empirical Study of Spatial Attention Mechanisms in Deep Networks	6687
<i>Xizhou Zhu (University of Science and Technology of China), Dazhi Cheng (Beijing Institute of Technology), Zheng Zhang (MSRA), Stephen Lin (Microsoft Research), and Jifeng Dai (Microsoft Research Asia)</i>	
Attribute Attention for Semantic Disambiguation in Zero-Shot Learning	6697
<i>Yang Liu (Zhejiang University), Jishun Guo (GAC R&D Center), Deng Cai (Zhejiang University), and Xiaofei He (Zhejiang University)</i>	
CIIDefence: Defeating Adversarial Attacks by Fusing Class-Specific Image Inpainting and Image Denoising	6707
<i>Puneet Gupta (IIT Indore) and Esa Rahtu (Tampere University of Technology)</i>	

ThunderNet: Towards Real-Time Generic Object Detection on Mobile Devices	6717
<i>Zheng Qin (National University of Defense Technology), Zeming Li (Megvii Face++ Inc), Zhaoning Zhang (National University of Defense Technology), Yiping Bao (Megvii Face++ Inc), Gang Yu (Megvii Inc), Yuxing Peng (National University of Defense Technology), and Jian Sun (Megvii Technology)</i>	
Dual Student: Breaking the Limits of the Teacher in Semi-Supervised Learning	6727
<i>Zhanghan Ke (City University of Hong Kong), Daoye Wang (SenseTime Research), Qiong Yan (SenseTime Group Limited), Jimmy Ren (SenseTime Research), and Rynson Lau (City University of Hong Kong)</i>	
MVP Matching: A Maximum-Value Perfect Matching for Mining Hard Samples, With Application to Person Re-Identification	6736
<i>Han Sun (Institute of Advanced Artificial Intelligence. IAAI), Zhiyuan Chen (Institute of Advanced Artificial Intelligence. IAAI), Shiyang Yan (Queen's University Belfast), and Lin Xu (Institute of Advanced Artificial Intelligence. IAAI)</i>	

Segmentation, Grouping, & Shape

Adaptive Context Network for Scene Parsing	6747
<i>Jun Fu (National Laboratory of Pattern Recognition. Institute of Automation. Chinese Academy of Sciences and University of Chinese Academy of Sciences), Jing Liu (National Lab of Pattern Recognition. Institute of Automation), Yuhang Wang (Institute of Automation. Chinese Academy of Sciences), Yong Li (Business Growth BU. JD.com), Yongjun Bao (JD.com), Jinhui Tang (Nanjing University of Science and Technology), and Hanqing Lu (NLPR. Institute of Automation. CAS)</i>	
Constructing Self-Motivated Pyramid Curriculums for Cross-Domain Semantic Segmentation: A Non-Adversarial Approach	6757
<i>Qing Lian (University of Electronic Science and Technology of China), Lixin Duan (University of Electronic Science and Technology of China), Fengmao Lv (University of Electronic Science and Technology of China), and Boqing Gong (Google / ICSI Berkeley)</i>	
SparseMask: Differentiable Connectivity Learning for Dense Image Prediction	6767
<i>Huikai Wu (CASIA), Junge Zhang (CASIA), and Kaiqi Huang (Institute of Automation. Chinese Academy of Sciences)</i>	
Significance-Aware Information Bottleneck for Domain Adaptive Semantic Segmentation	6777
<i>Yawei Luo (University of Technology Sydney), Ping Liu (UTS), Tao Guan (Huazhong University of Science and Technology), Junqing Yu (Huazhong University of Science & Technology), and Yi Yang (UTS)</i>	
Relational Attention Network for Crowd Counting	6787
<i>Anran Zhang (Beihang University), Jiayi Shen (Beihang University), Zehao Xiao (Beihang University), Fan Zhu (Inception Institute of Artificial Intelligence), Xiantong Zhen (Inception Institute of Artificial Intelligence), Xianbin Cao (Beihang University. China), and Ling Shao (Inception Institute of Artificial Intelligence)</i>	

ACFNet: Attentional Class Feature Network for Semantic Segmentation	6797
<i>Fan Zhang (Institute of Software. Chinese Academy of Sciences), Yanqin Chen (Baidu Inc.), Zhihang Li (Institute of Automation. Chinese Academy of Science), Zhibin Hong (Baidu Inc.), Jingtuo Liu (baidu), Feifei Ma (Institute of Software. Chinese Academy of Sciences), Junyu Han (Baidu Inc.), and Errui Ding (Baidu Inc.)</i>	
Frame-to-Frame Aggregation of Active Regions in Web Videos for Weakly Supervised Semantic Segmentation	6807
<i>Jungbeom Lee (Seoul National University), Eunji Kim (Seoul National University), Sungmin Lee (Seoul National University), Jangho Lee (Seoul National University), and Sungroh Yoon (Seoul National University)</i>	
Boundary-Aware Feature Propagation for Scene Segmentation	6818
<i>Henghui Ding (Nanyang Technological University), Xudong Jiang (Nanyang Technological University), Ai Qun Liu (Nanyang Technological University), Nadia Magnenat Thalmann (Nanyang Technological University), and Gang Wang (Alibaba Group)</i>	
Self-Ensembling With GAN-Based Data Augmentation for Domain Adaptation in Semantic Segmentation	6829
<i>Jaehoon Choi (KAIST), Taekyung Kim (KAIST), and Changick Kim (KAIST)</i>	

3D From Single View & RGBD

Explaining the Ambiguity of Object Detection and 6D Pose From Visual Data	6840
<i>Fabian Manhardt (TU Munich), Diego Martín Arroyo (TU Munich), Christian Rupprecht (University of Oxford), Benjamin Busam (Huawei Noah's Ark Lab), Tolga Birdal (TU Munich), Nassir Navab (TU Munich. Germany), and Federico Tombari (TUM. Google)</i>	
Accurate Monocular 3D Object Detection via Color-Embedded 3D Reconstruction for Autonomous Driving .	6850
<i>Xinzhu Ma (Dalian University of Technology), Zhihui Wang (Dalian University of Technology), Haojie Li (Dalian University of Technology), Pengbo Zhang (Dalian University of Technology), Wanli Ouyang (The University of Sydney), and Xin Fan (Dalian University of Technology)</i>	
MonoLoco: Monocular 3D Pedestrian Localization and Uncertainty Estimation	6860
<i>Lorenzo Bertoni (EPFL), Sven Kreiss (EPFL), and Alexandre Alahi (EPFL)</i>	
Unsupervised High-Resolution Depth Learning From Videos With Dual Networks	6871
<i>Junsheng Zhou (Tsinghua University), Yuwang Wang (Microsoft Research), Kaihuai Qin (Tsinghua University), and Wenjun Zeng (Microsoft Research)</i>	

Face & Body

Bayesian Graph Convolution LSTM for Skeleton Based Action Recognition	6881
<i>Rui Zhao (Amazon), Kang Wang (RPI), Hui Su (IBM), and Qiang Ji (Rensselaer Polytechnic Institute)</i>	

DeCaFA: Deep Convolutional Cascade for Face Alignment in the Wild	6892
<i>Arnaud Dapogny (Pierre and Marie Curie University. UPMC), Matthieu Cord (Sorbonne University), and Kevin Bailly (Datakalab)</i>	
Probabilistic Face Embeddings	6901
<i>Yichun Shi (Michigan State University) and Anil Jain (Michigan State University)</i>	
Gaze360: Physically Unconstrained Gaze Estimation in the Wild	6911
<i>Petr Kellnhofer (MIT), Adria Recasens (Massachusetts Institute of Technology), Simon Stent (Toyota Research Institute), Wojciech Matusik (MIT), and Antonio Torralba (MIT)</i>	
Unsupervised Person Re-Identification by Camera-Aware Similarity Consistency Learning	6921
<i>Ancong Wu (Sun Yat-sen University), Wei-Shi Zheng (Sun Yat-sen University. China), and Jian-Huang Lai (Sun Yat-sen University)</i>	
Photo-Realistic Monocular Gaze Redirection Using Generative Adversarial Networks	6931
<i>Zhe He (ETH Zürich and University of Zürich), Adrian Spurr (ETH), Xucong Zhang (ETH Zurich), and Otmar Hilliges (ETH Zurich)</i>	
Dynamic Kernel Distillation for Efficient Pose Estimation in Videos	6941
<i>Xuecheng Nie (NUS), Yuncheng Li (Snap), Linjie Luo (Snap Inc), Ning Zhang (UC Berkeley), and Jiashi Feng (NUS)</i>	
Single-Stage Multi-Person Pose Machines	6950
<i>Xuecheng Nie (NUS), Jiashi Feng (NUS), Jianfeng Zhang (National University of Singapore), and Shuicheng Yan (Qihoo/360)</i>	
SO-HandNet: Self-Organizing Network for 3D Hand Pose Estimation With Semi-Supervised Learning	6960
<i>Yujin Chen (Wuhan University), Zhigang Tu (Wuhan University), Liuhaog Ge (Nanyang Technological University), Dejun Zhang (China University of Geosciences. Wuhan), Ruizhi Chen (Wuhan University), and Junsong Yuan (State University of New York at Buffalo. USA)</i>	
Adaptive Wing Loss for Robust Face Alignment via Heatmap Regression	6970
<i>Xinyao Wang (Oregon State University), Liefeng Bo (JD Finance), and Li Fuxin (Oregon State University)</i>	
Single-Network Whole-Body Pose Estimation	6981
<i>Gines Hidalgo Martinez (Carnegie Mellon University), Yaadhav Raaj (CMU), Haroon Idrees (RetailNext), Donglai Xiang (Carnegie Mellon University), Hanbyul Joo (Facebook AI Research), Tomas Simon (Carnegie Mellon University), and Yaser Sheikh (CMU)</i>	
Face Alignment With Kernel Density Deep Neural Network	6991
<i>Lisha Chen (RENSELAER POLYTECHNIC INST), Hui Su (IBM), and Qiang Ji (Rensselaer Polytechnic Institute)</i>	

Action & Video

Spatiotemporal Feature Residual Propagation for Action Prediction	7002
<i>He Zhao (York University) and Rick Wildes (York University)</i>	

Identity From Here, Pose From There: Self-Supervised Disentanglement and Generation of Objects Using Unlabeled Videos	7012
<i>Fanyi Xiao (University of California Davis), Haotian Liu (Zhejiang University), and Yong Jae Lee (University of California. Davis)</i>	
Relation Distillation Networks for Video Object Detection	7022
<i>Jiajun Deng (University of Science and Technology of China), Yingwei Pan (JD AI Research), Ting Yao (JD AI Research), Wengang Zhou (University of Science and Technology of China), Houqiang Li (University of Science and Technology of China), and Tao Mei (AI Research of JD.com)</i>	
Video Compression With Rate-Distortion Autoencoders	7032
<i>Amirhossein Habibian (Qualcomm AI Research), Ties Van Rozendaal (Qualcomm AI Research), Jakub Tomczak (Qualcomm AI Research), and Taco Cohen (Qualcomm AI Research)</i>	
Non-Local ConvLSTM for Video Compression Artifact Reduction	7042
<i>Yi Xu (Fudan University), Longwen Gao (Bilibili), Kai Tian (Fudan University), Shuigeng Zhou (Fudan University), and Huyang Sun (Bilibili)</i>	
Self-Supervised Moving Vehicle Tracking With Stereo Sound	7052
<i>Chuang Gan (MIT-Watson AI Lab), Hang Zhao (MIT), Peihao Chen (SCUT), David Cox (MIT-IBM Watson AI Lab), and Antonio Torralba (MIT)</i>	
Self-Supervised Learning With Geometric Constraints in Monocular Video: Connecting Flow, Depth, and Camera	7062
<i>Yuhua Chen (ETH Zurich), Cordelia Schmid (Google), and Cristian Sminchisescu (Lund University)</i>	
Learning Temporal Action Proposals With Fewer Labels	7072
<i>Jingwei Ji (Stanford University), Kaidi Cao (Stanford University), and Juan Carlos Niebles (Stanford University)</i>	
TSM: Temporal Shift Module for Efficient Video Understanding	7082
<i>Ji Lin (MIT), Chuang Gan (MIT-Watson AI Lab), and Song Han (MIT)</i>	
Graph Convolutional Networks for Temporal Action Localization	7093
<i>Runhao Zeng (South China University of Technology), Wenbing Huang (Tencent AI Lab), Chuang Gan (MIT-Watson AI Lab), Mingkui Tan (South China University of Technology), Yu Rong (Tencent AI Lab), Peilin Zhao (Tencent AI Lab), and Junzhou Huang (University of Texas at Arlington)</i>	
Fast Object Detection in Compressed Video	7103
<i>Shiyao Wang (Tsinghua University), Alibaba Group, Hongchao Lu (Tsinghua University), and Zhidong Deng (Tsinghua University)</i>	

Motion & Tracking

Predicting 3D Human Dynamics From Video	7113
<i>Jason Zhang (University of California. Berkeley), Panna Felsen (University of California Berkeley), Angjoo Kanazawa (University of California Berkeley), and Jitendra Malik (University of California at Berkley)</i>	

Imitation Learning for Human Pose Prediction	7123
<i>Borui Wang (Stanford University), Ehsan Adeli (Stanford University), Hsu-Kuang Chiu (Stanford University), De-An Huang (Stanford University), and Juan Carlos Niebles (Stanford University)</i>	
Human Motion Prediction via Spatio-Temporal Inpainting	7133
<i>Alejandro Hernandez (Institut de Robotica i Informatica Industrial), Jürgen Gall (University of Bonn), and Francesc Moreno (IRI)</i>	
Structured Prediction Helps 3D Human Motion Modelling	7143
<i>Emre Aksan (ETH Zurich), Manuel Kaufmann (ETH Zurich), and Otmar Hilliges (ETH Zurich)</i>	

Computational Photography & Graphics

Learning Shape Templates With Structured Implicit Functions	7153
<i>Kyle Genova (Princeton University), Forrester Cole (Google Research), Daniel Vlasic (Google), Aaron Sarna (Google), William Freeman (Google), and Thomas Funkhouser (Princeton University)</i>	
CompenNet++: End-to-End Full Projector Compensation	7164
<i>Bingyao Huang (Temple University) and Haibin Ling (Temple University)</i>	
Deep Parametric Indoor Lighting Estimation	7174
<i>Marc-André Gardner (Université Laval), Yannick Hold-Geoffroy (Adobe Research), Kalyan Sunkavalli (Adobe Research), Christian Gagné (Université Laval), and Jean-Francois Lalonde (Université Laval)</i>	
FSGAN: Subject Agnostic Face Swapping and Reenactment	7183
<i>Yuval Nirkin (Bar-Ilan University), Yosi Keller (Bar Ilan University), and Tal Hassner (Open University of Israel)</i>	
Deep Single-Image Portrait Relighting	7193
<i>Hao Zhou (UMD), Sunil Hadap (Adobe), Kalyan Sunkavalli (Adobe Research), and David Jacobs (University of Maryland. USA)</i>	
PU-GAN: A Point Cloud Upsampling Adversarial Network	7202
<i>Ruihui Li (The Chinese University of Hong Kong), Xianzhi Li (The Chinese University of Hong Kong), Chi-Wing Fu (The Chinese University of Hong Kong), Daniel Cohen-Or (Tel Aviv University), and Pheng-Ann Heng (The Chinese University of Hong Kong)</i>	
Neural 3D Morphable Models: Spiral Convolutional Networks for 3D Shape Representation Learning and Generation	7212
<i>Giorgos Bouritsas (Imperial College London), Sergiy Bokhnyak (Università della Svizzera italiana), Stylianos Ploumpis (Imperial College London), Stefanos Zafeiriou (Imperial College London), and Michael Bronstein (Imperial College)</i>	

Low-Level & Optimization

Joint Learning of Saliency Detection and Weakly Supervised Semantic Segmentation	7222
<i>Zeng Yu (Dalian University of Technology), Yunzhi Zhuge (Dalian University of Technology), Huchuan Lu (Dalian University of Technology), and Lihe Zhang (Dalian University of Technology)</i>	

Towards High-Resolution Salient Object Detection	7233
<i>Yi Zeng (Dalian University of Technology), Pingping Zhang (Dalian University of Technology), Zhe Lin (Adobe Research), Jianming Zhang (Adobe Research), and Huchuan Lu (Dalian University of Technology)</i>	
Event-Based Motion Segmentation by Motion Compensation	7243
<i>Timo Stoffregen (Monash University), Guillermo Gallego (University of Zurich and ETH Zurich), Tom Drummond (Monash University), Lindsay Kleeman (Monash University), and Davide Scaramuzza (University of Zurich & ETH Zurich. Switzerland)</i>	
Depth-Induced Multi-Scale Recurrent Attention Network for Saliency Detection	7253
<i>Yongri Piao (Dalian University of Technology), Wei Ji (Dalian University of Technology), Jingjing Li (Dalian University of Technology), Miao Zhang (Dalian University of Technology), and Huchuan Lu (Dalian University of Technology)</i>	
Stacked Cross Refinement Network for Edge-Aware Salient Object Detection	7263
<i>Zhe Wu (University of Chinese Academy of Sciences), Li Su (University of Chinese Academy of Sciences), and Qingming Huang (University of Chinese Academy of Sciences)</i>	
Motion Guided Attention for Video Salient Object Detection	7273
<i>Haofeng Li (The University of Hong Kong), Guanqi Chen (Sun Yat-sen University), Guanbin Li (Sun Yat-sen University), and Yizhou Yu (Deepwise AI Lab)</i>	
Semi-Supervised Video Salient Object Detection Using Pseudo-Labels	7283
<i>Pengxiang Yan (Sun Yat-sen University), Guanbin Li (Sun Yat-sen University), Yuan Xie (Sun Yat-sen University), Zhen Li (Shenzhen Research Institute of Big Data. the Chinese University of Hong Kong. Shenzhen), Chuan Wang (Face++. Megvii), Tianshui Chen (Sun Yat-Sen University), and Liang Lin (Sun Yat-sen University)</i>	
Joint Learning of Semantic Alignment and Object Landmark Detection	7293
<i>Sangryul Jeon (Yonsei university), Dongbo Min (Ewha Womans University), Seungryong Kim (École Polytechnique Fédérale de Lausanne. EPFL), and Kwanghoon Sohn (Yonsei Univ.)</i>	
RainFlow: Optical Flow Under Rain Streaks and Rain Veiling Effect	7303
<i>Ruoteng Li (National University of Singapore), Robby Tan (Yale-NUS College. Singapore), Loong Fah Cheong (NUS), Angelica Aviles-Rivero (University of Cambridge), Qingnan Fan (Shandong University), and Carola Schoenlieb (Cambridge University)</i>	
GridDehazeNet: Attention-Based Multi-Scale Network for Image Dehazing	7313
<i>Xiaohong Liu (McMaster University), Yongrui Ma (McMaster University), Zhihao Shi (McMaster University), and Jun Chen (McMaster University)</i>	
Learning to See Moving Objects in the Dark	7323
<i>Haiyang Jiang (National Institute of Informatics) and Yinqiang Zheng (National Institute of Informatics)</i>	

Scene Understanding

SegSort: Segmentation by Discriminative Sorting of Segments	7333
<i>Jyh-Jing Hwang (International Computer Science Institute), Stella Yu (UC Berkeley / ICSI), Jianbo Shi (University of Pennsylvania), Maxwell Collins (Google Inc.), Tien-Ju Yang (Massachusetts Institute of Technology), Xiao Zhang (Google), and Liang-Chieh Chen (Google Inc.)</i>	
What Synthesis Is Missing: Depth Adaptation Integrated With Weak Supervision for Indoor Scene Parsing	7344
<i>Keng-Chi Liu (National Taiwan University), Yi-Ting Shen (National Taiwan University), Jan Klopp (National Taiwan University), and Liang-Gee Chen (National Taiwan University)</i>	
AdaptIS: Adaptive Instance Selection Network	7354
<i>Konstantin Sofiiuk (Samsung AI Center Moscow), Konstantin Sofiyuk (Samsung AI Center Moscow), Olga Barinova (Samsung AI Center Moscow), Anton Konushin (Samsung AI Center Moscow), and Olga Barinova (Samsung AI Center Moscow)</i>	
DADA: Depth-Aware Domain Adaptation in Semantic Segmentation	7363
<i>Tuan-Hung Vu (Valeo.ai), Himalaya Jain (Valeo.ai), Maxime Bucher (Valeo.ai), Matthieu Cord (Sorbonne University), and Patrick Pérez (Valeo.ai)</i>	
Guided Curriculum Model Adaptation and Uncertainty-Aware Evaluation for Semantic Nighttime Image Segmentation	7373
<i>Christos Sakaridis (ETH Zurich), Dengxin Dai (ETH Zurich), and Luc Van Gool (ETH Zurich)</i>	
SceneGraphNet: Neural Message Passing for 3D Indoor Scene Augmentation	7383
<i>Yang Zhou (UMass Amherst), Zachary While (UMass Amherst), and Evangelos Kalogerakis (UMass)</i>	
SkyScapes - Fine-Grained Semantic Understanding of Aerial Scenes	7392
<i>Seyed Majid Azimi (German Aerospace Center. DLR), Corentin Henry (German Aerospace Center. DLR), Lars Sommer (Fraunhofer IOSB. Karlsruhe. Germany), Arne Schumann (Fraunhofer IOSB), and Eleonora Vig (DLR)</i>	

Language & Reasoning

Transferable Representation Learning in Vision-and-Language Navigation	7403
<i>Haoshuo Huang (Google), Vihan Jain (Google Research), Harsh Mehta (Google Inc), Alexander Ku (Google), Gabriel Magalhaes (Google AI), Jason Baldridge (Google), and Eugene Ie (Google Research)</i>	
Towards Unsupervised Image Captioning With Shared Multimodal Embeddings	7413
<i>Iro Laina (Technical University of Munich), Christian Rupprecht (University of Oxford), and Nassir Navab (TU Munich. Germany)</i>	
ViCo: Word Embeddings From Visual Co-Occurrences	7424
<i>Tanmay Gupta (UIUC), Alexander Schwing (UIUC), and Derek Hoiem (University of Illinois at Urbana-Champaign)</i>	

Seq-SG2SL: Inferring Semantic Layout From Scene Graph Through Sequence to Sequence Learning	7434
<i>Boren Li (Alibaba AI Labs), Boyu Zhuang (Alibaba AI Labs), Mingyang Li (Alibaba A.I. Labs), and Jian Gu (Alibaba AI Labs)</i>	
U-CAM: Visual Explanation Using Uncertainty Based Class Activation Maps	7443
<i>Badri Patro (IIT Kanpur), Mayank Lunayach (IIT Kanpur), Shivansh Patel (IIT Kanpur), and Vinay Namboodiri (IIT Kanpur)</i>	
See-Through-Text Grouping for Referring Image Segmentation	7453
<i>Ding-Jie Chen (Academia Sinica), Songhao Jia (National Tsing Hua University), Yi-Chen Lo (National Tsing Hua University), Hwann-Tzong Chen (National Tsing Hua University), and Tyng-Luh Liu (Academia Sinica)</i>	
VideoBERT: A Joint Model for Video and Language Representation Learning	7463
<i>Chen Sun (Google), Austin Myers (Google), Carl Vondrick (Columbia University), Kevin Murphy (Google), and Cordelia Schmid (Google)</i>	
Language Features Matter: Effective Language Representations for Vision-Language Tasks	7473
<i>Andrea Burns (Boston University), Reuben Tan (Boston University), Kate Saenko (Boston University), Stan Sclaroff (Boston University), and Bryan Plummer (Boston University)</i>	

3D From Multiview & Sensors

Semantic Stereo Matching With Pyramid Cost Volumes	7483
<i>Zhenyao Wu (University of South Carolina), Xinyi Wu (University of South Carolina), Xiaoping Zhang (Wuhan University), Song Wang (University of South Carolina), and Lili Ju (University of South Carolina)</i>	
Spatial Correspondence With Generative Adversarial Network: Learning Depth From Monocular Videos	7493
<i>Zhenyao Wu (University of South Carolina), Xinyi Wu (University of South Carolina), Xiaoping Zhang (Wuhan University), Song Wang (University of South Carolina), and Lili Ju (University of South Carolina)</i>	
Learning Relationships for Multi-View 3D Object Recognition	7504
<i>Ze Yang (Peking University) and Liwei Wang (Peking University)</i>	
View N-Gram Network for 3D Object Retrieval	7514
<i>Xinwei He (Huazhong University of Science and Technology), Teng Teng Huang (Huazhong University of Science and Technology), Song Bai (University of Oxford), and Xiang Bai (Huazhong University of Science and Technology)</i>	
Expert Sample Consensus Applied to Camera Re-Localization	7524
<i>Eric Brachmann (Heidelberg University) and Carsten Rother (University of Heidelberg)</i>	
Semantic Part Detection via Matching: Learning to Generalize to Novel Viewpoints From Limited Training Data	7534
<i>Yutong Bai (Northwestern Polytechnical University), Qing Liu (Johns Hopkins University), Lingxi Xie (Huawei Noah's Ark Lab), Yan Zheng (Beihang University), Weichao Qiu (Johns Hopkins University), and Alan Yuille (Johns Hopkins University)</i>	

Dynamic Points Agglomeration for Hierarchical Point Sets Learning	7545
<i>Jinxian Liu (Shanghai Jiao Tong University), Bingbing Ni (Shanghai Jiao Tong University), Caiyuan Li (Shanghai Jiao Tong University), Jiancheng Yang (Shanghai Jiao Tong University), and Qi Tian (Huawei Noah's Ark Lab)</i>	

Image & Video Synthesis

Attributing Fake Images to GANs: Learning and Analyzing GAN Fingerprints	7555
<i>Ning Yu (Max Planck Institute for Informatics), Larry Davis (University of Maryland), and Mario Fritz (CISPA Helmholtz Center for Information Security)</i>	
Dual Adversarial Inference for Text-to-Image Synthesis	7566
<i>Qicheng Lao (Concordia University), Mohammad Havaei (Imagia), Ahmad Pesaranghader (Dalhousie University), Francis Dutil (Imagia), Lisa Di Jorio (Imagia Cybernetics Inc), and Thomas Fevens (Concordia University)</i>	
View-LSTM: Novel-View Video Synthesis Through View Decomposition	7576
<i>Mohamed Lakhel (Queen Mary University of London), Oswald Lanz (Fondazione Bruno Kessler. Italy), and Andrea Cavallaro (Queen Mary University of London. UK)</i>	
HoloGAN: Unsupervised Learning of 3D Representations From Natural Images	7587
<i>Thu Nguyen-Phuoc (University of Bath), Chuan Li (Lambda Labs), Lucas Theis (Twitter), Christian Richardt (University of Bath), and Yongliang Yang (University of Bath)</i>	
Unpaired Image-to-Speech Synthesis With Multimodal Information Bottleneck	7597
<i>Shuang Ma (SUNY Buffalo), Daniel Mcduff (Microsoft Research), and Yale Song (Microsoft)</i>	
Improved Conditional VRNNs for Video Prediction	7607
<i>Lluís Castrejón (Mila - Université de Montréal), Nicolas Ballas (Facebook FAIR), and Aaron Courville (MILA. Université de Montréal)</i>	
Visualizing the Invisible: Occluded Vehicle Segmentation and Recovery	7617
<i>Xiaosheng Yan (Fuzhou University), Yuanlong Yu (Fuzhou University), Feigege Wang (Fuzhou University), Wenxi Liu (Fuzhou University), Shengfeng He (South China University of Technology), and Jia Pan (City University of Hong Kong)</i>	

Oral 4.1A

Single-View 3D Modeling, Pose Estimation

Learning Single Camera Depth Estimation Using Dual-Pixels	7627
<i>Rahul Garg (Google), Neal Wadhwa (Google), Sameer Ansari (Google), and Jonathan Barron (-)</i>	

Domain-Adaptive Single-View 3D Reconstruction	7637
<i>Pedro O. Pinheiro (Element AI), Negar Rostamzadeh (Element AI), and Sungjin Ahn (Rutgers University)</i>	
Transformable Bottleneck Networks	7647
<i>Kyle Olszewski (University of Southern California), Sergey Tulyakov (Snap Inc), Oliver Woodford (Snap Inc.), Hao Li (Pinscreen/University of Southern California/USC ICT), and Linjie Luo (Snap Inc)</i>	
RIO: 3D Object Instance Re-Localization in Changing Indoor Environments	7657
<i>Johanna Wald (Technical University of Munich), Armen Avetisyan (Technical University of Munich), Nassir Navab (TU Munich. Germany), Federico Tombari (TUM. Google), and Matthias Niessner (Technical University of Munich)</i>	
Pix2Pose: Pixel-Wise Coordinate Regression of Objects for 6D Pose Estimation	7667
<i>Kiru Park (TU Wien), Timothy Patten (TU Wien), and Markus Vincze (TU Wien)</i>	
CDPN: Coordinates-Based Disentangled Pose Network for Real-Time RGB-Based 6-DoF Object Pose Estimation	7677
<i>Zhigang Li (Tsinghua University), Gu Wang (Tsinghua University), and Xiangyang Ji (Tsinghua University)</i>	
C3DPPO: Canonical 3D Pose Networks for Non-Rigid Structure From Motion	7687
<i>David Novotny (Facebook AI Research), Nikhila Ravi (Facebook AI Research), Ben Graham (Facebook Research), Natalia Neverova (Facebook AI Research), and Andrea Vedaldi (University of Oxford / Facebook AI Research)</i>	
Learning to Reconstruct 3D Manhattan Wireframes From a Single Image	7697
<i>Yichao Zhou (UC Berkeley), Haozhi Qi (UC Berkeley), Yuexiang Zhai (UC Berkeley), Qi Sun (Adobe Research), Zhili Chen (ByteDance AI Lab), Li-Yi Wei (Adobe Research), and Yi Ma (UC Berkeley)</i>	
Soft Rasterizer: A Differentiable Renderer for Image-Based 3D Reasoning	7707
<i>Shichen Liu (University of Southern California), Weikai Chen (USC Institute for Creative Technology), Tianye Li (University of Southern California), and Hao Li (Pinscreen/University of Southern California/USC ICT)</i>	
Learnable Triangulation of Human Pose	7717
<i>Karim Iskakov (Samsung), Egor Burkov (Samsung AI Center Moscow), Victor Lempitsky (Samsung), and Yury Malkov (Samsung AI Center Moscow)</i>	
xR-EgoPose: Egocentric 3D Human Pose From an HMD Camera	7727
<i>Denis Tome (UCL), Patrick Peluse (Facebook), Lourdes Agapito (University College London), and Hernan Badino (Facebook)</i>	
DeepHuman: 3D Human Reconstruction From a Single Image	7738
<i>Zerong Zheng (Tsinghua University), Tao Yu (Beihang University), Yixuan Wei (Tsinghua University), Qionghai Dai (Tsinghua University), and Yebin Liu (Tsinghua University)</i>	
A Neural Network for Detailed Human Depth Estimation From a Single Image	7749
<i>Sicong Tang (Simon Fraser University), Feitong Tan (Simon Fraser University), Kelvin Cheng (Simon Fraser University), Zhaoyang Li (Simon Fraser University), Siyu Zhu (Alibaba A.I. Labs), and Ping Tan (Simon Fraser University)</i>	

DenseRaC: Joint 3D Pose and Shape Estimation by Dense Render-and-Compare	7759
<i>Yuanlu Xu (University of California. Los Angeles), Song-Chun Zhu (UCLA), and Tony Tung (Facebook Reality Labs)</i>	
Not All Parts Are Created Equal: 3D Pose Estimation by Modeling Bi-Directional Dependencies of Body Parts	7770
<i>Jue Wang (University of Technology Sydney), Shaoli Huang (University of Sydney), Xinchao Wang (Stevens Institute of Technology), and Dacheng Tao (University of Sydney)</i>	

Oral 4.1B

Computational Photography

Extreme View Synthesis	7780
<i>Inchang Choi (KAIST), Orazio Gallo (NVIDIA Research), Alejandro Troccoli (NVIDIA), Min H. Kim (KAIST), and Jan Kautz (NVIDIA)</i>	
View Independent Generative Adversarial Network for Novel View Synthesis	7790
<i>Xiaogang Xu (The Chinese University of Hong Kong), Yingcong Chen (Chinese University of Hong Kong), and Jiaya Jia (Chinese University of Hong Kong)</i>	
Cascaded Context Pyramid for Full-Resolution 3D Semantic Scene Completion	7800
<i>Pingping Zhang (Dalian University of Technology), Wei Liu (University of Adelaide), Yinjie Lei (Sichuan University), Huchuan Lu (Dalian University of Technology), and Xiaoyun Yang (China Science IntelliCloud Technology Co.. Ltd)</i>	
View-Consistent 4D Light Field Superpixel Segmentation	7810
<i>Numair Khan (Brown University), Qian Zhang (Brown University), Lucas Kasser (Brown University), Henry Stone (Brown University), Min H. Kim (KAIST), and James Tompkin (Brown University)</i>	
GLoSH: Global-Local Spherical Harmonics for Intrinsic Image Decomposition	7819
<i>Hao Zhou (UMD), Xiang Yu (NEC Labs), and David Jacobs (University of Maryland. USA)</i>	
Surface Normals and Shape From Water	7829
<i>Satoshi Murai (Kyoto University), Meng-Yu Kuo (Kyoto University), Ryo Kawahara (Kyoto University), Shohei Nobuhara (Kyoto University), and Ko Nishino (Kyoto University)</i>	
Restoration of Non-Rigidly Distorted Underwater Images Using a Combination of Compressive Sensing and Local Polynomial Image Representations	7838
<i>Jerin Geo James (Indian Institute of Technology Bombay), Pranay Agrawal (IIT Bombay), and Ajit Rajwade (IIT Bombay)</i>	

Learning Perspective Undistortion of Portraits	7848
<i>Yajie Zhao (Institution for Creative Technologies. University of Southern California), Zeng Huang (University of Southern California), Tianye Li (University of Southern California), Weikai Chen (USC Institute for Creative Technology), Chloe Legendre (USC Institute for Creative Technology), Xinglei Ren (Institution for Creative Technologies. University of Southern California), Ari Shapiro (USC Institute for Creative Technologies), and Hao Li (Pinscreen/University of Southern California/USC ICT)</i>	
Towards Photorealistic Reconstruction of Highly Multiplexed Lensless Images	7859
<i>Salman Siddique Khan (IIT Madras), Adarsh V R (IIT Madras), Vivek Boominathan (Rice University), Jasper Tan (Rice University), Ashok Veeraraghavan (Rice University), and Kaushik Mitra (IIT Madras)</i>	
Unconstrained Motion Deblurring for Dual-Lens Cameras	7869
<i>Mahesh Mohan M R (Indian Institute of Technology Madras), Sharath Girish (IIT Madras), and Rajagopalan Ambasamudram (Indian Institute of Technology Madras)</i>	
Stochastic Exposure Coding for Handling Multi-ToF-Camera Interference	7879
<i>Jongho Lee (University of Wisconsin-Madison) and Mohit Gupta (University of Wisconsin-Madison. USA)</i>	
Convolutional Approximations to the General Non-Line-of-Sight Imaging Operator	7888
<i>Byeongjoo Ahn (Carnegie Mellon University), Akshat Dave (Rice University), Ashok Veeraraghavan (Rice University), Ioannis Gkioulekas (Carnegie Mellon University), and Aswin Sankaranarayanan (Carnegie Mellon University)</i>	
Agile Depth Sensing Using Triangulation Light Curtains	7899
<i>Joseph Bartels (Carnegie Mellon University), Jian Wang (Snap), William Whittaker (Carnegie Mellon University), and Srinivasa Narasimhan (Carnegie Mellon University)</i>	
Asynchronous Single-Photon 3D Imaging	7908
<i>Anant Gupta (University of Wisconsin Madison), Atul Ingle (University of Wisconsin-Madison), and Mohit Gupta (University of Wisconsin-Madison. USA)</i>	

Poster 4.1

Deep Learning

Cross-Dataset Person Re-Identification via Unsupervised Pose Disentanglement and Adaptation	7918
<i>Yu-Jhe Li (National Taiwan University), Ci-Siang Lin (National Taiwan University), Yan-Bo Lin (National Taiwan University), and Yu-Chiang Frank Wang (National Taiwan University)</i>	
A Learned Representation for Scalable Vector Graphics	7929
<i>Raphael Gontijo Lopes (Google Brain), David Ha (Google Brain), Douglas Eck (Google Brain), and Jonathon Shlens (Google)</i>	

ELF: Embedded Localisation of Features in Pre-Trained CNN	7939
<i>Assia Benbihi (GeorgiaTech Lorraine), Matthieu Geist (Google Brain), and Cedric Pradalier (GeorgiaTech Lorraine)</i>	
Joint Group Feature Selection and Discriminative Filter Learning for Robust Visual Object Tracking	7949
<i>Tianyang Xu (Jiangnan University), Zhen-Hua Feng (University of Surrey), Xiao-Jun Wu (Jiangnan University), and Josef Kittler (University of Surrey. UK)</i>	
Sampling Wisely: Deep Image Embedding by Top-K Precision Optimization	7960
<i>Jing Lu (Business Growth BU JD.com), Chaofan Xu (Harbin Institute of Technology), Wei Zhang (JD AI Research), Lingyu Duan (Peking University), and Tao Mei (AI Research of JD.com)</i>	
On the Global Optima of Kernelized Adversarial Representation Learning	7970
<i>Bashir Sadeghi (Michigan State University), Runyi Yu (Eastern Mediterranean University), and Vishnu Boddeti (Michigan State University)</i>	
Addressing Model Vulnerability to Distributional Shifts Over Image Transformation Sets	7979
<i>Riccardo Volpi (Istituto Italiano di Tecnologia) and Vittorio Murino (Istituto Italiano di Tecnologia)</i>	
Attract or Distract: Exploit the Margin of Open Set	7989
<i>Qianyu Feng (University of Technology Sydney), Guoliang Kang (UTS), Hehe Fan (UTS), and Yi Yang (University of Technology. Sydney)</i>	
MIC: Mining Interclass Characteristics for Improved Metric Learning	7999
<i>Biagio Brattoli (Heidelberg University), Karsten Roth (Heidelberg University), and Bjorn Ommers (Heidelberg University)</i>	
Self-Supervised Representation Learning via Neighborhood-Relational Encoding	8009
<i>Mohammad Sabokrou (Institute for Research in fundamental science. IPM), Mohammad Khalooei (Amirkabir University of Technology), and Ehsan Adeli (Stanford University)</i>	
AWSD: Adaptive Weighted Spatiotemporal Distillation for Video Representation	8019
<i>Mohammad Tavakolian (University of Oulu), Hamed Rezazadegan Tavakoli (Aalto University), and Abdenour Hadid (Finland)</i>	
Bilinear Attention Networks for Person Retrieval	8029
<i>Pengfei Fang (The Australian National University), Jieming Zhou (The Australian National University), Soumava Roy (AUSTRALIAN NATIONAL UNIVERSITY), Lars Petersson (Data61/CSIRO), and Mehrtash Harandi (Monash University)</i>	
Discriminative Feature Learning With Consistent Attention Regularization for Person Re-Identification	8039
<i>Sanping Zhou (Xi'an Jiaotong University), Fei Wang (Xi'an Jiaotong University), Zeyi Huang (Carnegie Mellon University), and Jinjun Wang (Xi'an Jiaotong University)</i>	
Semi-Supervised Domain Adaptation via Minimax Entropy	8049
<i>Kuniaki Saito (Boston University), Donghyun Kim (Boston University), Stan Sclaroff (Boston University), Trevor Darrell (UC Berkeley), and Kate Saenko (Boston University)</i>	

Boosting Few-Shot Visual Learning With Self-Supervision	8058
<i>Spyros Gidaris (valeo.ai), Andrei Bursuc (Valeo), Nikos Komodakis (ENPC. France), Patrick Pérez Pérez (Valeo.ai), and Matthieu Cord (Sorbonne University)</i>	
FDA: Feature Disruptive Attack	8068
<i>Aditya Ganeshan (Indian Institute of Science), Vivek B S (Indian Institute of Science), and Venkatesh Babu Radhakrishnan (Indian Institute of Science)</i>	
A Novel Unsupervised Camera-Aware Domain Adaptation Framework for Person Re-Identification	8079
<i>Lei Qi (Nanjing University), Lei Wang (University of Wollongong. Australia), Jing Huo (Nanjing University), Luping Zhou (University of Sydney), Yinghuan Shi (Nanjing University), and Yang Gao (Nanjing University)</i>	
Recover and Identify: A Generative Dual Model for Cross-Resolution Person Re-Identification	8089
<i>Yu-Jhe Li (National Taiwan University), Yun-Chun Chen (Academia Sinica), Yen-Yu Lin (National Chiao Tung University), Xiaofei Du (Umbo Computer Vision), and Yu-Chiang Frank Wang (National Taiwan University)</i>	
Cross-View Policy Learning for Street Navigation	8099
<i>Ang Li (DeepMind. Mountain View), Huiyi Hu (Google), Piotr Mirowski (DeepMind), and Mehrdad Farajtabar (DeepMind)</i>	
Learning Across Tasks and Domains	8109
<i>Pierluigi Zama Ramirez (University of Bologna), Alessio Tonioni (University of Bologna), Samuele Salti (University of Bologna), and Luigi Di Stefano (University of Bologna)</i>	
EMPNet: Neural Localisation and Mapping Using Embedded Memory Points	8119
<i>Gil Avraham (Monash University), Yan Zuo (Monash University), Thanuja Dharmasiri (Monash University), and Tom Drummond (Monash University)</i>	
AVT: Unsupervised Learning of Transformation Equivariant Representations by Autoencoding Variational Transformations	8129
<i>Guo-Jun Qi (Futurewei Technologies), Liheng Zhang (University of Central Florida), Chang Wen Chen (SUNY at Buffalo. USA), and Qi Tian (Huawei Noah's Ark Lab)</i>	
Composite Shape Modeling via Latent Space Factorization	8139
<i>Anastasia Dubrovina (Lyft), Fei Xia (Stanford University), Panos Achlioptas (Stanford University), Mira Shalah (Stanford University), Raphael Groskot (PSL Research University), and Leonidas Guibas (Stanford University)</i>	
Deep Comprehensive Correlation Mining for Image Clustering	8149
<i>Jianlong Wu (Peking University), Keyu Long (University of Chinese Academy of Sciences), Fei Wang (SenseTime), Chen Qian (SenseTime), Cheng Li (SenseTime Research), Zhouchen Lin (Peking University), and Hongbin Zha (Peking University. China)</i>	
Unsupervised Multi-Task Feature Learning on Point Clouds	8159
<i>Kaveh Hassani (Autodesk) and Mike Haley (Autodesk)</i>	

Reciprocal Multi-Layer Subspace Learning for Multi-View Clustering	8171
<i>Ruihuang Li (Tian Jin University), Changqing Zhang (Tianjin university), Huazhu Fu (Inception Institute of Artificial Intelligence), Xi Peng (College of Computer Science. Sichuan Univerisity), Joey Tianyi Zhou (IHPC. ASTAR), and Qinghua Hu (Tianjin University)</i>	
Geometric Disentanglement for Generative Latent Shape Models	8180
<i>Tristan Aumentado-Armstrong (University of Toronto), Stavros Tsogkas (University of Toronto), Allan Jepson (Samsung), and Sven Dickinson (University of Toronto)</i>	
GAN-Tree: An Incrementally Learned Hierarchical Generative Framework for Multi-Modal Data Distributions	8190
<i>Jogendra Nath Kundu (Indian Institute of Science), Maharshi Gor (Indian Institute of Science), Dakshit Agrawal (Indian Institute of Technology Roorkee. IITR), and Venkatesh Babu Radhakrishnan (Indian Institute of Science)</i>	
GODS: Generalized One-Class Discriminative Subspaces for Anomaly Detection	8200
<i>Jue Wang (ANU) and Anoop Cherian (MERL)</i>	
Neighborhood Preserving Hashing for Scalable Video Retrieval	8211
<i>Shuyan Li (Tsinghua University), Zhixiang Chen (Tsinghua University), Jiwen Lu (Tsinghua University), Xiu Li (Tsinghua University), and Jie Zhou (Tsinghua University)</i>	

Recognition

Self-Training With Progressive Augmentation for Unsupervised Cross-Domain Person Re-Identification	8221
<i>Xinyu Zhang (Tongji University), Jiewei Cao (The University of Adelaide), Chunhua Shen (University of Adelaide), and Mingyu You (Tongji University)</i>	
SCRDet: Towards More Robust Detection for Small, Cluttered and Rotated Objects	8231
<i>Xue Yang (Shanghai Jiao Tong University), Jirui Yang (Institute of Electrics. Chinese Academy of Sciences), Junchi Yan (Shanghai Jiao Tong University), Yue Zhang (Institute of Electrics. Chinese Academy of Sciences), Tengfei Zhang (Institute of Electrics. Chinese Academy of Sciences), Zhi Guo (Institute of Electrics. Chinese Academy of Sciences), Xian Sun (IECAS), and Kun Fu (Institute of Electrics. Chinese Academy of Sciences)</i>	
Cross-X Learning for Fine-Grained Visual Categorization	8241
<i>Wei Luo (South China Agricultural University), Xitong Yang (UMD), Xianjie Mo (South China Agricultural University), Yuheng Lu (University of Maryland), Larry Davis (University of Maryland), Jun Li (MIT), Jian Yang (Nanjing University of Science and Technology), and Ser-Nam Lim (Facebook AI)</i>	
Maximum-Margin Hamming Hashing	8251
<i>Rong Kang (Tsinghua University. China), Yue Cao (Tsinghua University), Mingsheng Long (Tsinghua University), Jianmin Wang (Tsinghua University. China), and Philip S Yu (UIC)</i>	

Conservative Wasserstein Training for Pose Estimation	8261
<i>Xiaofeng Liu (CMU), Yang Zou (Carnegie Mellon University), Tong Che (MILA), Ping Jia (CIOMP), Peng Ding (Changchun Institute of Optics. Fine Mechanics and Physics), Jane You (HK Poly U), and B. V. K. Vijaya Kumar (CMU. USA)</i>	
Learning to Rank Proposals for Object Detection	8272
<i>Zhiyu Tan (Alibaba Group), Xuecheng Nie (NUS), Qi Qian (Alibaba Group), Nan Li (Alibaba), and Hao Li (Alibaba Group)</i>	
Vehicle Re-Identification With Viewpoint-Aware Metric Learning	8281
<i>Ruihang Chu (Beihang Univeristy), Yifan Sun (Tsinghua University), Yadong Li (Megvii Inc. Face++), Zheng Liu (Megvii Research), Chi Zhang (Megvii Inc.), and Yichen Wei (Megvii Research Shanghai)</i>	
WSOD2: Learning Bottom-Up and Top-Down Objectness Distillation for Weakly-Supervised Object Detection	8291
<i>Zhaoyang Zeng (Sun Yat-sen University), Bei Liu (Microsoft Research), Jianlong Fu (Microsoft Research), Hongyang Chao (Sun Yat-sen University), and Lei Zhang (Microsoft)</i>	
Localization of Deep Inpainting Using High-Pass Fully Convolutional Network	8300
<i>Haodong Li (Shenzhen University) and Jiwu Huang (Shenzhen University)</i>	
Clustered Object Detection in Aerial Images	8310
<i>Fan Yang (Temple University), Heng Fan (Temple University), Peng Chu (Temple University), Erik Blasch (air force research lab), and Haibin Ling (Temple University)</i>	
Unsupervised Graph Association for Person Re-Identification	8320
<i>Jinlin Wu (Institute of Automation. Chinese Academy of Sciences. Beijing. China), Hao Liu (NLPR. CASIA), Yang Yang (Chinese Academy of Sciences), Zhen Lei (NLPR. CASIA. China), Shengcai Liao (Inception Institute of Artificial Intelligence), and Stan Li (National Lab. of Pattern Recognition. China)</i>	
Learning a Mixture of Granularity-Specific Experts for Fine-Grained Categorization	8330
<i>Lianbo Zhang (University of Technology Sydney), Shaoli Huang (University of Sydney), Wei Liu (UTS), and Dacheng Tao (University of Sydney)</i>	
advPattern: Physical-World Attacks on Deep Person Re-Identification via Adversarially Transformable Patterns	8340
<i>Zhibo Wang (Wuhan University), Siyan Zheng (Wuhan University), Mengkai Song (Wuhan University), Qian Wang (Wuhan University), Alireza Rahimpour (University of Tennessee-Knoxville), and Hairong Qi (University of Tennessee-Knoxville)</i>	
ABD-Net: Attentive but Diverse Person Re-Identification	8350
<i>Tianlong Chen (Texas A&M University), Shaojin Ding (Texas A&M University), Jingyi Xie (University of Science and Technology of China), Ye Yuan (Texas A&M University), Wuyang Chen (Texas A&M University), Yang Yang (Walmart Technology), Zhou Ren (Wormpex AI Research), and Zhangyang Wang (TAMU)</i>	

From Open Set to Closed Set: Counting Objects by Spatial Divide-and-Conquer	8361
<i>Haipeng Xiong (Huazhong University of Science and Technology), Hao Lu (The University of Adelaide), Chengxin Liu (Huazhong University of Science and Technology), Liang Liu (Huazhong University of Science and Technology), Zhiguo Cao (Huazhong Univ. of Sci.&Tech.), and Chunhua Shen (University of Adelaide)</i>	
Towards Precise End-to-End Weakly Supervised Object Detection Network	8371
<i>Ke Yang (NUDT), Dongsheng Li (School of Computer Science, National University of Defense Technology), and Yong Dou (National University of Defense Technology)</i>	
Learn to Scale: Generating Multipolar Normalized Density Maps for Crowd Counting	8381
<i>Chenfeng Xu (Huazhong University of Science and Technology), Kai Qiu (Microsoft Research), Jianlong Fu (Microsoft Research), Song Bai (University of Oxford), Yongchao Xu (Huazhong University of Science and Technology), and Xiang Bai (Huazhong University of Science and Technology)</i>	
Ground-to-Aerial Image Geo-Localization With a Hard Exemplar Reweighting Triplet Loss	8390
<i>Sudong Cai (Sun Yat-sen University, SYSU), Yulan Guo (National University of Defense Technology), Salman Khan (Australian National University, ANU), Jiwei Hu (Wuhan University of Technolgoy), and Gongjian Wen (National University of Defense Technology)</i>	
Learning to Discover Novel Visual Categories via Deep Transfer Clustering	8400
<i>Kai Han (University of Oxford), Andrea Vedaldi (Oxford University), and Andrew Zisserman (University of Oxford)</i>	
AM-LFS: AutoML for Loss Function Search	8409
<i>Chuming Li (sensenet), Xin Yuan (Tsinghua University), Chen Lin (Sensenet), Minghao Guo (Sensenet), Wei Wu (SenseTime Group Limited), Junjie Yan (Sensenet Group Limited), and Wanli Ouyang (The University of Sydney)</i>	
Few-Shot Object Detection via Feature Reweighting	8419
<i>Bingyi Kang (National University of Singapore), Zhuang Liu (UC Berkeley), Xin Wang (UC Berkeley), Fisher Yu (UC Berkeley), Jiashi Feng (NUS), and Trevor Darrell (UC Berkeley)</i>	
Objects365: A Large-Scale, High-Quality Dataset for Object Detection	8429
<i>Shuai Shao (Megvii, Face++), Zeming Li (Megvii Face++ Inc), Tianyuan Zhang (Peking University), Chao Peng (Megvii Face++ Inc), Gang Yu (Megvii Inc), Xiangyu Zhang (Megvii Technology), Jing Li (MEGVII), and Jian Sun (Megvii Technology)</i>	
Efficient and Accurate Arbitrary-Shaped Text Detection With Pixel Aggregation Network	8439
<i>Wenhai Wang (Nanjing university), Enze Xie (Tongji University), Xiaoge Song (Nanjing University), Yuhang Zang (University of Electronic Science and Technology of China), Wenjia Wang (Tongji University), Tong Lu (Nanjing University), Gang Yu (Megvii Inc), and Chunhua Shen (University of Adelaide)</i>	
Foreground-Aware Pyramid Reconstruction for Alignment-Free Occluded Person Re-Identification	8449
<i>He Lingxiao (casia), Yinggang Wang (JD AI Research), Wu Liu (AI Research of JD.com), He Zhao (JD AI Research), Zhenan Sun (Chinese of Academy of Sciences), and Jiashi Feng (NUS)</i>	

Collect and Select: Semantic Alignment Metric Learning for Few-Shot Learning	8459
<i>Fusheng Hao (Shenzhen Institutes of Advanced Technology. Chinese Academy of Sciences), Fengxiang He (The University of Sydney), Jun Cheng (Shenzhen Institutes of Advanced Technology. Chinese Academy of Sciences), Lei Wang (Shenzhen Institutes of Advanced Technology. Chinese Academy of Sciences), Jianzhong Cao (Xi'an Institute of Optics and Precision Mechanics), and Dacheng Tao (University of Sydney)</i>	

Segmentation, Grouping, & Shape

Bayesian Adaptive Superpixel Segmentation	8469
<i>Roy Uziel (Ben-Gurion University), Meitar Ronen (Ben-Gurion University), and Oren Freifeld (Ben-Gurion University)</i>	
CapsuleVOS: Semi-Supervised Video Object Segmentation Using Capsule Routing	8479
<i>Kevin Duarte (University of Central Florida), Yogesh Rawat (University of Central Florida), and Mubarak Shah (University of Central Florida)</i>	
BAE-NET: Branched Autoencoder for Shape Co-Segmentation	8489
<i>Zhiqin Chen (Simon Fraser University), Kangxue Yin (Simon Fraser University), Matthew Fisher (Adobe Research), Siddhartha Chaudhuri (Adobe Research), and Hao Zhang (Simon Fraser University)</i>	
VV-Net: Voxel VAE Net With Group Convolutions for Point Cloud Segmentation	8499
<i>Hsien-Yu Meng (Tsinghua University), Lin Gao (Institute of Computing Technology. Chinese Academy of Sciences), Yu-Kun Lai (Cardiff University), and Dinesh Manocha (UMD)</i>	
Miss Detection vs. False Alarm: Adversarial Learning for Small Object Segmentation in Infrared Images	8508
<i>Huan Wang (Nanjing University of Science & Technology), Luping Zhou (University of Sydney), and Lei Wang (University of Wollongong. Australia)</i>	
Group-Wise Deep Object Co-Segmentation With Co-Attention Recurrent Neural Network	8518
<i>Bo Li (Nanjing University), Zhengxing Sun (Nanjing University), Qian Li (Nanjing University), Yunjie Wu (Nanjing University), and Hu Anqi (Nanjing University)</i>	

Statistics, Physics, Theory & Datasets

Human Attention in Image Captioning: Dataset and Analysis	8528
<i>Sen He (University of Exeter), Hamed Rezazadegan Tavakoli (Aalto University), Ali Borji (University of Central Florida), and Nicolas Pugeault (Exeter)</i>	
Variational Uncalibrated Photometric Stereo Under General Lighting	8538
<i>Zhenzhang Ye (TU Munich), Bjoern Haefner (Technical University of Munich), Maolin Gao (Artisense), Tao Wu (TU Munich), Yvain Queau (CNRS), and Daniel Cremers (TU Munich)</i>	

SPLINE-Net: Sparse Photometric Stereo Through Lighting Interpolation and Normal Estimation Networks ...	8548
<i>Qian Zheng (Nanyang Technological University), Yiming Jia (Tsinghua University), Boxin Shi (Peking University), Xudong Jiang (Nanyang Technological University), Lingyu Duan (Peking University), and Alex Kot (Nanyang Technological University)</i>	
Hyperspectral Image Reconstruction Using Deep External and Internal Learning	8558
<i>Tao Zhang (Beijing Institute of Technology), Ying Fu (Beijing Institute of Technology), Lizhi Wang (Beijing Institute of Technology), and Hua Huang (Beijing Institute of Technology)</i>	
Gravity as a Reference for Estimating a Person’s Height From Video	8568
<i>Didier Bieler (EPFL), Semih Günel Günel (EPFL), Pascal Fua (EPFL, Switzerland), and Helge Rhodin (EPFL)</i>	
Shadow Removal via Shadow Image Decomposition	8577
<i>Hieu Le (Stony Brook University) and Dimitris Samaras (Stony Brook University)</i>	
OperatorNet: Recovering 3D Shapes From Difference Operators	8587
<i>Ruqi Huang (LIX, Ecole Polytechnique), Marie-Julie Rakotosaona (Ecole Polytechnique), Panos Achlioptas (Stanford University), Leonidas Guibas (Stanford University), and Maks Ovsjanikov (Ecole polytechnique)</i>	
Neural Inverse Rendering of an Indoor Scene From a Single Image	8597
<i>Soumyadip Sengupta (University of Washington), Jinwei Gu (NVIDIA), Kihwan Kim (NVIDIA), Guilin Liu (NVIDIA), David Jacobs (University of Maryland, USA), and Jan Kautz (NVIDIA)</i>	

3D From Single View & RGBD

ForkNet: Multi-Branch Volumetric Semantic Completion From a Single Depth Image	8607
<i>Yida Wang (Technische Universität München), David Joseph Tan (Google), Nassir Navab (TU Munich, Germany), and Federico Tombari (TUM, Google)</i>	
Moving Indoor: Unsupervised Video Depth Learning in Challenging Environments	8617
<i>Junsheng Zhou (Tsinghua University), Yuwang Wang (Microsoft Research), Kaihuai Qin (Tsinghua University), and Wenjun Zeng (Microsoft Research)</i>	
GraphX-Convolution for Point Cloud Deformation in 2D-to-3D Conversion	8627
<i>Duc Nguyen (Yonsei University), Seonghwa Choi (Yonsei University), Woojae Kim (Yonsei University), and Sanghoon Lee (Yonsei University, Korea)</i>	
FrameNet: Learning Local Canonical Frames of 3D Surfaces From a Single RGB Image	8637
<i>Jingwei Huang (Stanford University), Yichao Zhou (UC Berkeley), Thomas Funkhouser (Princeton University), and Leonidas Guibas (Stanford University)</i>	
Holistic++ Scene Understanding: Single-View 3D Holistic Scene Parsing and Human Pose Estimation With Human-Object Interaction and Physical Commonsense	8647
<i>Yixin Chen (UCLA), Siyuan Huang (UCLA), Tao Yuan (UCLA), Yixin Zhu (UCLA), Siyuan Qi (UCLA), and Song-Chun Zhu (UCLA)</i>	

Action & Video

- MMAct: A Large-Scale Dataset for Cross Modal Human Action Understanding 8657
Quan Kong (Hitachi), Ziming Wu (Hong Kong University of Science and Technology), Ziwei Deng (Hitachi), Martin Klinkigt (Hitachi. Ltd.), Bin Tong (HITACHI), and Tomokazu Murakami (Hitachi)
- HACS: Human Action Clips and Segments Dataset for Recognition and Temporal Localization 8667
Hang Zhao (MIT), Antonio Torralba (MIT), Lorenzo Torresani (Dartmouth College), and Zhicheng Yan (University of Illinois at Urbana-Champaign)
- 3C-Net: Category Count and Center Loss for Weakly-Supervised Action Localization 8678
Sanath Narayan (Inception Institute of Artificial Intelligence), Hisham Cholakkal (Inception Institute of Artificial Intelligence), Fahad Shahbaz Khan (Inception Institute of Artificial Intelligence), and Ling Shao (Inception Institute of Artificial Intelligence)
- Grounded Human-Object Interaction Hotspots From Video 8687
Tushar Nagarajan (UT Austin), Christoph Feichtenhofer (Facebook AI Research), and Kristen Grauman (Facebook AI Research & UT Austin)
- Hallucinating IDT Descriptors and I3D Optical Flow Features for Action Recognition With CNNs 8697
Lei Wang (University of Western Australia), Piotr Koniusz (Data61/CSIRO. ANU), and Du Huynh (University of Western Australia. Australia)

Computational Photography & Graphics

- Learning to Paint With Model-Based Deep Reinforcement Learning 8708
Zhewei Huang (Peking University), Shuchang Zhou (Megvii), and Wen Heng (Megvii inc.)
- Neural Re-Simulation for Generating Bounces in Single Images 8718
Carlo Innamorati (University College London), Bryan Russell (Adobe Research), Danny Kaufman (Adobe Research), and Niloy Mitra (university college london)
- Deep Appearance Maps 8728
Maxim Maximov (TUM), Tobias Ritschel (UCL), Laura Leal-Taixé (TUM), and Mario Fritz (CISPA Helmholtz Center for Information Security)
- GarNet: A Two-Stream Network for Fast and Accurate 3D Cloth Draping 8738
Erhan Gundogdu (EPFL), Victor Constantin (EPFL), Amrollah Seifoddini (Fision Technologies), Minh Dang (Fision Technologies), Mathieu Salzmann (EPFL), and Pascal Fua (EPFL. Switzerland)
- Joint Embedding of 3D Scan and CAD Objects 8748
Manuel Dahnert (Technical University of Munich), Angela Dai (Technical University of Munich), Leonidas Guibas (Stanford University), and Matthias Niessner (Technical University of Munich)

CompoNet: Learning to Generate the Unseen by Part Synthesis and Composition	8758
<i>Nadav Schor (Tel Aviv University), Oren Katzir (Tel Aviv University), Hao Zhang (Simon Fraser University), and Daniel Cohen-Or (Tel Aviv University)</i>	
DDSL: Deep Differentiable Simplex Layer for Learning Geometric Signals	8768
<i>Chiyu Jiang (UC Berkeley), Dana Lansigan (UC Berkeley), Philip Marcus (UC Berkeley), and Matthias Niessner (Technical University of Munich)</i>	

Low-Level & Optimization

EGNet: Edge Guidance Network for Salient Object Detection	8778
<i>Jiaxing Zhao (Nankai University), Jiang-Jiang Liu (Nankai University), Deng-Ping Fan (Inception Institute of Artificial Intelligence), Yang Cao (Nankai University), Jufeng Yang (Nankai University), and Ming-Ming Cheng (Nankai University)</i>	
SID4VAM: A Benchmark Dataset With Synthetic Images for Visual Attention Modeling	8788
<i>David Berge (Computer Vision Center), Xosé Ramón Fernández Vidal (University of Santiago de Compostela), Xavier Otazu (Computer Vision Center), and Xosé M. Pardo (Universidade de Santiago de Compostela)</i>	
Two-Stream Action Recognition-Oriented Video Super-Resolution	8798
<i>Haochen Zhang (University of Science and Technology of China), Dong Liu (University of Science and Technology of China), and Zhiwei Xiong (University of Science and Technology of China)</i>	
Where Is My Mirror?	8808
<i>Xin Yang (Dalian University of Technology), Haiyang Mei (Dalian University of Technology), Ke Xu (Dalian University of Technology; City University of Hong Kong), Xiaopeng Wei (Dalian University of Technology), Baocai Yin (Dalian University of Technology), and Rynson Lau (City University of Hong Kong)</i>	
Disentangled Image Matting	8818
<i>Shaofan Cai (Megvii Inc. face++), Xiaoshuai Zhang (Peking University), Haoqiang Fan (Megvii Inc face++), Haibin Huang (Megvii Technology), Jiangyu Liu (megvii inc), Jiaming Liu (Megvii), Jiaying Liu (Peking University), Jue Wang (Megvii Technology), and Jian Sun (Megvii Technology)</i>	
Guided Super-Resolution As Pixel-to-Pixel Transformation	8828
<i>Riccardo De Lutio (ETHZ), Stefano D'aronco (ETHZ), Jan Dirk Wegner (ETH Zurich), and Konrad Schindler (ETH)</i>	
Deep Learning for Light Field Saliency Detection	8837
<i>Tiantian Wang (Dalian University of Technology), Yongri Piao (Dalian University of Technology), Huchuan Lu (Dalian University of Technology), Xiao Li (Dalian University of Technology), and Lihe Zhang (Dalian University of Technology)</i>	
Optimizing the F-Measure for Threshold-Free Salient Object Detection	8848
<i>Kai Zhao (Nankai University), Shanghua Gao (Nankai University), Wenguan Wang (Inception Institute of Artificial Intelligence), and Ming-Ming Cheng (Nankai University)</i>	

Image Inpainting With Learnable Bidirectional Attention Maps	8857
<i>Chaohao Xie (Harbin Institute of Technology), Shaohui Liu (Harbin Institute of Technology), Chao Li (Baidu), Ming-Ming Cheng (Nankai University), Wangmeng Zuo (Harbin Institute of Technology, China), Xiao Liu (Baidu), Shilei Wen (Baidu Research), and Errui Ding (Baidu Inc.)</i>	
Joint Demosaicking and Denoising by Fine-Tuning of Bursts of Raw Images	8867
<i>Thibaud Ehret (CMLA, ENS Cachan), Axel Davy (ENS Paris-Saclay), Pablo Arias (ENS Paris-Saclay), and Gabriele Facciolo (ENS Paris - Saclay)</i>	
DeblurGAN-v2: Deblurring (Orders-of-Magnitude) Faster and Better	8877
<i>Orest Kupyn (Ukrainian Catholic University), Tetiana Martyniuk (Ukrainian Catholic University), Junru Wu (Texas A&M University), and Zhangyang Wang (TAMU)</i>	

Language & Reasoning

Reflective Decoding Network for Image Captioning	8887
<i>Lei Ke (HKUST), Wenjie Pei (Tencent), Ruiyu Li (Tencent), Xiaoyong Shen (Tencent), and Yu-Wing Tai (Tencent)</i>	
Joint Optimization for Cooperative Image Captioning	8897
<i>Gilad Vered (Bar-Ilan University), Gal Oren (Bar Ilan University), Yuval Atzmon (Bar-Ilan University, NVIDIA Research), and Gal Chechik (Bar Ilan University)</i>	
Watch, Listen and Tell: Multi-Modal Weakly Supervised Dense Event Captioning	8907
<i>Tanzila Rahman (University of British Columbia), Bicheng Xu (University of British Columbia), and Leonid Sigal (University of British Columbia)</i>	
Joint Syntax Representation Learning and Visual Cue Translation for Video Captioning	8917
<i>Jingyi Hou (Beijing Institute of Technology), Xinxiao Wu (Beijing Institute of Technology), Wentian Zhao (Beijing Institute of Technology), Jiebo Luo (University of Rochester), and Yunde Jia (Beijing Institute of Technology)</i>	
Entangled Transformer for Image Captioning	8927
<i>Guang Li (University of Technology Sydney), Linchao Zhu (University of Technology, Sydney), Ping Liu (UTS), and Yi Yang (UTS)</i>	
Shapeglot: Learning Language for Shape Differentiation	8937
<i>Panos Achlioptas (Stanford University), Leonidas Guibas (Stanford University), Noah Goodman (Stanford University), Judy Fan (Stanford), and Robert Hawkins (Stanford)</i>	
nocaps: novel object captioning at scale	8947
<i>Harsh Agrawal (Georgia Institute of Technology), Karan Desai (Georgia Tech), Yufei Wang (Macquarie University), Xinlei Chen (Facebook AI Research), Rishabh Jain (Georgia Tech), Mark Johnson (Macquarie University), Dhruv Batra (Georgia Tech & Facebook AI Research), Devi Parikh (Georgia Tech & Facebook AI Research), Stefan Lee (Georgia Institute of Technology), and Peter Anderson (Georgia Tech)</i>	

3D From Multiview & Sensors

Fully Convolutional Geometric Features	8957
<i>Christopher Choy (Stanford University), Jaesik Park (POSTECH), and Vladlen Koltun (Intel Labs)</i>	
Learning Local RGB-to-CAD Correspondences for Object Pose Estimation	8966
<i>Georgios Georgakis (GMU), Srikrishna Karanam (Siemens Corporate Technology, Princeton), Ziyang Wu (Siemens Corporation), and Jana Kosecka (George Mason University)</i>	
Depth From Videos in the Wild: Unsupervised Monocular Depth Learning From Unknown Cameras	8976
<i>Ariel Gordon (Google Research), Hanhan Li (Google AI), Rico Jonschkowski (Google), and Anelia Angelova (Google)</i>	
OmniMVS: End-to-End Learning for Omnidirectional Stereo Matching	8986
<i>Changhee Won (Hanyang University), Jongbin Ryu (Hanyang University), and Jongwoo Lim (Hanyang University)</i>	
On the Over-Smoothing Problem of CNN Based Disparity Estimation	8996
<i>Chuangrong Chen (Sun Yat-Sen University), Xiaozhi Chen (DJI), and Hui Cheng (Sun Yat-Sen University)</i>	

Image & Video Synthesis

Disentangling Propagation and Generation for Video Prediction	9005
<i>Hang Gao (UC Berkeley), Huazhe Xu (UC Berkeley), Qi-Zhi Cai (Sinovation Ventures AI Institute), Ruth Wang (UC Berkeley), Fisher Yu (UC Berkeley), and Trevor Darrell (UC Berkeley)</i>	
Guided Image-to-Image Translation With Bi-Directional Feature Transformation	9015
<i>Badour Albahar (Virginia Tech) and Jia-Bin Huang (Virginia Tech)</i>	
Towards Multi-Pose Guided Virtual Try-On Network	9025
<i>Haoye Dong (Sun Yat-sen University), Xiaodan Liang (Sun Yat-sen University), Xiaohui Shen (ByteDance AI Lab), Bochao Wang (Sun Yet-sen University), Hanjiang Lai (Sun Yat-Sen university), Jia Zhu (South China Normal University), Zhiting Hu (CMU), and Jian Yin (Sun Yat-Sen University)</i>	
Photorealistic Style Transfer via Wavelet Transforms	9035
<i>Jaeyun Yoo (Clova AI Research, NAVER Corp.), Youngjung Uh (Clova AI Research, NAVER Corp.), Sanghyuk Chun (Clova AI Research, NAVER Corp.), Byeongkyu Kang (Yonsei University), and Jung-Woo Ha (Clova AI Research, NAVER Corp.)</i>	
Personalized Fashion Design	9045
<i>Cong Yu (University of Electronic Science and Technology of China), Yang Hu (University of Electronic Science and Technology of China), Yan Chen (University of Electronic Science and Technology of China), and Bing Zeng (University of Electronic Science and Technology of China)</i>	
Tag2Pix: Line Art Colorization Using Text Tag With SECat and Changing Loss	9055
<i>Hyunsu Kim (Seoul National University), Ho Young Jhoo (Seoul National University), Eunhyeok Park (Seoul National University), and Sungjoo Yoo (Seoul National University)</i>	

Free-Form Video Inpainting With 3D Gated Convolution and Temporal PatchGAN	9065
<i>Ya-Liang Chang (National Taiwan University), Zhe Yu Liu (National Taiwan University), Kuan-Ying Lee (National Taiwan University), and Winston Hsu (National Taiwan University)</i>	

Applications, Medical & Robotics

TextDragon: An End-to-End Framework for Arbitrary Shaped Text Spotting	9075
<i>Wei Feng (CAS), Wenhao He (CASIA), Fei Yin (Institute of Automation of Chinese Academy of Sciences), Xu-Yao Zhang (Institute of Automation of Chinese Academy of Sciences), and Cheng-Lin Liu (Institute of Automation of Chinese Academy of Sciences)</i>	
Chinese Street View Text: Large-Scale Chinese Text Reading With Partially Supervised Learning	9085
<i>Yipeng Sun (Baidu Inc.), Jiaming Liu (Baidu Inc.), Wei Liu (The University of Hong Kong), Junyu Han (Baidu Inc.), Errui Ding (Baidu Inc.), and Jingtuo Liu (baidu)</i>	
Deep Floor Plan Recognition Using a Multi-Task Network With Room-Boundary-Guided Attention	9095
<i>Zhiliang Zeng (The Chinese University of Hong Kong), Xianzhi Li (The Chinese University of Hong Kong), Ying Kin Yu (The Chinese University of Hong Kong), and Chi-Wing Fu (The Chinese University of Hong Kong)</i>	
GA-DAN: Geometry-Aware Domain Adaptation Network for Scene Text Detection and Recognition	9104
<i>Fangneng Zhan (Nanyang Technological University), Chuhui Xue (Nanyang Technological University), and Shijian Lu (Nanyang Technological University)</i>	
Large-Scale Tag-Based Font Retrieval With Generative Feature Learning	9115
<i>Tianlang Chen (University of Rochester), Zhaowen Wang (Adobe Research), Ning Xu (Adobe Research), Hailin Jin (Adobe Research), and Jiebo Luo (U. Rochester)</i>	
Convolutional Character Networks	9125
<i>Linjie Xing (Malong Technologies), Zhi Tian (The University of Adelaide), Weilin Huang (Malong Technologies), and Matthew Scott (Malong Technologies)</i>	
Geometry Normalization Networks for Accurate Scene Text Detection	9136
<i>Jiaqi Duan (Huazhong University of Science and Technology), Youjiang Xu (SenseTime), Zhanghui Kuang (Sensetime Ltd.), Xiaoyu Yue (SenseTime), Hongbin Sun (Sensetime Ltd.), Yue Guan (Huazhong University of Science and Technology), and Wayne Zhang (SenseTime Research)</i>	
Symmetry-Constrained Rectification Network for Scene Text Recognition	9146
<i>Mingkun Yang (Huazhong University of Science and Technology), Yushuo Guan (Peking University), Minghui Liao (Huazhong University of Science and Technology), Xin He (Megvii), Kaigui Bian (Peking University), Song Bai (University of Oxford), Cong Yao (Megvii. Face++ Inc.), and Xiang Bai (Huazhong University of Science and Technology)</i>	

Oral 4.2A

Segmentation, Detection, 3D Scene Understanding

YOLACT: Real-Time Instance Segmentation	9156
<i>Daniel Bolya (University of California. Davis), Chong Zhou (University of California. Davis), Fanyi Xiao (University of California Davis), and Yong Jae Lee (University of California. Davis)</i>	
Expectation-Maximization Attention Networks for Semantic Segmentation	9166
<i>Xia Li (Peking University Shenzhen Graduate School), Zhisheng Zhong (Peking University), Jianlong Wu (Peking University), Yibo Yang (Peking University), Zhouchen Lin (Peking University), and Hong Liu (Peking University Shenzhen Graduate School)</i>	
Multi-Class Part Parsing With Joint Boundary-Semantic Awareness	9176
<i>Yifan Zhao (Beihang University), Jia Li (Beihang University), Yu Zhang (Beihang University), and Yonghong Tian (PKU)</i>	
Explaining Neural Networks Semantically and Quantitatively	9186
<i>Runjin Chen (Shanghai Jiao Tong University), Hao Chen (Huazhong University of Science and Technology), Ge Huang (Shanghai Jiao Tong University), Jie Ren (Shanghai Jiao Tong University), and Quanshi Zhang (Shanghai Jiao Tong University)</i>	
PANet: Few-Shot Image Semantic Segmentation With Prototype Alignment	9196
<i>Kaixin Wang (National University of Singapore), Jun Hao Liew (NUS), Yingtian Zou (National University of Singapore), Daquan Zhou (National University of Singapore), and Jiashi Feng (NUS)</i>	
ShapeMask: Learning to Segment Novel Objects by Refining Shape Priors	9206
<i>Weicheng Kuo (University of California Berkeley), Anelia Angelova (Google), Jitendra Malik (University of California at Berkley), and Tsung-Yi Lin (Google Brain)</i>	
Sequence Level Semantics Aggregation for Video Object Detection	9216
<i>Haiping Wu (McGill University), Yuntao Chen (CASIA), Naiyan Wang (TuSimple), and Zhao-Xiang Zhang (Chinese Academy of Sciences. China)</i>	
Video Object Segmentation Using Space-Time Memory Networks	9225
<i>Seoung Wug Oh (Yonsei Univeristy), Joon-Young Lee (Adobe Research), Ning Xu (Adobe Research), and Seon Joo Kim (Yonsei Univ.)</i>	
Zero-Shot Video Object Segmentation via Attentive Graph Neural Networks	9235
<i>Wenguan Wang (Inception Institute of Artificial Intelligence), Xiankai Lu (Inception Institute of Artificial Intelligence), Jianbing Shen (Inception Institute of Artificial Intelligence), David Crandall (Indiana University), and Ling Shao (Inception Institute of Artificial Intelligence)</i>	
MeteorNet: Deep Learning on Dynamic 3D Point Cloud Sequences	9245
<i>Xingyu Liu (Stanford University), Mengyuan Yan (Stanford University), and Jeannette Bohg (Stanford)</i>	
3D Instance Segmentation via Multi-Task Metric Learning	9255
<i>Jean Lahoud (KAUST), Bernard Ghanem (KAUST), Martin R. Oswald (ETH Zurich), and Marc Pollefeys (ETH Zurich / Microsoft)</i>	

DeepGCNs: Can GCNs Go As Deep As CNNs?	9266
<i>Guohao Li (King Abdullah University of Science and Technology. KAUST), Matthias Müller (King Abdullah University of Science and Technology. KAUST), Ali Thabet (KAUST), and Bernard Ghanem (KAUST)</i>	
Deep Hough Voting for 3D Object Detection in Point Clouds	9276
<i>Charles R. Qi (Facebook AI Research), Or Litany (Facebook AI Research), Kaiming He (Facebook AI Research), and Leonidas Guibas (Stanford University)</i>	
M3D-RPN: Monocular 3D Region Proposal Network for Object Detection	9286
<i>Garrick Brazil (Michigan State University) and Xiaoming Liu (Michigan State University)</i>	
SemanticKITTI: A Dataset for Semantic Scene Understanding of LiDAR Sequences	9296
<i>Jens Behley (University of Bonn), Martin Garbade (University of Bonn), Andres Milioto (University of Bonn), Jan Quenzel (University of Bonn), Sven Behnke (University of Bonn), Cyrill Stachniss (University of Bonn), and Jürgen Gall (University of Bonn)</i>	
WoodScape: A Multi-Task, Multi-Camera Fisheye Dataset for Autonomous Driving	9307
<i>Senthil Yogamani (Valeo Vision Systems), Ciaran Hughes (Valeo), Jonathan Horgan (Valeo), Ganesh Sistu (Valeo Vision Systems), Sumanth Chennupati (Valeo North America), Michal Uricar (Valeo), Stefan Milz (Valeo), Martin Simon (Valeo), Karl Amende (Valeo), Christian Witt (Valeo), Hazem Rashed (Valeo), Sanjaya Nayak (Valeo), Saquib Mansoor (Valeo), Padraig Varley (Valeo), Xavier Perrotton (Valeo), Derek Odea (Valeo), and Patrick Pérez (Valeo.ai)</i>	
Scalable Place Recognition Under Appearance Change for Autonomous Driving	9318
<i>Dzung Doan (The University of Adelaide), Yasir Latif (The University of Adelaide), Tat-Jun Chin (University of Adelaide), Yu Liu (The University of Adelaide), Thanh-Toan Do (The University of Liverpool), and Ian Reid (University of Adelaide. Australia)</i>	
Exploring the Limitations of Behavior Cloning for Autonomous Driving	9328
<i>Felipe Codevilla (UAB), Eder Santana (University of Florida), Antonio Lopez (CVC & UAB), and Adrien Gaidon (Toyota Research Institute)</i>	
Habitat: A Platform for Embodied AI Research	9338
<i>Manolis Savva (Simon Fraser University), Abhishek Kadian (Facebook AI Research), Oleksandr Maksymets (Facebook AI Research), Yili Zhao (Facebook AI Research), Erik Wijmans (Georgia Tech), Bhavana Jain (Facebook AI Research), Julian Straub (Facebook Reality Labs), Jia Liu (Facebook AI Research), Vladlen Koltun (Intel Labs), Jitendra Malik (University of California at Berkley), Devi Parikh (Georgia Tech & Facebook AI Research), and Dhruv Batra (Georgia Tech & Facebook AI Research)</i>	

Oral 4.2B

Face & Body Modeling

Towards Interpretable Face Recognition	9347
<i>Bangjie Yin (Michigan State University), Luan Tran (Michigan State University), Haoxiang Li (Wormpex AI Research), Xiaohui Shen (ByteDance AI Lab), and Xiaoming Liu (Michigan State University)</i>	
Co-Mining: Deep Face Recognition With Noisy Labels	9357
<i>Xiaobo Wang (JD AI Research), Shuo Wang (JD AI Research), Hailin Shi (JD AI Research), Jun Wang (JD AI Research), and Tao Mei (AI Research of JD.com)</i>	
Few-Shot Adaptive Gaze Estimation	9367
<i>Seonwook Park (ETH Zurich), Shalini De Mello (NVIDIA Research), Pavlo Molchanov (NVIDIA), Umar Iqbal (NVIDIA Research), Otmar Hilliges (ETH Zurich), and Jan Kautz (NVIDIA)</i>	
Live Face De-Identification in Video	9377
<i>Oran Gafni (Facebook AI Research), Lior Wolf (Tel Aviv University, Israel), and Yaniv Taigman (Facebook)</i>	
Face Video Deblurring Using 3D Facial Priors	9387
<i>Wenqi Ren (Institute of Information Engineering, Chinese Academy of Sciences), Jiaolong Yang (Microsoft Research Asia), Senyou Deng (Institute of Information Engineering, Chinese Academy of Sciences), David Wipf (Microsoft Research), Xiaochun Cao (Chinese Academy of Sciences), and Xin Tong (Microsoft)</i>	
Semi-Supervised Monocular 3D Face Reconstruction With End-to-End Shape-Preserved Domain Transfer ...	9397
<i>Jingtian Piao (CUHK MMLab), Chen Qian (SenseTime), and Hongsheng Li (Chinese University of Hong Kong)</i>	
3D Face Modeling From Diverse Raw Scan Data	9407
<i>Feng Liu (Michigan State University), Luan Tran (Michigan State University), and Xiaoming Liu (Michigan State University)</i>	
A Decoupled 3D Facial Shape Model by Adversarial Training	9418
<i>Victoria Fernandez Abrevaya (Inria), Adnane Boukhayma (University of Oxford), Stefanie Wuhler (Inria), and Edmond Boyer (Inria)</i>	
Photo-Realistic Facial Details Synthesis From Single Image	9428
<i>Anpei Chen (shanghaiTech), Zhang Chen (ShanghaiTech University), Guli Zhang (ShanghaiTech University), Kenny Mitchell (Edinburgh Napier University and Disney Research), and Jingyi Yu (Shanghai Tech University)</i>	
S2GAN: Share Aging Factors Across Ages and Share Aging Trends Among Individuals	9439
<i>Zhenliang He (Key Lab of Intelligent Information Processing of Chinese Academy of Sciences, CAS, Institute of Computing Technology, CAS), Meina Kan (Institute of Computing Technology, Chinese Academy of Sciences), Shiguang Shan (Chinese Academy of Sciences), and Xilin Chen (Institute of Computing Technology, Chinese Academy of Sciences)</i>	
PuppetGAN: Cross-Domain Image Manipulation by Demonstration	9449
<i>Ben Usman (Boston University), Nick Dufour (Google Research), Kate Saenko (Boston University), and Chris Bregler (Google Research)</i>	

Few-Shot Adversarial Learning of Realistic Neural Talking Head Models	9458
<i>Egor Zakharov (Skolkovo Institute of Science and Technology), Aliaksandra Shysheya (Samsung), Egor Burkov (Skoltech), and Victor Lempitsky (Samsung)</i>	
Pose-Aware Multi-Level Feature Network for Human Object Interaction Detection	9468
<i>Bo Wan (ShanghaiTech University), Desen Zhou (ShanghaiTech University), Yongfei Liu (ShanghaiTech), Rongjie Li (ShanghaiTech University), and Xuming He (ShanghaiTech University)</i>	
TRB: A Novel Triplet Representation for Understanding 2D Human Body	9478
<i>Haodong Duan (Peking University), Kwan-Yee Lin (Peking university), Sheng Jin (Tsinghua University), Wentao Liu (SenseTime), Chen Qian (SenseTime), and Wanli Ouyang (The University of Sydney)</i>	
Learning Trajectory Dependencies for Human Motion Prediction	9488
<i>Wei Mao (the Australian National University), Miaomiao Liu (The Australian National University), Mathieu Salzmann (EPFL), and Hongdong Li (Australian National University. Australia)</i>	
Cross-Domain Adaptation for Animal Pose Estimation	9497
<i>Jinkun Cao (Shanghai Jiao Tong University), Hongyang Tang (Shanghai Jiao Tong university), Hao-Shu Fang (SJTU), Xiaoyong Shen (Tencent), Yu-Wing Tai (Tencent), and Cewu Lu (Shanghai Jiao Tong University)</i>	

Poster 4.2

Recognition

NOTE-RCNN: NOise Tolerant Ensemble RCNN for Semi-Supervised Object Detection	9507
<i>Jiyang Gao (Waymo), Jiang Wang (Google Research), Shengyang Dai (-), Li-Jia Li (Stanford), and Ram Nevatia (U of Southern California)</i>	
Unsupervised Out-of-Distribution Detection by Maximum Classifier Discrepancy	9517
<i>Qing Yu (The University of Tokyo) and Kiyoharu Aizawa (The University of Tokyo)</i>	
SBSGAN: Suppression of Inter-Domain Background Shift for Person Re-Identification	9526
<i>Yan Huang (University of Technology Sydney), Qiang Wu (University of Technology Sydney), Jingsong Xu (University of Technology Sydney), and Yi Zhong (University of Technology Sydney)</i>	
Enriched Feature Guided Refinement Network for Object Detection	9536
<i>Jing Nie (Tianjin University), Rao Muhammad Anwer (Inception Institute of Artificial Intelligence), Hisham Cholakkal (Inception Institute of Artificial Intelligence), Fahad Shahbaz Khan (Inception Institute of Artificial Intelligence), Yanwei Pang (Tianjin University), and Ling Shao (Inception Institute of Artificial Intelligence)</i>	
Deep Meta Metric Learning	9546
<i>Guangyi Chen (Tsinghua University), Tianren Zhang (Tsinghua University), Jiwen Lu (Tsinghua University), and Jie Zhou (Tsinghua University)</i>	

Discriminative Feature Transformation for Occluded Pedestrian Detection	9556
<i>Chunluan Zhou (Nanyang Technological University), Ming Yang (Horizon Robotics), and Junsong Yuan (State University of New York at Buffalo, USA)</i>	
Contextual Attention for Hand Detection in the Wild	9566
<i>Supreeth Narasimhaswamy (Stony Brook University), Zhengwei Wei (Stony Brook University), Yang Wang (Stony Brook University), Justin Zhang (Caltech), and Minh Hoai Nguyen (Stony Brook University)</i>	
Meta R-CNN: Towards General Solver for Instance-Level Low-Shot Learning	9576
<i>Xiaopeng Yan (Sun Yat-sen University), Ziliang Chen (Sun Yat-sen University), Anni Xu (Sun Yat-sen University), Xiaoxi Wang (Sun Yat-Sen University), Xiaodan Liang (Sun Yat-sen University), and Liang Lin (Sun Yat-sen University)</i>	
Pyramid Graph Networks With Connection Attentions for Region-Based One-Shot Semantic Segmentation ..	9586
<i>Chi Zhang (Nanyang Technological University), Guosheng Lin (Nanyang Technological University), Fayao Liu (Institute for Infocomm Research, ASTAR), Jiushuang Guo (Stanford University), Qingyao Wu (South China University of Technology), and Rui Yao (China University of Mining and Technology)</i>	
Presence-Only Geographical Priors for Fine-Grained Image Classification	9595
<i>Oisin Mac Aodha (Caltech), Elijah Cole (Caltech), and Pietro Perona (California Institute of Technology)</i>	
POD: Practical Object Detection With Scale-Sensitive Network	9606
<i>Junran Peng (Chinese Academy of Sciences), Ming Sun (sensetime.com), Zhao-Xiang Zhang (Chinese Academy of Sciences, China), Tieniu Tan (NLPR, China), and Junjie Yan (Sensetime Group Limited)</i>	
Human Uncertainty Makes Classification More Robust	9616
<i>Joshua Peterson (Princeton University), Ruairidh Battleday (Princeton University), Thomas Griffiths (Princeton University), and Olga Russakovsky (Princeton University)</i>	
FCOS: Fully Convolutional One-Stage Object Detection	9626
<i>Zhi Tian (The University of Adelaide), Chunhua Shen (University of Adelaide), Hao Chen (The University of Adelaide), and Tong He (The University of Adelaide)</i>	
Self-Critical Attention Learning for Person Re-Identification	9636
<i>Guangyi Chen (Tsinghua University), Chunze Lin (Tsinghua University), Liangliang Ren (Tsinghua University), Jiwen Lu (Tsinghua University), and Jie Zhou (Tsinghua University)</i>	
Temporal Knowledge Propagation for Image-to-Video Person Re-Identification	9646
<i>Xinqian Gu (University of Chinese Academy of Sciences), Bingpeng Ma (University of Chinese Academy of Sciences), Hong Chang (Chinese Academy of Sciences), Shiguang Shan (Chinese Academy of Sciences), and Xilin Chen (Institute of Computing Technology, Chinese Academy of Sciences)</i>	
RepPoints: Point Set Representation for Object Detection	9656
<i>Ze Yang (Peking University), Shaohui Liu (Tsinghua University), Han Hu (Microsoft Research Asia), Liwei Wang (Peking University), and Stephen Lin (Microsoft Research)</i>	

SegEQA: Video Segmentation Based Visual Attention for Embodied Question Answering	9666
<i>Haonan Luo (Nanjing University of Science and Technology), Guosheng Lin (Nanyang Technological University), Zichuan Liu (Nanyang Technological University), Fayao Liu (Institute for Infocomm Research. ASTAR), Zhenmin Tang (Nanjing University of Science and Technology), and Yazhou Yao (Nanjing University of Science and Technology)</i>	
No-Frills Human-Object Interaction Detection: Factorization, Layout Encodings, and Training Techniques	9676
<i>Tanmay Gupta (UIUC), Alexander Schwing (UIUC), and Derek Hoiem (University of Illinois at Urbana-Champaign)</i>	
Cap2Det: Learning to Amplify Weak Caption Supervision for Object Detection	9685
<i>Keren Ye (University of Pittsburgh), Mingda Zhang (University of Pittsburgh), Adriana Kovashka (University of Pittsburgh), Wei Li (Facebook), Danfeng Qin (Google), and Jesse Berent (Google)</i>	
No Fear of the Dark: Image Retrieval Under Varying Illumination Conditions	9695
<i>Tomas Jenicek (CMP. FEE. Czech Technical University in Prague) and Ondrej Chum (Vision Recognition Group. Czech Technical University in Prague)</i>	
Hierarchical Shot Detector	9704
<i>Jiale Cao (Tianjin University), Yanwei Pang (Tianjin University), Jungong Han (Lancaster University), and Xuelong Li (Northwestern Polytechnical University)</i>	
Few-Shot Learning With Global Class Representations	9714
<i>Aoxue Li (Peking University), Tiange Luo (Peking University), Tao Xiang (University of Surrey), Weiran Huang (Noah's Ark Lab), and Liwei Wang (Peking University)</i>	
Better to Follow, Follow to Be Better: Towards Precise Supervision of Feature Super-Resolution for Small Object Detection	9724
<i>Junhyug Noh (Seoul National University), Wonho Bae (University of Massachusetts Amherst), Wonhee Lee (Seoul National University), Jinhwan Seo (Seoul National University), and Gunhee Kim (Seoul National University)</i>	
Weakly Supervised Object Detection With Segmentation Collaboration	9734
<i>Xiaoyan Li (University of the Chinese Academy of Sciences), Meina Kan (Institute of Computing Technology. Chinese Academy of Sciences), Shiguang Shan (Chinese Academy of Sciences), and Xilin Chen (Institute of Computing Technology. Chinese Academy of Sciences)</i>	
AutoFocus: Efficient Multi-Scale Inference	9744
<i>Mahyar Najibi (University of Maryland), Bharat Singh (Amazon), and Larry Davis (University of Maryland)</i>	
Leveraging Long-Range Temporal Relationships Between Proposals for Video Object Detection	9755
<i>Mykhailo Shvets (University of North Carolina at Chapel Hill), Wei Liu (Nuro Inc.), and Alexander Berg (University of North Carolina. USA)</i>	
Transferable Contrastive Network for Generalized Zero-Shot Learning	9764
<i>Huajie Jiang (ICT. CAS), Ruiping Wang (ICT. CAS), Shiguang Shan (Chinese Academy of Sciences), and Xilin Chen (Institute of Computing Technology. Chinese Academy of Sciences)</i>	

Fast Point R-CNN	9774
<i>Yilun Chen (Chinese University of Hong Kong), Shu Liu (Tencent), Xiaoyong Shen (Tencent), and Jiaya Jia (Chinese University of Hong Kong)</i>	
Mesh R-CNN	9784
<i>Georgia Gkioxari (Facebook), Justin Johnson (Facebook AI Research), and Jitendra Malik (University of California at Berkley)</i>	
Deep Supervised Hashing With Anchor Graph	9795
<i>Yudong Chen (Shenzhen University), Zhihui Lai (Shenzhen University), Yujuan Ding (the Hong Kong Polytechnic University), Kaiyi Lin (Peking University), and Waikeng Wong (Institute of Textiles and Clothing. The Hong Kong Polytechnic University)</i>	
Detecting 11K Classes: Large Scale Object Detection Without Fine-Grained Bounding Boxes	9804
<i>Hao Yang (Amazon), Hao Wu (Amazon), and Hao Chen (Amazon)</i>	
Re-ID Driven Localization Refinement for Person Search	9813
<i>Chuchu Han (HUST), Jiacheng Ye (PKU), Yunshan Zhong (Peking University), Xin Tan (Shanghai Jiao Tong University), Chi Zhang (Megvii Inc.), Changxin Gao (Huazhong University of Science and Technology), and Nong Sang (School of Automation. Huazhong University of Science and Technology)</i>	
Hierarchical Encoding of Sequential Data With Compact and Sub-Linear Storage Cost	9823
<i>Huu Le (Chalmers University of Technology), Ming Xu (University of New South Wales), Tuan Hoang (Singapore University of Technology and Design), and Michael Milford (ACRV and QUT. Australia)</i>	
C-MIDN: Coupled Multiple Instance Detection Network With Segmentation Guidance for Weakly Supervised Object Detection	9833
<i>Gao Yan (Institute Of Computing Technology Chinese Academy Of Sciences), Boxiao Liu (Institute of Computing Technology. Chinese Academy of Sciences), Nan Guo (institute of computing technology Chinese Academy of Sciences), Xiaochun Ye (Institute of Computing Technology Chinese Academy of Sciences), Fang Wan (University of Chinese Academy of Sciences), Haihang You (ICT. Chinese Academy of Sciences), and Dongrui Fan (ICT. Chinese Academy of Sciences)</i>	
Learning Feature-to-Feature Translator by Alternating Back-Propagation for Generative Zero-Shot Learning	9843
<i>Yizhe Zhu (Rutgers University), Jianwen Xie (Hikvision), Bingchen Liu (Rutgers. The State University of New Jersey), and Ahmed Elgammal (Rutgers University)</i>	
Deep Constrained Dominant Sets for Person Re-Identification	9854
<i>Leulseged Tesfaye Alemu (Ca' Foscari University of Venice), Mubarak Shah (University of Central Florida), and Marcello Pelillo (University of Venice)</i>	
Invariant Information Clustering for Unsupervised Image Classification and Segmentation	9864
<i>Xu Ji (Oxford University), Andrea Vedaldi (Oxford University), and Joao Henriques (University of Oxford)</i>	

Statistics, Physics, Theory & Datasets

Subspace Structure-Aware Spectral Clustering for Robust Subspace Clustering	9874
<i>Masataka Yamaguchi (NTT Corporation), Go Irie (NTT Communication Science Laboratories), Takahito Kawanishi (NTT Corporation), and Kunio Kashino (NTT Corporation)</i>	
Order-Preserving Wasserstein Discriminant Analysis	9884
<i>Bing Su (Institute of Software Chinese Academy of Sciences), Jiahuan Zhou (Northwestern University), and Ying Wu (Northwestern University)</i>	
LayoutVAE: Stochastic Scene Layout Generation From a Label Set	9894
<i>Akash Abdu Jyothi (Simon Fraser University), Thibaut Durand (Simon Fraser University), Jiawei He (Simon Fraser University), Leonid Sigal (University of British Columbia), and Greg Mori (Simon Fraser University)</i>	
Robust Variational Bayesian Point Set Registration	9904
<i>Jie Zhou (Yunnan Normal University), Xinke Ma (Yunnan Normal University), Li Liang (Yunnan Normal University), Liu Yuhe (312), Shijin Xu (Yunnan Normal University), Sim-Heng Ong (NUS), and Yang Yang (Yunnan Normal University)</i>	
Is an Affine Constraint Needed for Affine Subspace Clustering?	9914
<i>Chong You (University of California, Berkeley), Chun-Guang Li (Beijing University of Posts & Telecommunications), Daniel Robinson (Johns Hopkins University), and Rene Vidal (Johns Hopkins University)</i>	
Meta-Learning to Detect Rare Objects	9924
<i>Yu-Xiong Wang (Carnegie Mellon University), Deva Ramanan (Carnegie Mellon University), and Martial Hebert (Carnegie Mellon University)</i>	
New Convex Relaxations for MRF Inference With Unknown Graphs	9934
<i>Zhenhua Wang (Zhejiang University of Technology), Tong Liu (Zhejiang University of Technology), Qinfeng Shi (University of Adelaide), M. Pawan Kumar (University of Oxford), and Jianhua Zhang (Zhejiang University of Technology)</i>	
Cluster Alignment With a Teacher for Unsupervised Domain Adaptation	9943
<i>Zhijie Deng (Tsinghua University), Yucen Luo (Tsinghua University), and Jun Zhu (Tsinghua University)</i>	
Analyzing the Variety Loss in the Context of Probabilistic Trajectory Prediction	9953
<i>Luca Thiede (University of Göttingen) and Pratik Brahma (Volkswagen Electronics Research Lab)</i>	

3D From Single View & RGBD

Deep Mesh Reconstruction From Single RGB Images via Topology Modification Networks	9963
<i>Junyi Pan (South China University of Technology), Xiaoguang Han (Shenzhen Research Institute of Big Data, the Chinese University of Hong Kong, Shenzhen), Weikai Chen (USC Institute for Creative Technologies), Jiapeng Tang (South China University of Technology), and Kui Jia (South China University of Technology)</i>	

UprightNet: Geometry-Aware Camera Orientation Estimation From Single Images	9973
<i>Wenqi Xian (Cornell University), Zhengqi Li (Cornell University), Noah Snavely (Cornell University and Google AI), Matthew Fisher (Adobe Research), Jonathan Eisenman (Adobe Systems), and Eli Shechtman (Adobe Research. US)</i>	
Escaping Plato’s Cave: 3D Shape From Adversarial Rendering	9983
<i>Philipp Henzler (University College London), Niloy Mitra (University College London), and Tobias Ritschel (UCL)</i>	
Deep End-to-End Alignment and Refinement for Time-of-Flight RGB-D Module	9993
<i>Di Qiu (SenseTime Research), Jiahao Pang (SenseTime Research), Wenxiu Sun (SenseTime Research), and Chengxi Yang (SenseTime Group Limited)</i>	
GEOBIT: A Geodesic-Based Binary Descriptor Invariant to Non-Rigid Deformations for RGB-D Images ...	10003
<i>Erickson Nascimento (UFMG), Guilherme Potje (UFMG), Renato Martins (UFMG & INRIA), Felipe Chamone (UFMG), Mario Campos (UFMG), and Ruzena Bajcsy (UC Berkeley)</i>	
CDTB: A Color and Depth Visual Object Tracking Dataset and Benchmark	10012
<i>Alan Lukezic (University of Ljubljana), Ugur Kart (Tampere University of Technology), Jani Käpylä (Tampere University), Ahmed Durmush (Tampere University), Joni-Kristian Kamarainen (Tampere University), Jiri Matas (CMP CTU FEE), and Matej Kristan (University of Ljubljana)</i>	
Learning Joint 2D-3D Representations for Depth Completion	10022
<i>Yun Chen (Uber ATG), Bin Yang (Uber ATG & University of Toronto), Ming Liang (Uber ATG), and Raquel Urtasun (Uber ATG)</i>	

Face & Body

Make a Face: Towards Arbitrary High Fidelity Face Manipulation	10032
<i>Shengju Qian (Chinese University of Hong Kong), Kwan-Yee Lin (Peking university), Wayne Wu (Tsinghua University. SenseTime Research), Yangxiaokang Liu (University of Electronic Science and Technology of China), Quan Wang (Sensetime), Fumin Shen (UESTC), Chen Qian (SenseTime), and Ran He (Institute of Automation. Chinese Academy of Sciences)</i>	
M2FPA: A Multi-Yaw Multi-Pitch High-Quality Dataset and Benchmark for Facial Pose Analysis	10042
<i>Peipei Li (Institute of Automation Chinese Academy of Sciences), Xiang Wu (Institute of Automation. Chinese Academy of Science), Yibo Hu (Institute of Automation. Chinese Academy of Sciences), Ran He (Institute of Automation. Chinese Academy of Sciences), and Zhenan Sun (Chinese of Academy of Sciences)</i>	
Fair Loss: Margin-Aware Reinforcement Learning for Deep Face Recognition	10051
<i>Bingyu Liu (Beijing University of Posts and Telecommunications), Weihong Deng (Beijing University of Posts and Telecommunications), Yaoyao Zhong (Beijing University of Posts and Telecommunications), Mei Wang (Beijing University of Posts and Telecommunications), Jiani Hu (Beijing University of Posts and Telecommunications), Xunqiang Tao (Canon Information Technology. Beijing Co.. Ltd), and Yaohai Huang (Canon Information Technology. Beijing Co.. Ltd)</i>	

Face De-Occlusion Using 3D Morphable Model and Generative Adversarial Network	10061
<i>Xiaowei Yuan (Inha University) and In Kyu Park (Inha University)</i>	
Detecting Photoshopped Faces by Scripting Photoshop	10071
<i>Sheng-Yu Wang (University of California, Berkeley), Oliver Wang (Adobe Systems Inc), Richard Zhang (Adobe), Andrew Owens (UC Berkeley), and Alexei Efros (UC Berkeley)</i>	
Ego-Pose Estimation and Forecasting As Real-Time PD Control	10081
<i>Ye Yuan (Carnegie Mellon University) and Kris Kitani (CMU)</i>	
End-to-End Learning for Graph Decomposition	10092
<i>Jie Song (ETH Zurich), Bjoern Andres (University of Tübingen, Bosch Center for AI), Michael Black (Max Planck Institute for Intelligent Systems), Otmar Hilliges (ETH Zurich), and Siyu Tang (MPI for Intelligent Systems)</i>	
Laplace Landmark Localization	10102
<i>Joseph Robinson (Northeastern University), Yuncheng Li (Snap), Ning Zhang (UC Berkeley), Yun Fu (Northeastern University), and Sergey Tulyakov (Snap Inc)</i>	
Through-Wall Human Mesh Recovery Using Radio Signals	10112
<i>Mingmin Zhao (MIT), Yingcheng Liu (MIT), Aniruddh Raghu (MIT), Hang Zhao (MIT), Tianhong Li (MIT), Antonio Torralba (MIT), and Dina Katabi (Massachusetts Institute of Technology)</i>	
Discriminatively Learned Convex Models for Set Based Face Recognition	10122
<i>Hakan Cevikalp (Eskisehir Osmangazi University) and Golara Ghorban Dordinejad (E)</i>	
Camera Distance-Aware Top-Down Approach for 3D Multi-Person Pose Estimation From a Single RGB Image	10132
<i>Gyeongsik Moon (Seoul National University), Ju Yong Chang (Kwangwoon University), and Kyoung Mu Lee (Seoul National University)</i>	
Context-Aware Emotion Recognition Networks	10142
<i>Jiyoung Lee (Yonsei University), Seungryong Kim (École Polytechnique Fédérale de Lausanne, EPFL), Sunok Kim (Yonsei University), Jungin Park (Yonsei University), and Kwanghoon Sohn (Yonsei Univ.)</i>	
Aggregation via Separation: Boosting Facial Landmark Detector With Semi-Supervised Style Translation....	10152
<i>Shengju Qian (Chinese University of Hong Kong), Keqiang Sun (Tsinghua University), Wayne Wu (Tsinghua University, SenseTime Research), Chen Qian (SenseTime), and Jiaya Jia (Chinese University of Hong Kong)</i>	
Deep Head Pose Estimation Using Synthetic Images and Partial Adversarial Domain Adaption for Continuous Label Spaces	10163
<i>Felix Kuhnke (Leibniz Universität Hannover) and Joern Ostermann (Leibniz Universität Hannover)</i>	

Computational Photography & Graphics

Flare in Interference-Based Hyperspectral Cameras	10173
<i>Eden Sassoon (Technion), Tali Treibitz (University of Haifa), and Schechner Yoav (Technion)</i>	

Computational Hyperspectral Imaging Based on Dimension-Discriminative Low-Rank Tensor Recovery	10182
<i>Shipeng Zhang (Xi'an Jiaotong University), Lizhi Wang (Beijing Institute of Technology), Ying Fu (Beijing Institute of Technology), Xiaoming Zhong (Beijing Institute of Space Mechanics and Electricity), and Hua Huang (Beijing Institute of Technology)</i>	
Deep Optics for Monocular Depth Estimation and 3D Object Detection	10192
<i>Julie Chang (Stanford University) and Gordon Wetzstein (Stanford University)</i>	
Physics-Based Rendering for Improving Robustness to Rain	10202
<i>Shirsendu Halder (INRIA), Jean-Francois Lalonde (Université Laval), and Raoul De Charette (Inria)</i>	
ARGAN: Attentive Recurrent Generative Adversarial Network for Shadow Detection and Removal	10212
<i>Bin Ding (Wuhan University), Chengjiang Long (Kitware Inc), Ling Zhang (Wuhan University of Science and Technology), and Chunxia Xiao (Wuhan University)</i>	
Deep Tensor ADMM-Net for Snapshot Compressive Imaging	10222
<i>Jiawei Ma (Columbia University), Xiao-Yang Liu (Columbia University), Zheng Shou (Columbia University), and Xin Yuan (Bell Labs)</i>	

Low-Level & Optimization

Convex Relaxations for Consensus and Non-Minimal Problems in 3D Vision	10232
<i>Thomas Probst (ETH Zurich), Danda Pani Paudel (ETH Zürich), Ajad Chhatkuli (ETH Zurich), and Luc Van Gool (ETH Zurich)</i>	
Pareto Meets Huber: Efficiently Avoiding Poor Minima in Robust Estimation	10242
<i>Christopher Zach (Chalmers University) and Guillaume Bourmaud (University of Bordeaux)</i>	
K-Best Transformation Synchronization	10251
<i>Yifan Sun (UT Austin), Jiacheng Zhuo (The University of Texas at Austin), Arnav Mohan (Liberal Arts and Science Academy, LASA High School), and Qixing Huang (The University of Texas at Austin)</i>	
Parametric Majorization for Data-Driven Energy Minimization Methods	10261
<i>Jonas Geiping (University of Siegen) and Michael Moeller (University of Siegen)</i>	
A Bayesian Optimization Framework for Neural Network Compression	10273
<i>Xingchen Ma (KU Leuven), Amal Rannen Triki (KU Leuven), Maxim Berman (KU Leuven), Christos Sagonas (Onfido), Jacques Cali (Onfido), and Matthew Blaschko (KU Leuven)</i>	
HiPPI: Higher-Order Projected Power Iterations for Scalable Multi-Matching	10283
<i>Florian Bernard (Max Planck Institute for Informatics), Johan Thunberg (University of Luxembourg), Paul Swoboda (MPI fuer Informatik, Saarbruecken), and Christian Theobalt (MPI Informatik)</i>	

Language & Reasoning

Language-Conditioned Graph Networks for Relational Reasoning	10293
<i>Ronghang Hu (University of California, Berkeley), Anna Rohrbach (UC Berkeley), Trevor Darrell (UC Berkeley), and Kate Saenko (Boston University)</i>	
Tell, Draw, and Repeat: Generating and Modifying Images Based on Continual Linguistic Instruction	10303
<i>Alaaeldin El-Nouby (University of Guelph), Shikhar Sharma (Microsoft Research), Hannes Schulz (Microsoft), R Devon Hjelm (Microsoft Research), Layla El Asri (Microsoft), Samira Ebrahimi Kahou (McGill/Mila), Yoshua Bengio (Mila), and Graham Taylor (University of Guelph)</i>	
Relation-Aware Graph Attention Network for Visual Question Answering	10312
<i>Linjie Li (Microsoft), Zhe Gan (Microsoft), Yu Cheng (Microsoft), and Jingjing Liu (Microsoft)</i>	
Unpaired Image Captioning via Scene Graph Alignments	10322
<i>Jiuxiang Gu (Nanyang Technological University), Shafiq Joty (Nanyang Technological University), Jianfei Cai (Nanyang Technological University), Handong Zhao (Adobe Research), Xu Yang (Nanyang Technological University), and Gang Wang (Alibaba Group)</i>	
Modeling Inter and Intra-Class Relations in the Triplet Loss for Zero-Shot Learning	10332
<i>Yannick Le Cacheux (CEA LIST), Herve Le Borgne (CEA LIST), and Michel Crucianu (CEDRIC LAB. CNAM)</i>	
Occlusion-Shared and Feature-Separated Network for Occlusion Relationship Reasoning	10342
<i>Rui Lu (Beijing University of Posts and Telecommunications), Feng Xue (Beijing University of Posts and Telecommunications), Menghan Zhou (Beijing University of Posts and Telecommunications), Anlong Ming (Beijing University of Posts and Telecommunications), and Yu Zhou (Huazhong University of Science and Technology)</i>	
Compositional Video Prediction	10352
<i>Yufei Ye (Carnegie Mellon University), Maneesh Singh (Verisk Analytics), Abhinav Gupta (CMU/FAIR), and Shubham Tulsiani (Facebook AI Research)</i>	
Mixture-Kernel Graph Attention Network for Situation Recognition	10362
<i>Mohammed Suhail (University of British Columbia) and Leonid Sigal (University of British Columbia)</i>	
Learning Similarity Conditions Without Explicit Supervision	10372
<i>Reuben Tan (Boston University), Mariya Vasileva (University of Illinois at Urbana-Champaign), Kate Saenko (Boston University), and Bryan Plummer (Boston University)</i>	
Joint Prediction for Kinematic Trajectories in Vehicle-Pedestrian-Mixed Scenes	10382
<i>Huikun Bi (Institute of Computing Technology Chinese Academy of Sciences), Zhong Fang (Institute of Computing Technology Chinese Academy of Sciences), Tianlu Mao (Institute of Computing Technology Chinese Academy of Sciences), Zhaoqi Wang (Chinese Academy of Sciences), and Zhigang Deng (University of Houston)</i>	

Learning to Caption Images Through a Lifetime by Asking Questions	10392
<i>Tingke Shen (University of Toronto), Amlan Kar (University of Toronto), and Sanja Fidler (University of Toronto. NVIDIA)</i>	
VrR-VG: Refocusing Visually-Relevant Relationships	10402
<i>Yuanzhi Liang (Xi'an Jiaotong university), Yalong Bai (JD AI Research), Wei Zhang (JD AI Research), Xueming Qian (Xi'an Jiaotong University), Li Zhu (Xi'an Jiaotong University), and Tao Mei (AI Research of JD.com)</i>	

3D From Multiview & Sensors

TAPA-MVS: Textureless-Aware PAtchMatch Multi-View Stereo	10412
<i>Andrea Romanoni (Politecnico di Milano) and Matteo Matteucci (Politecnico di Milano)</i>	
U4D: Unsupervised 4D Dynamic Scene Understanding	10422
<i>Armin Mustafa (University of Surrey), Chris Russell (The Alan Turing Institute/University of Surrey), and Adrian Hilton (University of Surrey)</i>	
Hierarchical Point-Edge Interaction Network for Point Cloud Semantic Segmentation	10432
<i>Li Jiang (The Chinese University of Hong Kong), Hengshuang Zhao (The Chinese University of Hong Kong), Shu Liu (Tencent), Xiaoyong Shen (Tencent), Chi-Wing Fu (The Chinese University of Hong Kong), and Jiaya Jia (Chinese University of Hong Kong)</i>	
Multi-Angle Point Cloud-VAE: Unsupervised Feature Learning for 3D Point Clouds From Multiple Angles by Joint Self-Reconstruction and Half-to-Half Prediction	10441
<i>Zhizhong Han (University of Maryland. College Park), Xiyang Wang (Tsinghua University), Yu-Shen Liu (Tsinghua University), and Matthias Zwicker (University of Maryland)</i>	
P-MVSNet: Learning Patch-Wise Matching Confidence Aggregation for Multi-View Stereo	10451
<i>Keyang Luo (Huazhong University of Science and Technology), Tao Guan (Huazhong University of Science and Technology), Lili Ju (University of South Carolina), Haipeng Huang (Farsee2 Science and Technology Ltd), and Yawei Luo (University of Technology Sydney)</i>	

Image & Video Synthesis

SME-Net: Sparse Motion Estimation for Parametric Video Prediction Through Reinforcement Learning	10461
<i>Yung-Han Ho (NCTU), Chuan-Yuan Cho (NCTU), Guo-Lun Jin (NCTU), and Wen-Hsiao Peng (National Chiao Tung University)</i>	
ClothFlow: A Flow-Based Model for Clothed Person Generation	10470
<i>Xintong Han (Malong Technologies), Weilin Huang (Malong Technologies), Xiaojun Hu (Malong Technologies), and Matthew Scott (Malong Technologies)</i>	

LADN: Local Adversarial Disentangling Network for Facial Makeup and De-Makeup	10480
<i>Qiao Gu (Hong Kong University of Science and Technology), Guanzhi Wang (Hong Kong University of Science and Technology), Mang Tik Chiu (UIUC), Yu-Wing Tai (Tencent), and Chi-Keung Tang (Hong Kong University of Science and Technology)</i>	
Point-to-Point Video Generation	10490
<i>Tsun-Hsuan Wang (National Tsing Hua University), Yen-Chi Cheng (National Tsing Hua University), Chieh Hubert Lin (National Tsing Hua University), Hwann-Tzong Chen (National Tsing Hua University), and Min Sun (NTHU)</i>	
Semantics-Enhanced Adversarial Nets for Text-to-Image Synthesis	10500
<i>Hongchen Tan (Dalian University of Technology), Xiuping Liu (Dalian University of Technology), Xin Li (Louisiana State University), Yi Zhang (Dalian University of Technology), and Baocai Yin (Dalian University of Technology)</i>	
VTNFP: An Image-Based Virtual Try-On Network With Body and Clothing Feature Preservation	10510
<i>Ruiyun Yu (Northeastern University), Xiaoqi Wang (Northeastern University), and Xiaohui Xie (University of California, Irvine)</i>	
Boundless: Generative Adversarial Networks for Image Extension	10520
<i>Dilip Krishnan (Google), Piotr Teterwak (Google), Aaron Sarna (Google), Aaron Maschinot (Google Research), Ce Liu (Google), David Belanger (Google), and William Freeman (Google)</i>	
Image Synthesis From Reconfigurable Layout and Style	10530
<i>Wei Sun (NC state university) and Tianfu Wu (NC State University)</i>	
Attribute Manipulation Generative Adversarial Networks for Fashion Images	10540
<i>Kenan Ak (National University of Singapore), Ashraf Kassim (National University of Singapore, Singapore), Joo-Hwee Lim (Institute for Infocomm Research), and Jo Yew Tham (ESP xMedia Pte. Ltd.)</i>	
Few-Shot Unsupervised Image-to-Image Translation	10550
<i>Ming-Yu Liu (NVIDIA), Xun Huang (Cornell University), Arun Mallya (NVIDIA), Tero Karras (NVIDIA), Timo Aila (NVIDIA), Jaakko Lehtinen (NVIDIA & Aalto University), and Jan Kautz (NVIDIA)</i>	
Very Long Natural Scenery Image Prediction by Outpainting	10560
<i>Zongxin Yang (UTS; Qihoo/360), Jian Dong (Qihoo/360), Ping Liu (UTS), Yi Yang (UTS), and Shuicheng Yan (Qihoo/360)</i>	

Applications, Medical & Robotics

Scaling Recurrent Models via Orthogonal Approximations in Tensor Trains	10570
<i>Ronak Mehta (University of Wisconsin-Madison), Rudrasis Chakraborty (University of California, Berkeley), Vikas Singh (University of Wisconsin-Madison USA), and Yunyang Xiong (University of Wisconsin-Madison)</i>	
A Deep Cybersickness Predictor Based on Brain Signal Analysis for Virtual Reality Contents	10579
<i>Jinwoo Kim (Yonsei University), Woojae Kim (Yonsei University), Heeseok Oh (Electronics & Telecommunications Research Institute), Seongmin Lee (Yonsei University), and Sanghoon Lee (Yonsei University, Korea)</i>	

Learning With Unsure Data for Medical Image Diagnosis	10589
<i>Botong Wu (Peking University), Xinwei Sun (Peking University), Lingjing Hu (Capital Medical University), and Yizhou Wang (PKU)</i>	
Recursive Cascaded Networks for Unsupervised Medical Image Registration	10599
<i>Shengyu Zhao (Tsinghua University), Yue Dong (Tsinghua University), Eric Chang (Microsoft Asia), and Yan Xu (Beihang University)</i>	
DUAL-GLOW: Conditional Flow-Based Generative Model for Modality Transfer	10610
<i>Haoliang Sun (Shandong University), Ronak Mehta (University of Wisconsin-Madison), Hao Zhou (University of Wisconsin Madison), Zhichun Huang (University of Wisconsin-Madison), Sterling Johnson (University of Wisconsin-Madison), Vivek Prabhakaran (University of Wisconsin-Madison), and Vikas Singh (University of Wisconsin-Madison USA)</i>	
Dilated Convolutional Neural Networks for Sequential Manifold-Valued Data	10620
<i>Rudrasis Chakraborty (University of California, Berkeley), Xingjian Zhen (University of Wisconsin-Madison), Nicholas Vogt (University of Wisconsin-Madison), Barbara Bendlin (University of Wisconsin-Madison), and Vikas Singh (University of Wisconsin-Madison USA)</i>	
Align, Attend and Locate: Chest X-Ray Diagnosis via Contrast Induced Attention Network With Limited Supervision	10631
<i>Jingyu Liu (Peking University), Gangming Zhao (Deepwise AI Lab), Yu Fei (Peking University), Ming Zhang (Peking University, China), Yizhou Wang (PKU), and Yizhou Yu (Deepwise AI Lab)</i>	
Joint Acne Image Grading and Counting via Label Distribution Learning	10641
<i>Xiaoping Wu (Nankai University), Ni Wen (Beijing Tsinghua Changgung Hospital), Jie Liang (Nankai University), Yu-Kun Lai (Cardiff University), Dongyu She (Nankai University), Ming-Ming Cheng (Nankai University), and Jufeng Yang (Nankai University)</i>	
An Alarm System for Segmentation Algorithm Based on Shape Model	10651
<i>Fengze Liu (Johns Hopkins University), Yingda Xia (Johns Hopkins University), Dong Yang (NVIDIA Corporation), Alan Yuille (Johns Hopkins University), and Daguang Xu (NVIDIA Corporation)</i>	
HistoSegNet: Semantic Segmentation of Histological Tissue Type in Whole Slide Images	10661
<i>Lyndon Chan (University of Toronto), Mahdi Hosseini (University of Toronto), Corwyn Rowsell (st michael hospital), Konstantinos Plataniotis (UofT), and Savvas Damaskinos (Huron Digital Pathology)</i>	
Prior-Aware Neural Network for Partially-Supervised Multi-Organ Segmentation	10671
<i>Yuyin Zhou (Johns Hopkins University), Zhe Li (Google), Song Bai (University of Oxford), Xinlei Chen (Facebook AI Research), Mei Han (paii-labs.com), Chong Wang (ByteDance Inc.), Elliot Fishman (JHMI), and Alan Yuille (Johns Hopkins University)</i>	

CAMEL: A Weakly Supervised Learning Framework for Histopathology Image Segmentation	10681
<i>Gang Xu (School of Life Sciences. Tsinghua University), Zhigang Song (Department of Pathology. The Chinese PLA General Hospital), Zhuo Sun (Thorough Images), Calvin Ku (Thorough Images), Zhe Yang (Department of Physics. Tsinghua University), Cancheng Liu (Thorough Images), Shuhao Wang (Thorough Images), Jianpeng Ma (Multiscale Research Institute for Complex Systems. Fudan University), and Wei Xu (Institute for Interdisciplinary Information Sciences. Tsinghua University)</i>	
Conditional Recurrent Flow: Conditional Generation of Longitudinal Samples With Applications to Neuroimaging	10691
<i>Seong Jae Hwang (University of Wisconsin-Madison), Zirui Tao (University of Wisconsin-Madison), Vikas Singh (University of Wisconsin-Madison USA), and Won Hwa Kim (University of Texas at Arlington)</i>	
Multi-Stage Pathological Image Classification Using Semantic Segmentation	10701
<i>Shusuke Takahama (The University of Tokyo), Yusuke Kurose (The University of Tokyo), Yusuke Mukuta (The University of Tokyo), Hiroyuki Abe (The University of Tokyo), Masashi Fukayama (The University of Tokyo), Akihiko Yoshizawa (Kyoto University), Masanobu Kitagawa (Tokyo Medical and Dental University), and Tatsuya Harada (The University of Tokyo / RIKEN)</i>	
Semantic-Transferable Weakly-Supervised Endoscopic Lesions Segmentation	10711
<i>Jiahua Dong (Chinese Academy of Sciences), Yang Cong (Chinese Academy of Sciences), Gan Sun (State Key Laboratory of Robotics. Shenyang Institute of Automation. Institutes for Robotics and Intelligent Manufacturing. Chinese Academy of Sciences), and Dongdong Hou (University of Chinese Academy of Sciences)</i>	
Unsupervised Microvascular Image Segmentation Using an Active Contours Mimicking Neural Network	10721
<i>Shir Gur (Tel Aviv University), Lior Wolf (Tel Aviv University. Israel), Lior Golgher (Tel Aviv University), and Pablo Blinder (Tel Aviv University)</i>	
GLAMpoints: Greedily Learned Accurate Match Points	10731
<i>Prune Truong (RetinAI Medical AG / ETHZ), Stefanos Apostolopoulos (RetinAI Medical AG), Agata Mosinska (RetinAI Medical), Samuel Stucky (RetinAI Medical), Carlos Ciller (RetinAI Medical), and Sandro De Zanet (RetinAI Medical)</i>	
Author Index	10741