IEEE/CVF Conference on

Computer Vision and Pattern Recognition

Pocket Guide
(Main Conference)

CVPR

June 18 – 22, 2018

Salt Lake City, Utah

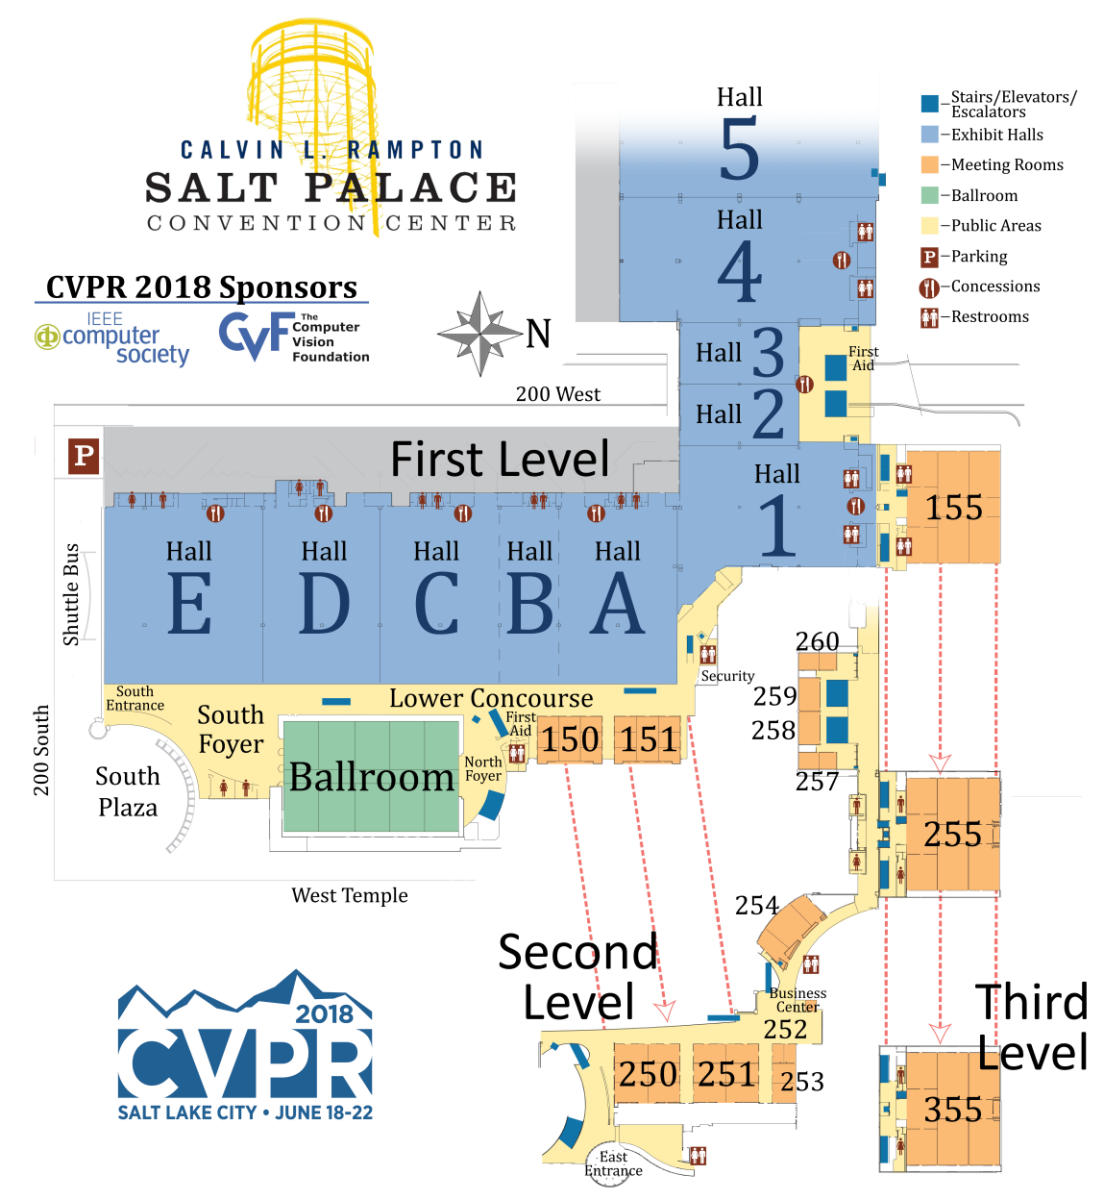
CALVIN L. RAMPTON
SALT PALACE
CONVENTION CENTER
CVPR 2018 Sponsors
IEEE computer society
CVF The Computer Vision Foundation
N
Stairs/Elevators/ Escalators
Exhibit Halls
Meeting Rooms
Ballroom
Public Areas
Parking
Concessions
Restrooms
Hall 5
Hall 4
Hall 3
Hall 2
Hall 1
Hall E
Hall D
Hall C
Hall B
Hall A
First Level
200 West
First Aid
155
260
259
258
257
255
254
Security
Shuttle Bus
200 South
South Entrance
South Foyer
South Plaza
Lower Concourse
Ballroom
North Foyer
First Aid
150
151
West Temple
Second Level
Business Center
252
250
251
253
East Entrance
Third Level
355
2018 CVPR
SALT LAKE CITY • JUNE 18-22

Message from the General and Program Chairs

Welcome to Salt Lake City and the 31st meeting of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2018). CVPR continues to be a showcase for some of the most exciting advances in computer vision, pattern recognition, machine learning, and artificial intelligence. With a full slate of oral presentations, poster sessions, workshops, tutorials, demos, social functions, and our growing industrial exhibition, this promises to be a week with something for everyone.

As our field continues to grow, so too do our conferences. CVPR 2014 first passed 2,000 attendees, and within three years CVPR 2017 had grown to nearly 5,000, with this year's meeting eclipsing that mark. As attendance has grown, so too has the number of submitted papers. This year's meeting received 3,359 submissions, of which 3,309 were fully reviewed after a handful of administrative rejections. To handle the reviewing process for that many papers, we invited 108 Area Chairs (ACs). This meant including many up-and-coming researchers who had not previously served as ACs and who were paired with more experienced ones. We recruited a pool of 2,385 reviewers, out of which 1,715 were invited to perform reviews. Of these, 575 reviewed the maximum number (10) of assigned papers. One hundred and sixty emergency reviewers graciously provided their time, usually on short notice.

The review process was similar to that used in previous years. Each paper was reviewed by at least three reviewers and then considered further by at least three ACs before a decision was made. Borderline papers and candidates for orals and spotlights were discussed in groups of three non-conflicted ACs. Oral and spotlight recommendations were then made by panels of ACs. Program Chairs did not submit papers, which allowed them to avoid conflicts throughout the process.

Past attendees know that as CVPR has grown it has become harder to see all of the amazing content one might hope to. It is tempting to try to rein in the growth of the conference program accordingly, but we feel strongly that maintaining the health of the research community is our highest priority. Paper decisions were made with the same high standards we have always had without trying to artificially cap the total number of accepted papers. This meant accepting 979 papers, 25% more than in 2017 and 52% more than in 2016. These include 70 papers selected for oral presentation (2.1%), 224 papers selected for spotlight oral presentation (6.6%), and 685 papers selected for poster presentation (20.4%). Overall, 29% of submitted papers were accepted for publication, with 30% of accepted papers selected for some form of oral presentation. As in recent years, all accepted papers appear in the poster sessions.

Prior to 2014, CVPR used a three-day, two-track format. As we've grown, the organizers of recent meetings have experimented with a variety of formats, including keeping the same format while including more papers (2015), adding a fourth day (2014, 2016), adding a fourth day with a half-day break during the meeting (2017), and adding a third parallel track (2017). In each case attendee feedback was gathered and used to guide future meetings. However, with our three-year (now four-year) planning window, some decisions had to be locked in years in advance. We made the decision back in 2015 for CVPR 2018 to be a three-day meeting, but the abundance of space in the facility makes it possible to again hold three parallel oral presentation tracks as well as to comfortably grow our space for posters alongside our industrial exhibition.

To accommodate both more attendees and more papers, we have experimented with a few changes this year. First, we made the decision to invite only one keynote speaker and to use the time for additional oral or poster sessions. When that speaker had to withdraw for family reasons, we went ahead and used that time similarly. As a result, each day of the main program includes one more poster or oral session than in previous years. Second, to continue to provide lunch to a growing number of attendees, we have extended the period during which lunch will be served and similarly extended the poster sessions immediately before and after. This overlap between extended lunch and extended poster sessions hopefully means greater flexibility and less crowding. Third, the exhibition will be open continuously from 10:00 a.m. to 6:30 p.m., allowing greater flexibility as well. Finally, we are

Message from the General and Program Chairs

experimenting with offering a few tutorials during the main program days in parallel with the paper sessions, providing more options for newer students and others attending the conference.

Our industrial exhibition also continues to grow, and we are grateful for the participation of our sponsors and other exhibitors. This exhibition provides a great opportunity for researchers, students, and other attendees to connect with companies ranging from promising startups to industry leaders. Exhibitor fees, plus generous donations from our corporate sponsors, allow us to improve the attendee experience while keeping CVPR the incredible value it has always been.

Two other changes are also worth noting: First, we have instituted a new "Code of Conduct", which can be found on the conference website. We trust that everyone involved in CVPR will continue to maintain the high degree of professionalism and collegiality we have always enjoyed here and hope this addition will be unnecessary, but we also feel it is important to be clear about our community expectations. Second, we have added a "Good Citizen of CVPR" panel on Friday for all who are new to our community and want to increase and improve their involvement.

Finally, we would like to thank everyone involved in making CVPR 2018 a success. This includes the organizing committee, the area chairs, the reviewers, emergency reviewers, authors, demo session participants, donors, exhibitors, and everyone else without whom this meeting would not be possible. We are grateful to Hall-Erickson International, Freeman, Presentation Technologies, and other vendors for their helpful coordination and experience. We are also grateful that Eric Mortensen, who has for many years managed the CVPR publication process, graciously agreed to reprise his role as publications chair one more time for us. A huge thank-you also goes to Nicole Finn, Liz Ryan, and C to C Events for their organization of the logistics of the conference.

Most of all, we thank all of you for attending CVPR and making it one of the top venues for computer vision research in the world. We hope that your travel plans permit you time before or after the conference to get out from the city and to enjoy the rich natural diversity of the mountains, lakes, rivers, deserts, and National Parks of Utah. Enjoy CVPR 2018!

Program Chairs: **David Forsyth, Ivan Laptev, Aude Oliva, and Deva Ramanan**

General Chairs: **Michael S. Brown, Bryan Morse, and Shmuel Peleg**

QR Codes for the IEEE CPS Mobile App

Apple iTunes App Store (iOS)

Google Play (Android)

CVPR 2018 Organizing Committee

General Chairs:	Michael S. Brown Bryan Morse Shmuel Peleg
Program Chairs:	David Forsyth Ivan Laptev Deva Ramanan Aude Oliva
Workshops Chairs:	Srikumar Ramalingam Mathieu Salzmann
Tutorials Chairs:	M. Pawan Kumar Andrea Vedaldi
Finance Chairs:	Walter Scheirer Ramin Zabih
Publications Chairs:	Eric Mortensen William Brendel
Demos Chair:	Seon Joo Kim
Website Chairs:	Dana Berman Abdelrahman Kamel
Corporate Relations Chairs:	Rogerio Feris Scott McCloskey Yu Wing Tai
Doctoral Consortium Chairs:	Peter Carr Xiaoming Liu
Student Volunteers Chairs:	Danna Gurari Marcus Brubaker
Local Arrangements Chair:	Ryan Farrell

CVPR 2018 Area Chairs

Sameer Agarwal
Aseem Agrawala
Karteek Alahari
Mathieu Aubry
Shai Avidan
Kavita Bala
Jon Barron
Kobus Barnard
Dhruv Batra
Tamara Berg
Alexander C Berg
Margrit Betke
Horst Bischof
Ayan Chakrabarti
Shih-Fu Chang
Tat-Jun Chin
Minsu Cho
Robert Collins
Daniel Cremers
Kristin Dana
Fernando de la Torre
Jia Deng
Alexey Dosovitskiy
Alexei (Alyosha) Efros
Ahmed Elgammal
Irfan Essa
Ali Farhadi
Vittorio Ferrari
Sanja Fidler
Katerina Fragkiadaki
Pascal Fua
Yasutaka Furukawa
Andreas Geiger
Bernard Ghanem
Ross Girshick
Georgia Gkioxari
Abhinav Gupta
Kaiming He
Wolfgang Heidrich
Aaron Hertzmann
Derek Hoiem
Anthony Hoogs
Nazli Ikizler-Cinbis
Nathan Jacobs
Hervé Jégou
Neel Joshi
Ira Kemelmacher-Shlizerman
Junmo Kim
Seon Joo Kim
Kris Kitani
Iasonas Kokkinos
Vladlen Koltun
Nikos Komodakis
Adriana Kovashka
Philipp Krähenbühl
Pawan Kumar
Kyros Kutulakos
Simon Lacoste-Julien
Christoph Lampert
Svetlana Lazebnik
Bastian Leibe
Victor Lempitsky
Zicheng Liao
Simon Lucey
Michael Maire
Subhransu Maji
B.S. Manjunath
R. Manmatha
Jiri Matas
Greg Mori
Srinivasa Narasimhan
Ram Nevatia
Juan Carlos Niebles
Ko Nishino
Sebastian Nowozin
Tomas Pajdla
Devi Parikh
Hamed Pirsiavash
Srikumar Ramalingam
Ian Reid
Xiaofeng Ren
Stefan Roth
Olga Russakovsky
Kate Saenko
Pinar Duygulu Sahin
Mathieu Salzmann
Dimitris Samaras
Torsten Sattler
Silvio Savarese
Alex Schwing
Stan Sclaroff
Thomas Serre
Greg Shakhnarovich
Yaser Sheikh
Abhinav Shrivastava
Josef Sivic
Cristian Sminchisescu
Noah Snavely
Richard Souvenir
Ganesh Sundaramoorthi
Ping Tan
Camillo J. Taylor
Carl Vondrick
Gang Wang
Gordon Wetzstein
Ming-Hsuan Yang
Lihi Zelnik-Manor
Todd Zickler

CVPR 2018 Outstanding Reviewers

We are pleased to recognize the following researchers as "CVPR 2018 Outstanding Reviewers". These reviewers were identified by one or more of the CVPR Area Chairs for their hard work in providing high quality and detailed reviews for their assigned papers.

Pulkit Agrawal
Pablo Arbelaez
Artem Babenko
Leonard Berrada
Gedas Bertasius
Adel Bibi
Piotr Bojanowski
Shyamal Buch
Rudy Bunel
Zoya Bylinskii
Cesar Caden
Peter Carr
Visesh Chari
Liang-Chieh Chen
Ming-Ming Cheng
Christopher Choy
Andrea Cohen
David Crandall
Dengxin Dai
Abir Das
Aditya Deshpande
Ferran Diego
Thanh-Toan Do
Simon Donné
Victor Escorcia
Chen Fang
Ryan Farrell
Chen Feng
Michael Figurnov
Michael Firman
Ruth Fong
David Fouhey
Thomas Funkhouser
Adrien Gaidon
Juergen Gall
Silvano Galliani
Orazio Gallo
Spyros Gidaris
Clement Godard
Yash Goyal
Alp Guler

Jinwei Gu
Saurabh Gupta
Fatma Güney
Kai Han
Bharath Hariharan
Lisa Anne Hendricks
Le Hou
Ronghang Hu
Wei-Chih Hung
Junhwa Hur
Christian Haene
Slobodan Ilic
Varun Jampani
Kushal Kafle
Vicky Kalogeiton
Evangelos Kalogerakis
Abhishek Kar
Svebor Karaman
Jan Kautz
Kihwan Kim
Alexander Kirillov
Nicholas Kolkin
Abhijit Kundu
Kaustav Kundu
Suha Kwak
Wei-Sheng Lai
Jean-Francois Lalonde
John Lambert
Diane Larlus
Katrin Lasinger
Christoph Lassner
Dangwei Li
Xin Li
Yijun Li
Weixin Li
Jie Liang
Tsung-Yu Lin
Ming-Yu Liu
Sifei Liu
Xianglong Liu
Yun Liu

Le Lu
Chao Ma
Oisin Mac Aodha
Clement Mallet
Arun Mallya
Julia Martinez
Gellert Mattyus
Nikolaus Mayer
Roey Mechrez
Lars Mescheder
Pavlo Molchanov
Philippos Mordohai
Matthias Mueller
Mahyar Najibi
David Nilsson
Jean-Marc Odobez
Mohamed Omran
Aleksis Pirinen
Bryan Plummer
Tobias Plötz
Matteo Poggi
Alin Popa
Andrea Prati
Maria Priisalu
Srikumar Ramalingam
Adria Recasens
Timo Rehfeld
Stephan Richter
Gernot Riegler
Jason Rock
Anna Rohrbach
Samuel Rota Bulò
Bryan Russell
Manolis Savva
Yoav Schechner
Bernt Schiele
Konrad Schindler
Johannes Schönberger
Pradeep Sen
Eli Shechtman
Vikas Singh

Kihyuk Sohn
Francesco Solera
Shuran Song
Yibing Song
Pratul Srinivasan
Julian Straub
Hang Su
Jong-Chyi Su
Chong Sun
Deqing Sun
Libin Sun
Yuichi Taguchi
Tatsunori Taniai
Lucas Theis
Giorgos Tolias
Federico Tombari
Yi-Hsuan Tsai
Zhuowen Tu
Shubham Tulsiani
Osman Ulusoy
Laurens van der Maaten
Christoph Vogel
Yalin Wang
Anne Wannenwetsch
Jan Wegner
Jiajun Wu
Tianfan Xue
Raymond Yeh
Fisher Yu
Hong-Xing Yu
Christopher Zach
Andrei Zanfir
Mihai Zanfir
Tianzhu Zhang
Yinda Zhang
Kun Zhou
Tinghui Zhou
Ji Zhu
Andrew Zisserman
Daniel Zoran
Laurens van der Maaten

CVPR 2018 Emergency Reviewers

We also want to recognize the following researchers as "CVPR 2018 Emergency Reviewers". These reviewers were willing to provide an "emergency" review on short notice within a very short time frame. Thank you for your service.

Pablo Arbelaez
Carlos Arteta
Mahsa Baktashmotlagh
Aayush Bansal
Pierre Baqué
Sarah Bargal
Boulbaba Benamor
Archith Bency
Lucas Beyer
Adel Bibi
Hakan Bilen
Vishnu Naresh Boddeti
Piotr Bojanowski
Lisa Brown
Marcus Brubaker
Rudy Bunel
Michal Bušta
Zoya Bylinskii
Visesh Chari
Jianhui Chen
Liang-Chieh Chen
Christopher Choy
Andrea Cohen
Dengxin Dai
Abir Das
Achal Dave
Shalini De Mello
Aditya Deshpande
Matthijs Douze
Mohamed Elhoseiny
Francis Engelmann
Aykut Erdem
Quanfu Fan
Chen Feng
Christopher Funk
Matheus Gadelha
Juergen Gall
Hamed Galoogahi
Ashwinkumar Ganesan
Utkarsh Gaur
Oleg Grinchuk
Alp Guler
Tanmay Gupta
Bharath Hariharan
Adam Harley
James Hays
Felix Heide
David Held
Lisa Anne Hendricks
Alexander Hermans
Minh Hoai
Hexiang Hu
Shell Hu
Dong Huang
Sheng Huang
Younghyun Jo
Shantanu Joshi
Naeemullah Khan
Jaechul Kim
Alexander Kolesnikov
Nicholas Kolkin
Chen Kong
Hilde Kuehne
Sebastian Kurtek
Ilja Kuzborskij
Tien Lan
Dong Lao
Longin Jan Latecki
Jose Lezama
Jie Liang
Tsung-Yu Lin
Feng Liu
Ming-Yu Liu
Xianglong Liu
Xiaoming Liu
Yebin Liu
Yadong MU
Chao Ma
Pablo Marquez Neila
Mohammadreza Mostajabi
Matthias Mueller
Calvin Murdock
Ana Murillo
Lakshmanan Nataraj
Thuyen Ngo
Erfan Noury
Michael Opitz
Maxime Oquab
Roy Or-El
Hyun Soo Park
Omkar Parkhi
Deepak Pathak
Giorgio Patrini
Ramtin Pedarsani
Selen Pehlivan
Anastasia Pentina
Robert Pless
Georg Poier
Horst Possegger
Zhen Qian
Chao Qu
Kandan Ramakrishnan
Timo Rehfeld
Tal Remez
Wenqi Ren
Jason Rock
Peter Roth
Aruni RoyChowdhury
Amelie Royer
Artem Rozantsev
Bryan Russell
Alexandre Sablayroll
Hideo Saito
Elham Saraee
Pradeep Sen
Sohil Shah
Gaurav Sharma
Mohit Sharma
Chunhua Shen
Miaojing Shi
Takashi Shibata
Gunnar Sigurdsson
Bharat Singh
Saurabh Singh
Cees Snoek
Pratul Srinivasan
Hang Su
Jong-Chyi Su
Shuochen Su
Chen Sun
Chong Sun
Deqing Sun
Libin Sun
Ying Sun
Xin Tao
Mehmet Tek
Pavel Tokmakov
Matthew Trager
Quoc-Huy Tran
Son Tran
Oytun Ulutan
Paul Upchurch
Ben Usman
Patricio Vela
Minh Vo
Jayakorn Vongkulbhisal
Chaoyang Wang
Zijun Wei
Hui Wu
Lu Xia
Tao Xu
Tianfan Xue
Jimei Yang
Zhenheng Yang
Raymond Yeh
Christopher Zach
Sergey Zagoruyko
Ji Zhu
Rui Zhu
Laurens van der Maaten

Tuesday, June 19

0730–1830 Registration (South Lobby)

0730–0900 Breakfast (Hall A; Halls 1-4)

0800–1000 Poster Setup (Halls C-E)

0830–0850 Opening Remarks & Paper Awards (Ballroom)

0850-0940 Special Session: Workshop Competitions (Ballroom)

0850 Introduction, *Ramin Zabih*

0855 The Role of Competitions in Machine Learning, *Anthony Goldbloom*

0915 Selected CVPR18 Workshop Competitions

Format (5 min. for presentation; no questions)

1. The 2018 NVIDIA AI City Challenge, *Milind Naphade*
2. iNaturalist + iMaterialist 2018 (FGVC5), *Ryan Farrell*
3. Autonomous Driving Perception Challenges, *Fisher Yu*
4. The Visual Question Answering Challenge, *Yash Goyal*

0940-1010 Session 1-1A: Object Recognition & Scene Understanding I (Ballroom)

Papers in this session are also in Poster Session P1-1. Poster tag in square brackets (e.g., [***A1***])

Chairs: Ali Farhadi *(Univ. of Washington)*
Tamara Berg *(UNC Chapel Hill)*

0940 Orals (O1-1A)

Format (12 min. for presentation + 2 min. for questions)

1. [***A1***] Embodied Question Answering, *Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, Dhruv Batra*
2. [***A4***] Learning by Asking Questions, *Ishan Misra, Ross Girshick, Rob Fergus, Martial Hebert, Abhinav Gupta, Laurens van der Maaten*

0850-1010 Session 1-1B: Analyzing Humans in Images I (Room 155)

Papers in this session are also in Poster Session P1-1. Poster tag in square brackets (e.g., [***A7***])

Chairs: Fernando de la Torre *(Carnegie Mellon Univ.)*
Ira Kemelmacher *(Univ. of Washington)*

0850 Orals (O1-1B)

Format (12 min. for presentation + 2 min. for questions)

1. [***A7***] Finding Tiny Faces in the Wild With Generative Adversarial Network, *Yancheng Bai, Yongqiang Zhang, Mingli Ding, Bernard Ghanem*
2. [***A10***] Learning Face Age Progression: A Pyramid Architecture of GANs, *Hongyu Yang, Di Huang, Yunhong Wang, Anil K. Jain*
3. [***B2***] PairedCycleGAN: Asymmetric Style Transfer for Applying and Removing Makeup, *Huiwen Chang, Jingwan Lu, Fisher Yu, Adam Finkelstein*

0934 Spotlights (S1-1B)

Format (4 min. for presentation; no questions)

1. [***B5***] GANerated Hands for Real-Time 3D Hand Tracking From Monocular RGB, *Franziska Mueller, Florian Bernard, Oleksandr Sotnychenko, Dushyant Mehta, Srinath Sridhar, Dan Casas, Christian Theobalt*
2. [***B8***] Learning Pose Specific Representations by Predicting Different Views, *Georg Poier, David Schinagl, Horst Bischof*
3. [***B11***] Weakly and Semi Supervised Human Body Part Parsing via Pose-Guided Knowledge Transfer, *Hao-Shu Fang, Guansong Lu, Xiaolin Fang, Jianwen Xie, Yu-Wing Tai, Cewu Lu*
4. [***B14***] Person Transfer GAN to Bridge Domain Gap for Person Re-Identification, *Longhui Wei, Shiliang Zhang, Wen Gao, Qi Tian*
5. [***B17***] Cross-Modal Deep Variational Hand Pose Estimation, *Adrian Spurr, Jie Song, Seonwook Park, Otmar Hilliges*
6. [***B20***] Disentangled Person Image Generation, *Liqian Ma, Qianru Sun, Stamatios Georgoulis, Luc Van Gool, Bernt Schiele, Mario Fritz*
7. [***C1***] Super-FAN: Integrated Facial Landmark Localization and Super-Resolution of Real-World Low Resolution Faces in Arbitrary Poses With GANs, *Adrian Bulat, Georgios Tzimiropoulos*
8. [***C4***] Multistage Adversarial Losses for Pose-Based Human Image Synthesis, *Chenyang Si, Wei Wang, Liang Wang, Tieniu Tan*

0850-1010 Session 1-1C: 3D Vision I (Room 255)

Papers in this session are also in Poster Session P1-1.
Poster tag in square brackets (e.g., [***C7***])

Chairs: Tomas Pajdla *(Czech Technical Univ.)*
Torsten Sattler *(ETH Zürich)*

0850 Orals (O1-1C)

Format (12 min. for presentation + 2 min. for questions)

1. [***C7***] Rotation Averaging and Strong Duality, *Anders Eriksson, Carl Olsson, Fredrik Kahl, Tat-Jun Chin*
2. [***C10***] Hybrid Camera Pose Estimation, *Federico Camposeco, Andrea Cohen, Marc Pollefeys, Torsten Sattler*
3. [***C13***] A Certifiably Globally Optimal Solution to the Non-Minimal Relative Pose Problem, *Jesus Briales, Laurent Kneip, Javier Gonzalez-Jimenez*

0934 Spotlights (S1-1C)

Format (4 min. for presentation; no questions)

1. [***C16***] Single View Stereo Matching, *Yue Luo, Jimmy Ren, Mude Lin, Jiahao Pang, Wenxiu Sun, Hongsheng Li, Liang Lin*
2. [***C19***] Fight Ill-Posedness With Ill-Posedness: Single-Shot Variational Depth Super-Resolution From Shading, *Bjoern Haefner, Yvain Quéau, Thomas Möllenhoff, Daniel Cremers*
3. [***C22***] Deep Depth Completion of a Single RGB-D Image, *Yinda Zhang, Thomas Funkhouser*
4. [***D3***] Multi-View Harmonized Bilinear Network for 3D Object Recognition, *Tan Yu, Jingjing Meng, Junsong Yuan*
5. [***D6***] PPFNet: Global Context Aware Local Features for Robust 3D Point Matching, *Haowen Deng, Tolga Birdal, Slobodan Ilic*
6. [***D9***] FoldingNet: Point Cloud Auto-Encoder via Deep Grid Deformation, *Yaoqing Yang, Chen Feng, Yiru Shen, Dong Tian*
7. [***D12***] A Papier-Mâché Approach to Learning 3D Surface Generation, *Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan C. Russell, Mathieu Aubry*
8. [***D15***] LEGO: Learning Edge With Geometry All at Once by Watching Videos, *Zhenheng Yang, Peng Wang, Yang Wang, Wei Xu, Ram Nevatia*

1000–1045 Morning Break (Halls A-C)

1000-1830 Exhibits (Halls D-E)

- See Exhibits map for list of exhibitors.

1010-1230 Demos (Hall C)

- A Face-To-Face Neural Conversation Model, *Hang Chu, Daiqing Li, Sanja Fidler (Univ. of Toronto)*
- Real-Time Denoising of Range Data Sensed by a ToF Device Working in Low-Sensing Mode, *Mihail Georgiev, Robert Bregovic, Atanas Gotchev (Tampere Univ. of Technology)*
- OpenPose: From 2D Human Keypoint Estimation to 3D Avatar Retargeting, *Gines Hidalgo, Donglai Xiang, Aayush Bansal, Hanbyul Joo, and Yaser Sheikh (Carnegie Mellon Univ.)*
- KCNN: Extremely-Efficient Hardware Keypoint Detection With a Compact Convolutional Neural Network, *Paolo Di Febbo, Carlo Dal Mutto (Aquifi Inc.)*
- Interactive Classification for Deep Learning Interpretation, *Angel Cabrera, Fred Hohman, Jason Lin, Duen Horng Chau (Georgia Tech)*
- Gibson Environment: Real-World Perception for Embodied Agents, *Amir Zamir, Fei Xia, Jerry He, Sasha Sax, Jitendra Malik, Silvio Savarese (Stanford, Berkeley)*
- DoubleFusion: Real-Time Capture of Human Performances With Inner Body Shapes From a Single Depth Sensor, *Tao Yu, Yebin Liu, Lan Xu (Tsinghua Univ., Beihang Univ.)*

1010-1230 Poster Session P1-1 (Halls C-E)

Poster tag in square brackets (e.g., [***D18***])

3D Vision

1. [***D18***] Five-Point Fundamental Matrix Estimation for Uncalibrated Cameras, *Daniel Barath*
2. [***D21***] PointFusion: Deep Sensor Fusion for 3D Bounding Box Estimation, *Danfei Xu, Dragomir Anguelov, Ashesh Jain*
3. [***E2***] Scalable Dense Non-Rigid Structure-From-Motion: A Grassmannian Perspective, *Suryansh Kumar, Anoop Cherian, Yuchao Dai, Hongdong Li*
4. [***E5***] GVCNN: Group-View Convolutional Neural Networks for 3D Shape Recognition, *Yifan Feng, Zizhao Zhang, Xibin Zhao, Rongrong Ji, Yue Gao*
5. [***E8***] Depth and Transient Imaging With Compressive SPAD Array Cameras, *Qilin Sun, Xiong Dun, Yifan Peng, Wolfgang Heidrich*

6. [***E11***] GeoNet: Geometric Neural Network for Joint Depth and Surface Normal Estimation, *Xiaojuan Qi, Renjie Liao, Zhengzhe Liu, Raquel Urtasun, Jiaya Jia*
7. [***E14***] Real-Time Seamless Single Shot 6D Object Pose Prediction, *Bugra Tekin, Sudipta N. Sinha, Pascal Fua*
8. [***E17***] Factoring Shape, Pose, and Layout From the 2D Image of a 3D Scene, *Shubham Tulsiani, Saurabh Gupta, David F. Fouhey, Alexei A. Efros, Jitendra Malik*
9. [***E20***] Monocular Relative Depth Perception With Web Stereo Data Supervision, *Ke Xian, Chunhua Shen, Zhiguo Cao, Hao Lu, Yang Xiao, Ruibo Li, Zhenbo Luo*
10. [***F1***] Spline Error Weighting for Robust Visual-Inertial Fusion, *Hannes Ovrén, Per-Erik Forssén*
11. [***F4***] Single-Image Depth Estimation Based on Fourier Domain Analysis, *Jae-Han Lee, Minhyeok Heo, Kyung-Rae Kim, Chang-Su Kim*
12. [***F7***] Unsupervised Learning of Monocular Depth Estimation and Visual Odometry With Deep Feature Reconstruction, *Huangying Zhan, Ravi Garg, Chamara Saroj Weerasekera, Kejie Li, Harsh Agarwal, Ian Reid*

Analyzing Humans in Images

13. [***F10***] Detect-and-Track: Efficient Pose Estimation in Videos, *Rohit Girdhar, Georgia Gkioxari, Lorenzo Torresani, Manohar Paluri, Du Tran*
14. [***F13***] Supervision-by-Registration: An Unsupervised Approach to Improve the Precision of Facial Landmark Detectors, *Xuanyi Dong, Shoou-I Yu, Xinshuo Weng, Shih-En Wei, Yi Yang, Yaser Sheikh*
15. [***F16***] Diversity Regularized Spatiotemporal Attention for Video-Based Person Re-Identification, *Shuang Li, Slawomir Bak, Peter Carr, Xiaogang Wang*
16. [***F19***] Style Aggregated Network for Facial Landmark Detection, *Xuanyi Dong, Yan Yan, Wanli Ouyang, Yi Yang*
17. [***F22***] Learning Deep Models for Face Anti-Spoofing: Binary or Auxiliary Supervision, *Yaojie Liu, Amin Jourabloo, Xiaoming Liu*
18. [***G3***] Deep Cost-Sensitive and Order-Preserving Feature Learning for Cross-Population Age Estimation, *Kai Li, Junliang Xing, Chi Su, Weiming Hu, Yundong Zhang, Stephen Maybank*
19. [***G6***] First-Person Hand Action Benchmark With RGB-D Videos and 3D Hand Pose Annotations, *Guillermo Garcia-Hernando, Shanxin Yuan, Seungryul Baek, Tae-Kyun Kim*
20. [***G9***] A Pose-Sensitive Embedding for Person Re-Identification With Expanded Cross Neighborhood Re-Ranking, *M. Saquib Sarfraz, Arne Schumann, Andreas Eberle, Rainer Stiefelhagen*
21. [***G12***] Disentangling 3D Pose in a Dendritic CNN for Unconstrained 2D Face Alignment, *Amit Kumar, Rama Chellappa*
22. [***G15***] A Hierarchical Generative Model for Eye Image Synthesis and Eye Gaze Estimation, *Kang Wang, Rui Zhao, Qiang Ji*
23. [***G18***] MiCT: Mixed 3D/2D Convolutional Tube for Human Action Recognition, *Yizhou Zhou, Xiaoyan Sun, Zheng-Jun Zha, Wenjun Zeng*
24. [***G21***] Learning to Estimate 3D Human Pose and Shape From a Single Color Image, *Georgios Pavlakos, Luyang Zhu, Xiaowei Zhou, Kostas Daniilidis*
25. [***H2***] Glimpse Clouds: Human Activity Recognition From Unstructured Feature Points, *Fabien Baradel, Christian Wolf, Julien Mille, Graham W. Taylor*

Image Motion & Tracking

26. [***H5***] Context-Aware Deep Feature Compression for High-Speed Visual Tracking, *Jongwon Choi, Hyung Jin Chang, Tobias Fischer, Sangdoo Yun, Kyuewang Lee, Jiyeoup Jeong, Yiannis Demiris, Jin Young Choi*
27. [***H8***] Correlation Tracking via Joint Discrimination and Reliability Learning, *Chong Sun, Dong Wang, Huchuan Lu, Ming-Hsuan Yang*
28. [***H11***] PhaseNet for Video Frame Interpolation, *Simone Meyer, Abdelaziz Djelouah, Brian McWilliams, Alexander Sorkine-Hornung, Markus Gross, Christopher Schroers*
29. [***H14***] The Best of Both Worlds: Combining CNNs and Geometric Constraints for Hierarchical Motion Segmentation, *Pia Bideau, Aruni RoyChowdhury, Rakesh R. Menon, Erik Learned-Miller*
30. [***H17***] Hyperparameter Optimization for Tracking With Continuous Deep Q-Learning, *Xingping Dong, Jianbing Shen, Wenguan Wang, Yu Liu, Ling Shao, Fatih Porikli*
31. [***H20***] Scale-Transferrable Object Detection, *Peng Zhou, Bingbing Ni, Cong Geng, Jianguo Hu, Yi Xu*
32. [***I1***] A Prior-Less Method for Multi-Face Tracking in Unconstrained Videos, *Chung-Ching Lin, Ying Hung*
33. [***I4***] End-to-End Flow Correlation Tracking With Spatial-Temporal Attention, *Zheng Zhu, Wei Wu, Wei Zou, Junjie Yan*

Low-level & Mid-level Vision

34. [***I7***] Deep Texture Manifold for Ground Terrain Recognition, *Jia Xue, Hang Zhang, Kristin Dana*
35. [***I10***] Learning Superpixels With Segmentation-Aware Affinity Loss, *Wei-Chih Tu, Ming-Yu Liu, Varun Jampani, Deqing Sun, Shao-Yi Chien, Ming-Hsuan Yang, Jan Kautz*
36. [***I13***] Interactive Image Segmentation With Latent Diversity, *Zhuwen Li, Qifeng Chen, Vladlen Koltun*
37. [***I16***] The Unreasonable Effectiveness of Deep Features as a Perceptual Metric, *Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, Oliver Wang*
38. [***I19***] Local Descriptors Optimized for Average Precision, *Kun He, Yan Lu, Stan Sclaroff*
39. [***I22***] Recovering Realistic Texture in Image Super-Resolution by Deep Spatial Feature Transform, *Xintao Wang, Ke Yu, Chao Dong, Chen Change Loy*
40. [***J3***] Deep Extreme Cut: From Extreme Points to Object Segmentation, *Kevis-Kokitsi Maninis, Sergi Caelles, Jordi Pont-Tuset, Luc Van Gool*
41. [***J6***] Learning to Parse Wireframes in Images of Man-Made Environments, *Kun Huang, Yifan Wang, Zihan Zhou, Tianjiao Ding, Shenghua Gao, Yi Ma*
42. [***J9***] Occlusion-Aware Rolling Shutter Rectification of 3D Scenes, *Subeesh Vasu, Mahesh Mohan M. R., A. N. Rajagopalan*
43. [***J12***] Content-Sensitive Supervoxels via Uniform Tessellations on Video Manifolds, *Ran Yi, Yong-Jin Liu, Yu-Kun Lai*
44. [***J15***] Intrinsic Image Transformation via Scale Space Decomposition, *Lechao Cheng, Chengyi Zhang, Zicheng Liao*
45. [***J18***] Learned Shape-Tailored Descriptors for Segmentation, *Naeemullah Khan, Ganesh Sundaramoorthi*
46. [***J21***] PAD-Net: Multi-Tasks Guided Prediction-and-Distillation Network for Simultaneous Depth Estimation and Scene Parsing, *Dan Xu, Wanli Ouyang, Xiaogang Wang, Nicu Sebe*
47. [***K2***] Multi-Image Semantic Matching by Mining Consistent Features, *Qianqian Wang, Xiaowei Zhou, Kostas Daniilidis*
48. [***K5***] Density-Aware Single Image De-Raining Using a Multi-Stream Dense Network, *He Zhang, Vishal M. Patel*
49. [***K8***] Joint Cuts and Matching of Partitions in One Graph, *Tianshu Yu, Junchi Yan, Jieyi Zhao, Baoxin Li*
50. [***K11***] Progressive Attention Guided Recurrent Network for Salient Object Detection, *Xiaoning Zhang, Tiantian Wang, Jinqing Qi, Huchuan Lu, Gang Wang*
51. [***K14***] Fast and Accurate Single Image Super-Resolution via Information Distillation Network, *Zheng Hui, Xiumei Wang, Xinbo Gao*
52. [***K17***] Hallucinated-IQA: No-Reference Image Quality Assessment via Adversarial Learning, *Kwan-Yee Lin, Guanxiang Wang*

Machine Learning for Computer Vision

53. [***K20***] NAG: Network for Adversary Generation, *Konda Reddy Mopuri, Utkarsh Ojha, Utsav Garg, R. Venkatesh Babu*
54. [***L1***] Dynamic-Structured Semantic Propagation Network, *Xiaodan Liang, Hongfei Zhou, Eric Xing*
55. [***L4***] Cross-Domain Self-Supervised Multi-Task Feature Learning Using Synthetic Imagery, *Zhongzheng Ren, Yong Jae Lee*
56. [***L7***] A Two-Step Disentanglement Method, *Naama Hadad, Lior Wolf, Moni Shahar*
57. [***L10***] Robust Facial Landmark Detection via a Fully-Convolutional Local-Global Context Network, *Daniel Merget, Matthias Rock, Gerhard Rigoll*
58. [***L13***] Decorrelated Batch Normalization, *Lei Huang, Dawei Yang, Bo Lang, Jia Deng*
59. [***L16***] Learning to Sketch With Shortcut Cycle Consistency, *Jifei Song, Kaiyue Pang, Yi-Zhe Song, Tao Xiang, Timothy M. Hospedales*
60. [***L19***] Towards a Mathematical Understanding of the Difficulty in Learning With Feedforward Neural Networks, *Hao Shen*
61. [***L22***] FaceID-GAN: Learning a Symmetry Three-Player GAN for Identity-Preserving Face Synthesis, *Yujun Shen, Ping Luo, Junjie Yan, Xiaogang Wang, Xiaoou Tang*
62. [***M3***] A Constrained Deep Neural Network for Ordinal Regression, *Yanzhu Liu, Adams Wai Kin Kong, Chi Keong Goh*
63. [***M6***] Modulated Convolutional Networks, *Xiaodi Wang, Baochang Zhang, Ce Li, Rongrong Ji, Jungong Han, Xianbin Cao, Jianzhuang Liu*
64. [***M9***] Learning Steerable Filters for Rotation Equivariant CNNs, *Maurice Weiler, Fred A. Hamprecht, Martin Storath*

65. [***M12***] Efficient Interactive Annotation of Segmentation Datasets With Polygon-RNN++, *David Acuna, Huan Ling, Amlan Kar, Sanja Fidler*
66. [***M15***] SplineCNN: Fast Geometric Deep Learning With Continuous B-Spline Kernels, *Matthias Fey, Jan Eric Lenssen, Frank Weichert, Heinrich Müller*
67. [***M18***] GAGAN: Geometry-Aware Generative Adversarial Networks, *Jean Kossaifi, Linh Tran, Yannis Panagakis, Maja Pantic*
68. [***M21***] On the Robustness of Semantic Segmentation Models to Adversarial Attacks, *Anurag Arnab, Ondrej Miksik, Philip H.S. Torr*
69. [***N2***] Feedback-Prop: Convolutional Neural Network Inference Under Partial Evidence, *Tianlu Wang, Kota Yamaguchi, Vicente Ordonez*
70. [***N5***] Super-Resolving Very Low-Resolution Face Images With Supplementary Attributes, *Xin Yu, Basura Fernando, Richard Hartley, Fatih Porikli*

Object Recognition & Scene Understanding

71. [***N8***] Frustum PointNets for 3D Object Detection From RGB-D Data, *Charles R. Qi, Wei Liu, Chenxia Wu, Hao Su, Leonidas J. Guibas*
72. [***N11***] W2F: A Weakly-Supervised to Fully-Supervised Framework for Object Detection, *Yongqiang Zhang, Yancheng Bai, Mingli Ding, Yongqiang Li, Bernard Ghanem*
73. [***N14***] 3D Object Detection With Latent Support Surfaces, *Zhile Ren, Erik B. Sudderth*
74. [***N17***] Towards Faster Training of Global Covariance Pooling Networks by Iterative Matrix Square Root Normalization, *Peihua Li, Jiangtao Xie, Qilong Wang, Zilin Gao*
75. [***N20***] Recurrent Scene Parsing With Perspective Understanding in the Loop, *Shu Kong, Charless C. Fowlkes*
76. [***O1***] Improving Occlusion and Hard Negative Handling for Single-Stage Pedestrian Detectors, *Junhyug Noh, Soochan Lee, Beomsu Kim, Gunhee Kim*
77. [***O4***] Learning to Act Properly: Predicting and Explaining Affordances From Images, *Ching-Yao Chuang, Jiaman Li, Antonio Torralba, Sanja Fidler*
78. [***O7***] Pointwise Convolutional Neural Networks, *Binh-Son Hua, Minh-Khoi Tran, Sai-Kit Yeung*
79. [***O10***] Image-Image Domain Adaptation With Preserved Self-Similarity and Domain-Dissimilarity for Person Re-Identification, *Weijian Deng, Liang Zheng, Qixiang Ye, Guoliang Kang, Yi Yang, Jianbin Jiao*
80. [***O13***] A Generative Adversarial Approach for Zero-Shot Learning From Noisy Texts, *Yizhe Zhu, Mohamed Elhoseiny, Bingchen Liu, Xi Peng, Ahmed Elgammal*
81. [***O16***] Tensorize, Factorize and Regularize: Robust Visual Relationship Learning, *Seong Jae Hwang, Sathya N. Ravi, Zirui Tao, Hyunwoo J. Kim, Maxwell D. Collins, Vikas Singh*
82. [***O19***] Transductive Unbiased Embedding for Zero-Shot Learning, *Jie Song, Chengchao Shen, Yezhou Yang, Yang Liu, Mingli Song*
83. [***O22***] Hierarchical Novelty Detection for Visual Object Recognition, *Kibok Lee, Kimin Lee, Kyle Min, Yuting Zhang, Jinwoo Shin, Honglak Lee*
84. [***P3***] Zero-Shot Visual Recognition Using Semantics-Preserving Adversarial Embedding Networks, *Long Chen, Hanwang Zhang, Jun Xiao, Wei Liu, Shih-Fu Chang*
85. [***P6***] Learning Rich Features for Image Manipulation Detection, *Peng Zhou, Xintong Han, Vlad I. Morariu, Larry S. Davis*
86. [***P9***] Human Semantic Parsing for Person Re-Identification, *Mahdi M. Kalayeh, Emrah Basaran, Muhittin Gökmen, Mustafa E. Kamasak, Mubarak Shah*
87. [***P12***] Stacked Latent Attention for Multimodal Reasoning, *Haoqi Fan, Jiatong Zhou*
88. [***P15***] R-FCN-3000 at 30fps: Decoupling Detection and Classification, *Bharat Singh, Hengduo Li, Abhishek Sharma, Larry S. Davis*
89. [***P18***] CSRNet: Dilated Convolutional Neural Networks for Understanding the Highly Congested Scenes, *Yuhong Li, Xiaofan Zhang, Deming Chen*
90. [***P21***] Revisiting Knowledge Transfer for Training Object Class Detectors, *Jasper Uijlings, Stefan Popov, Vittorio Ferrari*

Theory of Computer Vision

91. [***Q2***] Deep Sparse Coding for Invariant Multimodal Halle Berry Neurons, *Edward Kim, Darryl Hannan, Garrett Kenyon*
92. [***Q5***] On the Convergence of PatchMatch and Its Variants, *Thibaud Ehret, Pablo Arias*

Video Analytics

93. [***Q8***] Rethinking the Faster R-CNN Architecture for Temporal Action Localization, *Yu-Wei Chao, Sudheendra Vijayanarasimhan, Bryan Seybold, David A. Ross, Jia Deng, Rahul Sukthankar*
94. [***Q11***] MoNet: Deep Motion Exploitation for Video Object Segmentation, *Huaxin Xiao, Jiashi Feng, Guosheng Lin, Yu Liu, Maojun Zhang*
95. [***Q14***] Video Representation Learning Using Discriminative Pooling, *Jue Wang, Anoop Cherian, Fatih Porikli, Stephen Gould*
96. [***Q17***] Recognizing Human Actions as the Evolution of Pose Estimation Maps, *Mengyuan Liu, Junsong Yuan*
97. [***Q20***] Video Person Re-Identification With Competitive Snippet-Similarity Aggregation and Co-Attentive Snippet Embedding, *Dapeng Chen, Hongsheng Li, Tong Xiao, Shuai Yi, Xiaogang Wang*
98. [***R1***] Mask-Guided Contrastive Attention Model for Person Re-Identification, *Chunfeng Song, Yan Huang, Wanli Ouyang, Liang Wang*
99. [***R4***] Blazingly Fast Video Object Segmentation With Pixel-Wise Metric Learning, *Yuhua Chen, Jordi Pont-Tuset, Alberto Montes, Luc Van Gool*

1130–1330 Lunch (Hall A; Halls 1-4)

Notes:

1230-1450 Poster Session P1-2 (Halls C-E)

Poster tag in square brackets (e.g., [***A2***])

Object Recognition & Scene Understanding

1. [**A2**] Learning to Compare: Relation Network for Few-Shot Learning, Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip H.S. Torr, Timothy M. Hospedales
2. [**A5**] COCO-Stuff: Thing and Stuff Classes in Context, Holger Caesar, Jasper Uijlings, Vittorio Ferrari
3. [**A8**] Image Generation From Scene Graphs, Justin Johnson, Agrim Gupta, Li Fei-Fei
4. [**A22**] Deep Cauchy Hashing for Hamming Space Retrieval, Yue Cao, Mingsheng Long, Bin Liu, Jianmin Wang
5. [**B3**] Learning to Look Around: Intelligently Exploring Unseen Environments for Unknown Tasks, Dinesh Jayaraman, Kristen Grauman
6. [**B6**] Multi-Scale Location-Aware Kernel Representation for Object Detection, Hao Wang, Qilong Wang, Mingqi Gao, Peihua Li, Wangmeng Zuo
7. [**B9**] Clinical Skin Lesion Diagnosis Using Representations Inspired by Dermatologist Criteria, Jufeng Yang, Xiaoxiao Sun, Jie Liang, Paul L. Rosin
8. [**B12**] Compare and Contrast: Learning Prominent Visual Differences, Steven Chen, Kristen Grauman
9. [**B15**] Multi-Evidence Filtering and Fusion for Multi-Label Classification, Object Detection and Semantic Segmentation Based on Weakly Supervised Learning, Weifeng Ge, Sibei Yang, Yizhou Yu
10. [**B18**] HashGAN: Deep Learning to Hash With Pair Conditional Wasserstein GAN, Yue Cao, Bin Liu, Mingsheng Long, Jianmin Wang
11. [**B21**] Min-Entropy Latent Model for Weakly Supervised Object Detection, Fang Wan, Pengxu Wei, Jianbin Jiao, Zhenjun Han, Qixiang Ye
12. [**C2**] MAttNet: Modular Attention Network for Referring Expression Comprehension, Licheng Yu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, Mohit Bansal, Tamara L. Berg
13. [**C5**] AttnGAN: Fine-Grained Text to Image Generation With Attentional Generative Adversarial Networks, Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, Xiaodong He
14. [**C8**] Adversarial Complementary Learning for Weakly Supervised Object Localization, Xiaolin Zhang, Yunchao Wei, Jiashi Feng, Yi Yang, Thomas S. Huang
15. [**C11**] Conditional Generative Adversarial Network for Structured Domain Adaptation, Weixiang Hong, Zhenzhen Wang, Ming Yang, Junsong Yuan
16. [**C14**] GroupCap: Group-Based Image Captioning With Structured Relevance and Diversity Constraints, Fuhai Chen, Rongrong Ji, Xiaoshuai Sun, Yongjian Wu, Jinsong Su
17. [**C17**] Weakly-Supervised Semantic Segmentation by Iteratively Mining Common Object Features, Xiang Wang, Shaodi You, Xi Li, Huimin Ma
18. [**C20**] Bootstrapping the Performance of Webly Supervised Semantic Segmentation, Tong Shen, Guosheng Lin, Chunhua Shen, Ian Reid
19. [**D1**] DeepVoting: A Robust and Explainable Deep Network for Semantic Part Detection Under Partial Occlusion, Zhishuai Zhang, Cihang Xie, Jianyu Wang, Lingxi Xie, Alan L. Yuille
20. [**D4**] Geometry-Aware Scene Text Detection With Instance Transformation Network, Fangfang Wang, Liming Zhao, Xi Li, Xinchao Wang, Dacheng Tao

Video Analytics

21. [***D7***] Optical Flow Guided Feature: A Fast and Robust Motion Representation for Video Action Recognition, *Shuyang Sun, Zhanghui Kuang, Lu Sheng, Wanli Ouyang, Wei Zhang*
22. [***D10***] Motion-Guided Cascaded Refinement Network for Video Object Segmentation, *Ping Hu, Gang Wang, Xiangfei Kong, Jason Kuen, Yap-Peng Tan*
23. [***D13***] A Memory Network Approach for Story-Based Temporal Summarization of 360° Videos, *Sangho Lee, Jinyoung Sung, Youngjae Yu, Gunhee Kim*
24. [***D16***] Cube Padding for Weakly-Supervised Saliency Prediction in 360° Videos, *Hsien-Tzu Cheng, Chun-Hung Chao, Jin-Dong Dong, Hao-Kai Wen, Tyng-Luh Liu, Min Sun*
25. [***D19***] Appearance-and-Relation Networks for Video Classification, *Limin Wang, Wei Li, Wen Li, Luc Van Gool*
26. [***D22***] Excitation Backprop for RNNs, *Sarah Adel Bargal, Andrea Zunino, Donghyun Kim, Jianming Zhang, Vittorio Murino, Stan Sclaroff*

27. [***E3***] One-Shot Action Localization by Learning Sequence Matching Network, *Hongtao Yang, Xuming He, Fatih Porikli*
28. [***E6***] Structure Preserving Video Prediction, *Jingwei Xu, Bingbing Ni, Zefan Li, Shuo Cheng, Xiaokang Yang*
29. [***E9***] Person Re-Identification With Cascaded Pairwise Convolutions, *Yicheng Wang, Zhenzhong Chen, Feng Wu, Gang Wang*

Machine Learning for Computer Vision

30. [**E12**] On the Importance of Label Quality for Semantic Segmentation, Aleksandar Zlateski, Ronnachai Jaroensri, Prafull Sharma, Frédo Durand
31. [**E15**] Scalable and Effective Deep CCA via Soft Decorrelation, Xiaobin Chang, Tao Xiang, Timothy M. Hospedales
32. [**E18**] Duplex Generative Adversarial Network for Unsupervised Domain Adaptation, Lanqing Hu, Meina Kan, Shiguang Shan, Xilin Chen
33. [**E21**] Edit Probability for Scene Text Recognition, Fan Bai, Zhanzhan Cheng, Yi Niu, Shiliang Pu, Shuigeng Zhou
34. [**F2**] Global Versus Localized Generative Adversarial Nets, Guo-Jun Qi, Liheng Zhang, Hao Hu, Marzieh Edraki, Jingdong Wang, Xian-Sheng Hua
35. [**F5**] MoCoGAN: Decomposing Motion and Content for Video Generation, Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, Jan Kautz
36. [**F8**] Recurrent Residual Module for Fast Inference in Videos, Bowen Pan, Wuwei Lin, Xiaolin Fang, Chaoqin Huang, Bolei Zhou, Cewu Lu
37. [**F11**] Improving Landmark Localization With Semi-Supervised Learning, Sina Honari, Pavlo Molchanov, Stephen Tyree, Pascal Vincent, Christopher Pal, Jan Kautz
38. [**F14**] Adversarial Data Programming: Using GANs to Relax the Bottleneck of Curated Labeled Data, Arghya Pal, Vineeth N. Balasubramanian
39. [**F17**] Stochastic Variational Inference With Gradient Linearization, Tobias Plötz, Anne S. Wannenwetsch, Stefan Roth
40. [**F20**] Multi-Label Zero-Shot Learning With Structured Knowledge Graphs, Chung-Wei Lee, Wei Fang, Chih-Kuan Yeh, Yu-Chiang Frank Wang
41. [**G1**] MorphNet: Fast & Simple Resource-Constrained Structure Learning of Deep Networks, Ariel Gordon, Elad Eban, Ofir Nachum, Bo Chen, Hao Wu, Tien-Ju Yang, Edward Choi
42. [**G4**] Deep Adversarial Subspace Clustering, Pan Zhou, Yunqing Hou, Jiashi Feng
43. [**G7**] Towards Human-Machine Cooperation: Self-Supervised Sample Mining for Object Detection, Keze Wang, Xiaopeng Yan, Dongyu Zhang, Lei Zhang, Liang Lin
44. [**G10**] Discrete-Continuous ADMM for Transductive Inference in Higher-Order MRFs, Emanuel Laude, Jan-Hendrik Lange, Jonas Schüpfer, Csaba Domokos, Laura Leal-Taixé, Frank R. Schmidt, Bjoern Andres, Daniel Cremers
45. [**G13**] Robust Physical-World Attacks on Deep Learning Visual Classification, Kevin Eykholt, Ivan Evtimov, Earlence Fernandes, Bo Li, Amir Rahmati, Chaowei Xiao, Atul Prakash, Tadayoshi Kohno, Dawn Song
46. [**G16**] Generating a Fusion Image: One's *Identity* and Another's *Shape*, DongGyu Joo, Doyeon Kim, Junmo Kim

Low-level & Mid-level Vision

47. [**G19**] Learning to Promote Saliency Detectors, Yu Zeng, Huchuan Lu, Lihe Zhang, Mengyang Feng, Ali Borji
48. [**G22**] Image Super-Resolution via Dual-State Recurrent Networks, Wei Han, Shiyu Chang, Ding Liu, Mo Yu, Michael Witbrock, Thomas S. Huang
49. [**H3**] Deep Back-Projection Networks for Super-Resolution, Muhammad Haris, Gregory Shakhnarovich, Norimichi Ukita
50. [**H6**] Focus Manipulation Detection via Photometric Histogram Analysis, Can Chen, Scott McCloskey, Jingyi Yu
51. [**H9**] Compassionately Conservative Balanced Cuts for Image Segmentation, Nathan D. Cahill, Tyler L. Hayes, Renee T. Meinhold, John F. Hamilton
52. [**H12**] A High-Quality Denoising Dataset for Smartphone Cameras, Abdelrahman Abdelhamed, Stephen Lin, Michael S. Brown
53. [**H15**] Context-Aware Synthesis for Video Frame Interpolation, Simon Niklaus, Feng Liu
54. [**H18**] Salient Object Detection Driven by Fixation Prediction, Wenguan Wang, Jianbing Shen, Xingping Dong, Ali Borji

55. [**H21**] Enhancing the Spatial Resolution of Stereo Images Using a Parallax Prior, Daniel S. Jeon, Seung-Hwan Baek, Inchang Choi, Min H. Kim
56. [**I2**] HATS: Histograms of Averaged Time Surfaces for Robust Event-Based Object Classification, Amos Sironi, Manuele Brambilla, Nicolas Bourdis, Xavier Lagorce, Ryad Benosman
57. [**I5**] A Bi-Directional Message Passing Model for Salient Object Detection, Lu Zhang, Ju Dai, Huchuan Lu, You He, Gang Wang
58. [**I8**] Matching Pixels Using Co-Occurrence Statistics, Rotal Kat, Roy Jevnisek, Shai Avidan
59. [**I11**] SeedNet: Automatic Seed Generation With Deep Reinforcement Learning for Robust Interactive Segmentation, Gwangmo Song, Heesoo Myeong, Kyoung Mu Lee
60. [**I14**] Jerk-Aware Video Acceleration Magnification, Shoichiro Takeda, Kazuki Okami, Dan Mikami, Megumi Isogai, Hideaki Kimata
61. [**I17**] Defense Against Adversarial Attacks Using High-Level Representation Guided Denoiser, Fangzhou Liao, Ming Liang, Yinpeng Dong, Tianyu Pang, Xiaolin Hu, Jun Zhu
62. [**I20**] Stacked Conditional Generative Adversarial Networks for Jointly Learning Shadow Detection and Shadow Removal, Jifeng Wang, Xiang Li, Jian Yang
63. [**J1**] Image Correction via Deep Reciprocating HDR Transformation, Xin Yang, Ke Xu, Yibing Song, Qiang Zhang, Xiaopeng Wei, Rynson W.H. Lau
64. [**J4**] PieAPP: Perceptual Image-Error Assessment Through Pairwise Preference, Ekta Prashnani, Hong Cai, Yasamin Mostofi, Pradeep Sen
65. [**J7**] Normalized Cut Loss for Weakly-Supervised CNN Segmentation, Meng Tang, Abdelaziz Djelouah, Federico Perazzi, Yuri Boykov, Christopher Schroers
66. [**J10**] ISTA-Net: Interpretable Optimization-Inspired Deep Network for Image Compressive Sensing, Jian Zhang, Bernard Ghanem
67. [**J13**] Fast End-to-End Trainable Guided Filter, Huikai Wu, Shuai Zheng, Junge Zhang, Kaiqi Huang
68. [**J16**] Disentangling Structure and Aesthetics for Style-Aware Image Completion, Andrew Gilbert, John Collomosse, Hailin Jin, Brian Price
69. [**J19**] Learning a Discriminative Feature Network for Semantic Segmentation, Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, Nong Sang
70. [**J22**] Kernelized Subspace Pooling for Deep Local Descriptors, Xing Wei, Yue Zhang, Yihong Gong, Nanning Zheng

3D Vision

71. [***K3***] pOSE: Pseudo Object Space Error for Initialization-Free Bundle Adjustment, *Je Hyeong Hong, Christopher Zach*
72. [***K6***] Deformable Shape Completion With Graph Convolutional Autoencoders, *Or Litany, Alex Bronstein, Michael Bronstein, Ameesh Makadia*
73. [***K9***] Learning From Millions of 3D Scans for Large-Scale 3D Face Recognition, *Syed Zulqarnain Gilani, Ajmal Mian*
74. [***K12***] CarFusion: Combining Point Tracking and Part Detection for Dynamic 3D Reconstruction of Vehicles, *N. Dinesh Reddy, Minh Vo, Srinivasa G. Narasimhan*
75. [***K15***] Deep Material-Aware Cross-Spectral Stereo Matching, *Tiancheng Zhi, Bernardo R. Pires, Martial Hebert, Srinivasa G. Narasimhan*
76. [***K18***] Augmenting Crowd-Sourced 3D Reconstructions Using Semantic Detections, *True Price, Johannes L. Schönberger, Zhen Wei, Marc Pollefeys, Jan-Michael Frahm*
77. [***K21***] Matryoshka Networks: Predicting 3D Geometry via Nested Shape Layers, *Stephan R. Richter, Stefan Roth*
78. [***L2***] Triplet-Center Loss for Multi-View 3D Object Retrieval, *Xinwei He, Yang Zhou, Zhichao Zhou, Song Bai, Xiang Bai*
79. [***L5***] Learning 3D Shape Completion From Laser Scan Data With Weak Supervision, *David Stutz, Andreas Geiger*
80. [***L8***] End-to-End Learning of Keypoint Detector and Descriptor for Pose Invariant 3D Matching, *Georgios Georgakis, Srikrishna Karanam, Ziyan Wu, Jan Ernst, Jana Košecká*
81. [***L11***] ICE-BA: Incremental, Consistent and Efficient Bundle Adjustment for Visual-Inertial SLAM, *Haomin Liu, Mingyu Chen, Guofeng Zhang, Hujun Bao, Yingze Bao*
82. [***L14***] GeoNet: Unsupervised Learning of Dense Depth, Optical Flow and Camera Pose, *Zhichao Yin, Jianping Shi*
83. [***L17***] Radially-Distorted Conjugate Translations, *James Pritts, Zuzana Kukelova, Viktor Larsson, Ondřej Chum*

84. [***L20***] Deep Ordinal Regression Network for Monocular Depth Estimation, *Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, Dacheng Tao*
85. [***M1***] Analytical Modeling of Vanishing Points and Curves in Catadioptric Cameras, *Pedro Miraldo, Francisco Eiras, Srikumar Ramalingam*
86. [***M4***] Learning Depth From Monocular Videos Using Direct Methods, *Chaoyang Wang, José Miguel Buenaposada, Rui Zhu, Simon Lucey*
87. [***M7***] Salience Guided Depth Calibration for Perceptually Optimized Compressive Light Field 3D Display, *Shizheng Wang, Wenjuan Liao, Phil Surman, Zhigang Tu, Yuanjin Zheng, Junsong Yuan*
88. [***M10***] MegaDepth: Learning Single-View Depth Prediction From Internet Photos, *Zhengqi Li, Noah Snavely*
89. [***M13***] LayoutNet: Reconstructing the 3D Room Layout From a Single RGB Image, *Chuhang Zou, Alex Colburn, Qi Shan, Derek Hoiem*
90. [***M16***] CBMV: A Coalesced Bidirectional Matching Volume for Disparity Estimation, *Konstantinos Batsos, Changjiang Cai, Philippos Mordohai*
91. [***M19***] Zoom and Learn: Generalizing Deep Stereo Matching to Novel Domains, *Jiahao Pang, Wenxiu Sun, Chengxi Yang, Jimmy Ren, Ruichao Xiao, Jin Zeng, Liang Lin*

Analyzing Humans in Images

92. [***M22***] Exploring Disentangled Feature Representation Beyond Face Identification, *Yu Liu, Fangyin Wei, Jing Shao, Lu Sheng, Junjie Yan, Xiaogang Wang*
93. [***N3***] Learning Facial Action Units From Web Images With Scalable Weakly Supervised Clustering, *Kaili Zhao, Wen-Sheng Chu, Aleix M. Martinez*
94. [***N6***] Human Pose Estimation With Parsing Induced Learner, *Xuecheng Nie, Jiashi Feng, Yiming Zuo, Shuicheng Yan*
95. [***N9***] Multi-Level Factorisation Net for Person Re-Identification, *Xiaobin Chang, Timothy M. Hospedales, Tao Xiang*
96. [***N12***] Attention-Aware Compositional Network for Person Re-Identification, *Jing Xu, Rui Zhao, Feng Zhu, Huaming Wang, Wanli Ouyang*
97. [**N15**] Look at Boundary: A Boundary-Aware Face Alignment Algorithm, Wayne Wu, Chen Qian, Shuo Yang, Quan Wang, Yici Cai, Qiang Zhou
98. [***N18***] Demo2Vec: Reasoning Object Affordances From Online Videos, *Kuan Fang, Te-Lin Wu, Daniel Yang, Silvio Savarese, Joseph J. Lim*
99. [***N21***] Monocular 3D Pose and Shape Estimation of Multiple People in Natural Scenes - The Importance of Multiple Scene Constraints, *Andrei Zanfir, Elisabeta Marinoiu, Cristian Sminchisescu*
100. [***O2***] 3D Human Sensing, Action and Emotion Recognition in Robot Assisted Therapy of Children With Autism, *Elisabeta Marinoiu, Mihai Zanfir, Vlad Olaru, Cristian Sminchisescu*
101. [***O5***] Facial Expression Recognition by De-Expression Residue Learning, *Huiyuan Yang, Umur Ciftci, Lijun Yin*
102. [***O8***] A Causal And-Or Graph Model for Visibility Fluent Reasoning in Tracking Interacting Objects, *Yuanlu Xu, Lei Qin, Xiaobai Liu, Jianwen Xie, Song-Chun Zhu*
103. [***O11***] Weakly Supervised Facial Action Unit Recognition Through Adversarial Training, *Guozhu Peng, Shangfei Wang*
104. [***O14***] Non-Linear Temporal Subspace Representations for Activity Recognition, *Anoop Cherian, Suvrit Sra, Stephen Gould, Richard Hartley*
105. [***O17***] Towards Pose Invariant Face Recognition in the Wild, *Jian Zhao, Yu Cheng, Yan Xu, Lin Xiong, Jianshu Li, Fang Zhao, Karlekar Jayashree, Sugiri Pranata, Shengmei Shen, Junliang Xing, Shuicheng Yan, Jiashi Feng*
106. [***O20***] Unifying Identification and Context Learning for Person Recognition, *Qingqiu Huang, Yu Xiong, Dahua Lin*
107. [***P1***] Jointly Optimize Data Augmentation and Network Training: Adversarial Data Augmentation in Human Pose Estimation, *Xi Peng, Zhiqiang Tang, Fei Yang, Rogerio S. Feris, Dimitris Metaxas*
108. [***P4***] Wing Loss for Robust Facial Landmark Localisation With Convolutional Neural Networks, *Zhen-Hua Feng, Josef Kittler, Muhammad Awais, Patrik Huber, Xiao-Jun Wu*
109. [***P7***] Multiple Granularity Group Interaction Prediction, *Taiping Yao, Minsi Wang, Bingbing Ni, Huawei Wei, Xiaokang Yang*
110. [***P10***] Social GAN: Socially Acceptable Trajectories With Generative Adversarial Networks, *Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, Alexandre Alahi*
111. [***P13***] Deep Group-Shuffling Random Walk for Person Re-Identification, *Yantao Shen, Hongsheng Li, Tong Xiao, Shuai Yi, Dapeng Chen, Xiaogang Wang*

112. [***P16***] Transferable Joint Attribute-Identity Deep Learning for Unsupervised Person Re-Identification, *Jingya Wang, Xiatian Zhu, Shaogang Gong, Wei Li*
113. [***P19***] Harmonious Attention Network for Person Re-Identification, *Wei Li, Xiatian Zhu, Shaogang Gong*
114. [***P22***] Real-Time Rotation-Invariant Face Detection With Progressive Calibration Networks, *Xuepeng Shi, Shiguang Shan, Meina Kan, Shuzhe Wu, Xilin Chen*
115. [***Q3***] Deep Regression Forests for Age Estimation, *Wei Shen, Yilu Guo, Yan Wang, Kai Zhao, Bo Wang, Alan L. Yuille*
116. [***Q6***] Weakly-Supervised Deep Convolutional Neural Network Learning for Facial Action Unit Intensity Estimation, *Yong Zhang, Weiming Dong, Bao-Gang Hu, Qiang Ji*
117. [***Q9***] Memory Based Online Learning of Deep Representations From Video Streams, *Federico Pernici, Federico Bartoli, Matteo Bruni, Alberto Del Bimbo*
118. [***Q12***] Efficient and Deep Person Re-Identification Using Multi-Level Similarity, *Yiluan Guo, Ngai-Man Cheung*

Applications

119. [***Q15***] Multi-Level Fusion Based 3D Object Detection From Monocular Images, *Bin Xu, Zhenzhong Chen*
120. [***Q18***] A Perceptual Measure for Deep Single Image Camera Calibration, *Yannick Hold-Geoffroy, Kalyan Sunkavalli, Jonathan Eisenmann, Matthew Fisher, Emiliano Gambaretto, Sunil Hadap, Jean-François Lalonde*
121. [***Q21***] Learning to Generate Time-Lapse Videos Using Multi-Stage Dynamic Generative Adversarial Networks, *Wei Xiong, Wenhan Luo, Lin Ma, Wei Liu, Jiebo Luo*
122. [***R2***] Document Enhancement Using Visibility Detection, *Netanel Kligler, Sagi Katz, Ayellet Tal*
123. [***R5***] A Weighted Sparse Sampling and Smoothing Frame Transition Approach for Semantic Fast-Forward First-Person Videos, *Michel Silva, Washington Ramos, João Ferreira, Felipe Chamone, Mario Campos, Erickson R. Nascimento*

1420–1450 Afternoon Break (Halls A-C)

1450-1630 Session 1-2A: Machine Learning for Computer Vision I (Ballroom)

Papers in this session are also in Poster Session P1-3. Poster tag in square brackets (e.g., [***A3***])

Chairs: Sanja Fidler *(Univ. of Toronto)*
Simon Lucey *(Carnegie Mellon Univ.)*

1450 Orals (O1-2A)

Format (12 min. for presentation + 2 min. for questions)

1. [***A3***] Context Contrasted Feature and Gated Multi-Scale Aggregation for Scene Segmentation, *Henghui Ding, Xudong Jiang, Bing Shuai, Ai Qun Liu, Gang Wang*
1. [**A6**] Deep Layer Aggregation, *Fisher Yu, Dequan Wang, Evan Shelhamer, Trevor Darrell*
2. [***A9***] Convolutional Neural Networks With Alternately Updated Clique, *Yibo Yang, Zhisheng Zhong, Tiancheng Shen, Zhouchen Lin*
3. [***B1***] Practical Block-Wise Neural Network Architecture Generation, *Zhao Zhong, Junjie Yan, Wei Wu, Jing Shao, Cheng-Lin Liu*

1548 Spotlights (S1-2A)

Format (4 min. for presentation; no questions)

1. [***B4***] xUnit: Learning a Spatial Activation Function for Efficient Image Restoration, *Idan Kligvasser, Tamar Rott Shaham, Tomer Michaeli*
2. [***B7***] Crafting a Toolchain for Image Restoration by Deep Reinforcement Learning, *Ke Yu, Chao Dong, Liang Lin, Chen Change Loy*
3. [***B10***] Deformation Aware Image Compression, *Tamar Rott Shaham, Tomer Michaeli*
4. [***B13***] Distributable Consistent Multi-Object Matching, *Nan Hu, Qixing Huang, Boris Thibert, Leonidas J. Guibas*
5. [***B16***] Residual Dense Network for Image Super-Resolution, *Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, Yun Fu*
6. [***B19***] Attentive Generative Adversarial Network for Raindrop Removal From a Single Image, *Rui Qian, Robby T. Tan, Wenhan Yang, Jiajun Su, Jiaying Liu*
7. [***B22***] FSRNet: End-to-End Learning Face Super-Resolution With Facial Priors, *Yu Chen, Ying Tai, Xiaoming Liu, Chunhua Shen, Jian Yang*
8. [***C3***] Burst Denoising With Kernel Prediction Networks, *Ben Mildenhall, Jonathan T. Barron, Jiawen Chen, Dillon Sharlet, Ren Ng, Robert Carroll*

9. [***C6***] Unsupervised Sparse Dirichlet-Net for Hyperspectral Image Super-Resolution, *Ying Qu, Hairong Qi, Chiman Kwan*
10. [***C9***] Dynamic Scene Deblurring Using Spatially Variant Recurrent Neural Networks, *Jiawei Zhang, Jinshan Pan, Jimmy Ren, Yibing Song, Linchao Bao, Rynson W.H. Lau, Ming-Hsuan Yang*

1450-1630 Session 1-2B: 3D Vision II (Room 155)

Papers in this session are also in Poster Session P1-3. Poster tag in square brackets (e.g., [***C12***])

Chairs: Derek Hoiem *(UIUC)*
Mathieu Aubry *(École des Ponts ParisTech)*

1450 Orals (O1-2B)

Format (12 min. for presentation + 2 min. for questions)

1. [***C12***] SPLATNet: Sparse Lattice Networks for Point Cloud Processing, *Hang Su, Varun Jampani, Deqing Sun, Subhransu Maji, Evangelos Kalogerakis, Ming-Hsuan Yang, Jan Kautz*
2. [***C15***] Surface Networks, *Ilya Kostrikov, Zhongshi Jiang, Daniele Panozzo, Denis Zorin, Joan Bruna*
3. [***C18***] Self-Supervised Multi-Level Face Model Learning for Monocular Reconstruction at Over 250 Hz, *Ayush Tewari, Michael Zollhöfer, Pablo Garrido, Florian Bernard, Hyeongwoo Kim, Patrick Pérez, Christian Theobalt*
4. [***C21***] CodeSLAM — Learning a Compact, Optimisable Representation for Dense Visual SLAM, *Michael Bloesch, Jan Czarnowski, Ronald Clark, Stefan Leutenegger, Andrew J. Davison*

1548 Spotlights (S1-2B)

Format (4 min. for presentation; no questions)

1. [***D2***] SGPN: Similarity Group Proposal Network for 3D Point Cloud Instance Segmentation, *Weiyue Wang, Ronald Yu, Qiangui Huang, Ulrich Neumann*
2. [***D5***] PlaneNet: Piece-Wise Planar Reconstruction From a Single RGB Image, *Chen Liu, Jimei Yang, Duygu Ceylan, Ersin Yumer, Yasutaka Furukawa*
3. [***D8***] Deep Parametric Continuous Convolutional Neural Networks, *Shenlong Wang, Simon Suo, Wei-Chiu Ma, Andrei Pokrovsky, Raquel Urtasun*
4. [***D11***] FeaStNet: Feature-Steered Graph Convolutions for 3D Shape Analysis, *Nitika Verma, Edmond Boyer, Jakob Verbeek*
5. [***D14***] Image Collection Pop-Up: 3D Reconstruction and Clustering of Rigid and Non-Rigid Categories, *Antonio Agudo, Melcior Pijoan, Francesc Moreno-Noguer*
6. [***D17***] Geometry-Aware Learning of Maps for Camera Localization, *Samarth Brahmbhatt, Jinwei Gu, Kihwan Kim, James Hays, Jan Kautz*
7. [***D20***] Recurrent Slice Networks for 3D Segmentation of Point Clouds, *Qiangui Huang, Weiyue Wang, Ulrich Neumann*
8. [***E1***] Depth-Based 3D Hand Pose Estimation: From Current Achievements to Future Goals, *Shanxin Yuan, Guillermo Garcia-Hernando, Björn Stenger, Gyeongsik Moon, Ju Yong Chang, Kyoung Mu Lee, Pavlo Molchanov, Jan Kautz, Sina Honari, Liuhao Ge, Junsong Yuan, Xinghao Chen, Guijin Wang, Fan Yang, Kai Akiyama, Yang Wu, Qingfu Wan, Meysam Madadi, Sergio Escalera, Shile Li, Dongheui Lee, Iason Oikonomidis, Antonis Argyros, Tae-Kyun Kim*
9. [***E4***] SobolevFusion: 3D Reconstruction of Scenes Undergoing Free Non-Rigid Motion, *Miroslava Slavcheva, Maximilian Baust, Slobodan Ilic*
10. [***E7***] AdaDepth: Unsupervised Content Congruent Adaptation for Depth Estimation, *Jogendra Nath Kundu, Phani Krishna Uppala, Anuj Pahuja, R. Venkatesh Babu*

1450-1630 Session 1-2C: Machine Learning for Computer Vision II (Room 255)

Papers in this session are also in Poster Session P1-3. Poster tag in square brackets (e.g., [***E10***])

Chairs: Minsu Cho *(Pohang Univ. of Science & Technology)*
Barath Hariharan *(Cornell Univ.)*

1450 Orals (O1-2C)

Format (12 min. for presentation + 2 min. for questions)

1. [***E10***] Learning to Find Good Correspondences, *Kwang Moo Yi, Eduard Trulls, Yuki Ono, Vincent Lepetit, Mathieu Salzmann, Pascal Fua*
2. [***E13***] OATM: Occlusion Aware Template Matching by Consensus Set Maximization, *Simon Korman, Mark Milam, Stefano Soatto*
3. [***E16***] Deep Learning of Graph Matching, *Andrei Zanfir, Cristian Sminchisescu*
4. [***E19***] Unsupervised Discovery of Object Landmarks as Structural Representations, *Yuting Zhang, Yijie Guo, Yixin Jin, Yijun Luo, Zhiyuan He, Honglak Lee*

1548 Spotlights (S1-2C)

Format (4 min. for presentation; no questions)

1. [***E22***] Quantization and Training of Neural Networks for Efficient Integer-Arithmetic-Only Inference, *Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, Dmitry Kalenichenko*
2. [***F3***] Lean Multiclass Crowdsourcing, *Grant Van Horn, Steve Branson, Scott Loarie, Serge Belongie, Pietro Perona*
3. [***F6***] Partial Transfer Learning With Selective Adversarial Networks, *Zhangjie Cao, Mingsheng Long, Jianmin Wang, Michael I. Jordan*
4. [***F9***] Self-Supervised Feature Learning by Learning to Spot Artifacts, *Simon Jenni, Paolo Favaro*
5. [***F12***] LDMNet: Low Dimensional Manifold Regularized Neural Networks, *Wei Zhu, Qiang Qiu, Jiaji Huang, Robert Calderbank, Guillermo Sapiro, Ingrid Daubechies*
6. [***F15***] CondenseNet: An Efficient DenseNet Using Learned Group Convolutions, *Gao Huang, Shichen Liu, Laurens van der Maaten, Kilian Q. Weinberger*
7. [***F18***] Learning Deep Descriptors With Scale-Aware Triplet Networks, *Michel Keller, Zetao Chen, Fabiola Maffra, Patrik Schmuck, Margarita Chli*
8. [***F21***] Decoupled Networks, *Weiyang Liu, Zhen Liu, Zhiding Yu, Bo Dai, Rongmei Lin, Yisen Wang, James M. Rehg, Le Song*
9. [***G2***] Deep Adversarial Metric Learning, *Yueqi Duan, Wenzhao Zheng, Xudong Lin, Jiwen Lu, Jie Zhou*

1630-1830 Demos (Hall C)

- A Face-To-Face Neural Conversation Model, *Hang Chu, Daiqing Li, Sanja Fidler (Univ. of Toronto)*
- Real-Time Denoising of Range Data Sensed by a ToF Device Working in Low-Sensing Mode, *Mihail Georgiev, Robert Bregovic, Atanas Gotchev (Tampere Univ. of Technology)*
- OpenPose: From 2D Human Keypoint Estimation to 3D Avatar Retargeting, *Gines Hidalgo, Donglai Xiang, Aayush Bansal, Hanbyul Joo, and Yaser Sheikh (Carnegie Mellon Univ.)*
- Through-Wall Human Pose Estimation Using Radio Signals, *Mingmin Zhao, Rumen Hristov, Hang Zhao, Antonio Torralba, Dina Katabi (MIT)*
- Watching Soccer in AR/VR, *Konstantinos Rematas, Ira Kemelmacher-Shlizerman, Brian Curless, Steve Seitz (Univ. of Washington)*
- Touchscreen Biometrics: Signature and Password Authentication, *Ruben Vera-Rodriguez, Ruben Tolosana, Julian Fierrez and Javier Ortega-Garcia (Universidad Autonoma de Madrid)*
- Visual Turing: Human-In-The-Loop Evaluation for Computer Vision Challenges, *Deshraj Yadav, Rishabh Jain, Abhishek Das, Satwik Kottur, Devi Parikh, Dhruv Batra (Georgia Tech)*
- Real-Time LCR-Net: Detecting Multi-Person 2D and 3D Poses in Real-World Images, *Gregory Rogez, Philippe Weinzaepfel, Xavier Martin, Cordelia Schmid (INRIA)*

1630-1830 Poster Session P1-3 (Halls C-E)

Poster tag in square brackets (e.g., [***G5***])

3D Vision

1. [**G5**] PU-Net: Point Cloud Upsampling Network, *Lequan Yu, Xianzhi Li, Chi-Wing Fu, Daniel Cohen-Or, Pheng-Ann Heng*
2. [***G8***] Real-Time Monocular Depth Estimation Using Synthetic Data With Domain Adaptation via Image Style Transfer, *Amir Atapour-Abarghouei, Toby P. Breckon*
3. [***G11***] Learning for Disparity Estimation Through Feature Constancy, *Zhengfa Liang, Yiliu Feng, Yulan Guo, Hengzhu Liu, Wei Chen, Linbo Qiao, Li Zhou, Jianfeng Zhang*
4. [***G14***] DeepMVS: Learning Multi-View Stereopsis, *Po-Han Huang, Kevin Matzen, Johannes Kopf, Narendra Ahuja, Jia-Bin Huang*
5. [***G17***] Self-Calibrating Polarising Radiometric Calibration, *Daniel Teo, Boxin Shi, Yinqiang Zheng, Sai-Kit Yeung*
6. [***G20***] Coding Kendall's Shape Trajectories for 3D Action Recognition, *Amor Ben Tanfous, Hassen Drira, Boulbaba Ben Amor*
7. [***H1***] Efficient, Sparse Representation of Manifold Distance Matrices for Classical Scaling, *Javier S. Turek, Alexander G. Huth*
8. [***H4***] Motion Segmentation by Exploiting Complementary Geometric Models, *Xun Xu, Loong Fah Cheong, Zhuwen Li*
9. [***H7***] Estimation of Camera Locations in Highly Corrupted Scenarios: All About That Base, No Shape Trouble, *Yunpeng Shi, Gilad Lerman*
10. [***H10***] 4D Human Body Correspondences From Panoramic Depth Maps, *Zhong Li, Minye Wu, Wangyiteng Zhou, Jingyi Yu*

11. [***H13***] Reconstructing Thin Structures of Manifold Surfaces by Integrating Spatial Curves, *Shiwei Li, Yao Yao, Tian Fang, Long Quan*
12. [***H16***] Multi-View Consistency as Supervisory Signal for Learning Shape and Pose Prediction, *Shubham Tulsiani, Alexei A. Efros, Jitendra Malik*
13. [***H19***] Probabilistic Plant Modeling via Multi-View Image-to-Image Translation, *Takahiro Isokane, Fumio Okura, Ayaka Ide, Yasuyuki Matsushita, Yasushi Yagi*
14. [***H22***] Deep Marching Cubes: Learning Explicit Surface Representations, *Yiyi Liao, Simon Donné, Andreas Geiger*
15. [***I3***] Tags2Parts: Discovering Semantic Regions From Shape Tags, *Sanjeev Muralikrishnan, Vladimir G. Kim, Siddhartha Chaudhuri*
16. [***I6***] Uncalibrated Photometric Stereo Under Natural Illumination, *Zhipeng Mo, Boxin Shi, Feng Lu, Sai-Kit Yeung, Yasuyuki Matsushita*
17. [***I9***] Robust Depth Estimation From Auto Bracketed Images, *Sunghoon Im, Hae-Gon Jeon, In So Kweon*
18. [***I12***] Free Supervision From Video Games, *Philipp Krähenbühl*
19. [***I15***] Planar Shape Detection at Structural Scales, *Hao Fang, Florent Lafarge, Mathieu Desbrun*
20. [***I18***] Pix3D: Dataset and Methods for Single-Image 3D Shape Modeling, *Xingyuan Sun, Jiajun Wu, Xiuming Zhang, Zhoutong Zhang, Chengkai Zhang, Tianfan Xue, Joshua B. Tenenbaum, William T. Freeman*
21. [***I21***] Camera Pose Estimation With Unknown Principal Point, *Viktor Larsson, Zuzana Kukelova, Yinqiang Zheng*
22. [***J2***] Inverse Composition Discriminative Optimization for Point Cloud Registration, *Jayakorn Vongkulbhisal, Beñat Irastorza Ugalde, Fernando De la Torre, João P. Costeira*
23. [***J5***] SurfConv: Bridging 3D and 2D Convolution for RGBD Images, *Hang Chu, Wei-Chiu Ma, Kaustav Kundu, Raquel Urtasun, Sanja Fidler*
24. [***J8***] A Fast Resection-Intersection Method for the Known Rotation Problem, *Qianggong Zhang, Tat-Jun Chin, Huu Minh Le*
25. [***J11***] 3D Pose Estimation and 3D Model Retrieval for Objects in the Wild, *Alexander Grabner, Peter M. Roth, Vincent Lepetit*
26. [***J14***] Structure From Recurrent Motion: From Rigidity to Recurrency, *Xiu Li, Hongdong Li, Hanbyul Joo, Yebin Liu, Yaser Sheikh*
27. [***J17***] Learning Patch Reconstructability for Accelerating Multi-View Stereo, *Alex Poms, Chenglei Wu, Shoou-I Yu, Yaser Sheikh*
28. [***J20***] Progressively Complementarity-Aware Fusion Network for RGB-D Salient Object Detection, *Hao Chen, Youfu Li*
29. [***K1***] Pixels, Voxels, and Views: A Study of Shape Representations for Single View 3D Object Shape Prediction, *Daeyun Shin, Charless C. Fowlkes, Derek Hoiem*

Low-level & Mid-level Vision

30. [***K4***] Learning Dual Convolutional Neural Networks for Low-Level Vision, *Jinshan Pan, Sifei Liu, Deqing Sun, Jiawei Zhang, Yang Liu, Jimmy Ren, Zechao Li, Jinhui Tang, Huchuan Lu, Yu-Wing Tai, Ming-Hsuan Yang*
31. [***K7***] Defocus Blur Detection via Multi-Stream Bottom-Top-Bottom Fully Convolutional Network, *Wenda Zhao, Fan Zhao, Dong Wang, Huchuan Lu*
32. [***K10***] PiCANet: Learning Pixel-Wise Contextual Attention for Saliency Detection, *Nian Liu, Junwei Han, Ming-Hsuan Yang*
33. [***K13***] Curve Reconstruction via the Global Statistics of Natural Curves, *Ehud Barnea, Ohad Ben-Shahar*
34. [***K16***] What Do Deep Networks Like to See?, *Sebastian Palacio, Joachim Folz, Jörn Hees, Federico Raue, Damian Borth, Andreas Dengel*
35. [***K19***] "Zero-Shot" Super-Resolution Using Deep Internal Learning, *Assaf Shocher, Nadav Cohen, Michal Irani*
36. [***K22***] Detect Globally, Refine Locally: A Novel Approach to Saliency Detection, *Tiantian Wang, Lihe Zhang, Shuo Wang, Huchuan Lu, Gang Yang, Xiang Ruan, Ali Borji*
37. [***L3***] Beyond the Pixel-Wise Loss for Topology-Aware Delineation, *Agata Mosinska, Pablo Márquez-Neila, Mateusz Koziński, Pascal Fua*
38. [***L6***] KIPPI: KInetic Polygonal Partitioning of Images, *Jean-Philippe Bauchet, Florent Lafarge*
39. [***L9***] Image Blind Denoising With Generative Adversarial Network Based Noise Modeling, *Jingwen Chen, Jiawei Chen, Hongyang Chao, Ming Yang*

40. [***L12***] Multi-Scale Weighted Nuclear Norm Image Restoration, *Noam Yair, Tomer Michaeli*
41. [***L15***] MoNet: Moments Embedding Network, *Mengran Gou, Fei Xiong, Octavia Camps, Mario Sznaier*
42. [***L18***] Active Fixation Control to Predict Saccade Sequences, *Calden Wloka, Iuliia Kotseruba, John K. Tsotsos*
43. [***L21***] Densely Connected Pyramid Dehazing Network, *He Zhang, Vishal M. Patel*
44. [***M2***] Universal Denoising Networks : A Novel CNN Architecture for Image Denoising, *Stamatios Lefkimmiatis*
45. [***M5***] Learning Convolutional Networks for Content-Weighted Image Compression, *Mu Li, Wangmeng Zuo, Shuhang Gu, Debin Zhao, David Zhang*
46. [***M8***] Deep Video Super-Resolution Network Using Dynamic Upsampling Filters Without Explicit Motion Compensation, *Younghyun Jo, Seoung Wug Oh, Jaeyeon Kang, Seon Joo Kim*
47. [***M11***] Erase or Fill? Deep Joint Recurrent Rain Removal and Reconstruction in Videos, *Jiaying Liu, Wenhan Yang, Shuai Yang, Zongming Guo*
48. [***M14***] Flow Guided Recurrent Neural Encoder for Video Salient Object Detection, *Guanbin Li, Yuan Xie, Tianhao Wei, Keze Wang, Liang Lin*
49. [***M17***] Gated Fusion Network for Single Image Dehazing, *Wenqi Ren, Lin Ma, Jiawei Zhang, Jinshan Pan, Xiaochun Cao, Wei Liu, Ming-Hsuan Yang*
50. [***M20***] Learning a Single Convolutional Super-Resolution Network for Multiple Degradations, *Kai Zhang, Wangmeng Zuo, Lei Zhang*
51. [***N1***] Non-Blind Deblurring: Handling Kernel Uncertainty With CNNs, *Subeesh Vasu, Venkatesh Reddy Maligireddy, A. N. Rajagopalan*
52. [***N4***] Boundary Flow: A Siamese Network That Predicts Boundary Motion Without Training on Motion, *Peng Lei, Fuxin Li, Sinisa Todorovic*
53. [***N7***] Learning to See in the Dark, *Chen Chen, Qifeng Chen, Jia Xu, Vladlen Koltun*

Machine Learning for Computer Vision

54. [***N10***] BPGrad: Towards Global Optimality in Deep Learning via Branch and Pruning, *Ziming Zhang, Yuanwei Wu, Guanghui Wang*
55. [***N13***] Perturbative Neural Networks, *Felix Juefei-Xu, Vishnu Naresh Boddeti, Marios Savvides*
56. [***N16***] Unsupervised Correlation Analysis, *Yedid Hoshen, Lior Wolf*
57. [***N19***] A Biresolution Spectral Framework for Product Quantization, *Lopamudra Mukherjee, Sathya N. Ravi, Jiming Peng, Vikas Singh*
58. [***N22***] Domain Adaptive Faster R-CNN for Object Detection in the Wild, *Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, Luc Van Gool*
59. [***O3***] Low-Shot Learning With Large-Scale Diffusion, *Matthijs Douze, Arthur Szlam, Bharath Hariharan, Hervé Jégou*
60. [***O6***] Joint Pose and Expression Modeling for Facial Expression Recognition, *Feifei Zhang, Tianzhu Zhang, Qirong Mao, Changsheng Xu*
61. [***O9***] Lightweight Probabilistic Deep Networks, *Jochen Gast, Stefan Roth*
62. [***O12***] Adversarially Learned One-Class Classifier for Novelty Detection, *Mohammad Sabokrou, Mohammad Khalooei, Mahmood Fathy, Ehsan Adeli*
63. [***O15***] Defense Against Universal Adversarial Perturbations, *Naveed Akhtar, Jian Liu, Ajmal Mian*
64. [***O18***] Disentangling Factors of Variation by Mixing Them, *Qiyang Hu, Attila Szabó, Tiziano Portenier, Paolo Favaro, Matthias Zwicker*
65. [***O21***] Deformable GANs for Pose-Based Human Image Generation, *Aliaksandr Siarohin, Enver Sangineto, Stéphane Lathuilière, Nicu Sebe*
66. [***P2***] Hierarchical Recurrent Attention Networks for Structured Online Maps, *Namdar Homayounfar, Wei-Chiu Ma, Shrinidhi Kowshika Lakshmikanth, Raquel Urtasun*
67. [***P5***] Sliced Wasserstein Distance for Learning Gaussian Mixture Models, *Soheil Kolouri, Gustavo K. Rohde, Heiko Hoffmann*
68. [***P8***] Aligning Infinite-Dimensional Covariance Matrices in Reproducing Kernel Hilbert Spaces for Domain Adaptation, *Zhen Zhang, Mianzhi Wang, Yan Huang, Arye Nehorai*
69. [***P11***] CLEAR: Cumulative LEARning for One-Shot One-Class Image Recognition, *Jedrzej Kozerawski, Matthew Turk*
70. [***P14***] Local and Global Optimization Techniques in Graph-Based Clustering, *Daiki Ikami, Toshihiko Yamasaki, Kiyoharu Aizawa*

71. [***P17***] Multi-Task Learning by Maximizing Statistical Dependence, *Youssef A. Mejjati, Darren Cosker, Kwang In Kim*
72. [***P20***] Robust Classification With Convolutional Prototype Learning, *Hong-Ming Yang, Xu-Yao Zhang, Fei Yin, Cheng-Lin Liu*
73. [***Q1***] Generative Modeling Using the Sliced Wasserstein Distance, *Ishan Deshpande, Ziyu Zhang, Alexander G. Schwing*
74. [***Q4***] Learning Time/Memory-Efficient Deep Architectures With Budgeted Super Networks, *Tom Véniat, Ludovic Denoyer*
75. [***Q7***] Cross-View Image Synthesis Using Conditional GANs, *Krishna Regmi, Ali Borji*

Applications

76. [***Q10***] Sparse, Smart Contours to Represent and Edit Images, *Tali Dekel, Chuang Gan, Dilip Krishnan, Ce Liu, William T. Freeman*
77. [***Q13***] Anticipating Traffic Accidents With Adaptive Loss and Large-Scale Incident DB, *Tomoyuki Suzuki, Hirokatsu Kataoka, Yoshimitsu Aoki, Yutaka Satoh*
78. [***Q16***] A Minimalist Approach to Type-Agnostic Detection of Quadrics in Point Clouds, *Tolga Birdal, Benjamin Busam, Nassir Navab, Slobodan Ilic, Peter Sturm*
79. [***Q19***] Facelet-Bank for Fast Portrait Manipulation, *Ying-Cong Chen, Huaijia Lin, Michelle Shu, Ruiyu Li, Xin Tao, Xiaoyong Shen, Yangang Ye, Jiaya Jia*
80. [***Q22***] Visual to Sound: Generating Natural Sound for Videos in the Wild, *Yipin Zhou, Zhaowen Wang, Chen Fang, Trung Bui, Tamara L. Berg*

1930–2200 Social (Gallivan Center)

Directions: The Gallivan Center is located at 239 Main Street—two blocks away from the Salt Palace Convention Center. However, you will need to enter the plaza from the North side (on 200 South). See map below.

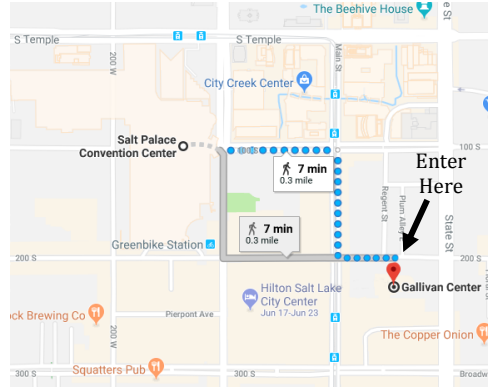

Notes:

Wednesday, June 20

0730–1830 Registration (South Lobby)

0730–0900 Breakfast (Hall A; Halls 1-4)

0800–1000 Poster Setup (Halls C-E)

0830-1010 Session 2-1A: Object Recognition & Scene Understanding II (Ballroom)

Papers in this session are also in Poster Session P2-1. Poster tag in square brackets (e.g., [***A1***])

Chairs: Michael Maire *(TTI Chicago)*
Subhransu Maji *(Univ. of Massachusetts, Amherst)*

0830 Orals (O2-1A)

Format (12 min. for presentation + 2 min. for questions)

1. [***A1***] 3D-RCNN: Instance-Level 3D Object Reconstruction via Render-and-Compare, *Abhijit Kundu, Yin Li, James M. Rehg*
2. [***A4***] Fast and Furious: Real Time End-to-End 3D Detection, Tracking and Motion Forecasting With a Single Convolutional Net, *Wenjie Luo, Bin Yang, Raquel Urtasun*
3. [***A7***] An Analysis of Scale Invariance in Object Detection - SNIP, *Bharat Singh, Larry S. Davis*
4. [***A10***] Relation Networks for Object Detection, *Han Hu, Jiayuan Gu, Zheng Zhang, Jifeng Dai, Yichen Wei*

0928 Spotlights (S2-1A)

Format (4 min. for presentation; no questions)

1. [***B2***] Zero-Shot Sketch-Image Hashing, *Yuming Shen, Li Liu, Fumin Shen, Ling Shao*
2. [***B5***] VizWiz Grand Challenge: Answering Visual Questions From Blind People, *Danna Gurari, Qing Li, Abigale J. Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, Jeffrey P. Bigham*
3. [***B8***] Divide and Grow: Capturing Huge Diversity in Crowd Images With Incrementally Growing CNN, *Deepak Babu Sam, Neeraj N. Sajjan, R. Venkatesh Babu, Mukundhan Srinivasan*
4. [***B11***] Structured Set Matching Networks for One-Shot Part Labeling, *Jonghyun Choi, Jayant Krishnamurthy, Aniruddha Kembhavi, Ali Farhadi*
5. [***B14***] Self-Supervised Learning of Geometrically Stable Features Through Probabilistic Introspection, *David Novotny, Samuel Albanie, Diane Larlus, Andrea Vedaldi*
6. [***B17***] Link and Code: Fast Indexing With Graphs and Compact Regression Codes, *Matthijs Douze, Alexandre Sablayrolles, Hervé Jégou*
7. [***B20***] Textbook Question Answering Under Instructor Guidance With Memory Networks, *Juzheng Li, Hang Su, Jun Zhu, Siyu Wang, Bo Zhang*
8. [***C1***] Unsupervised Deep Generative Adversarial Hashing Network, *Kamran Ghasedi Dizaji, Feng Zheng, Najmeh Sadoughi, Yanhua Yang, Cheng Deng, Heng Huang*
9. [***C4***] Vision-and-Language Navigation: Interpreting Visually-Grounded Navigation Instructions in Real Environments, *Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, Anton van den Hengel*
10. [***C7***] DenseASPP for Semantic Segmentation in Street Scenes, *Maoke Yang, Kun Yu, Chi Zhang, Zhiwei Li, Kuiyuan Yang*

0830-1010 Session 2-1B: Machine Learning for Computer Vision III (Room 155)

Papers in this session are also in Poster Session P2-1. Poster tag in square brackets (e.g., [***C10***])

Chairs: Alex Schwing *(UIUC)*
Judy Hoffman *(Univ. of Califonia, Berkeley)*

0830 Orals (O2-1B)

Format (12 min. for presentation + 2 min. for questions)

1. [***C10***] Efficient Optimization for Rank-Based Loss Functions, *Pritish Mohapatra, Michal Rolínek, C.V. Jawahar, Vladimir Kolmogorov, M. Pawan Kumar*
2. [***C13***] Wasserstein Introspective Neural Networks, *Kwonjoon Lee, Weijian Xu, Fan Fan, Zhuowen Tu*
3. [***C16***] Taskonomy: Disentangling Task Transfer Learning, *Amir R. Zamir, Alexander Sax, William Shen, Leonidas J. Guibas, Jitendra Malik, Silvio Savarese*
4. [***C19***] Maximum Classifier Discrepancy for Unsupervised Domain Adaptation, *Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, Tatsuya Harada*

0928 Spotlights (S2-1B)

Format (4 min. for presentation; no questions)

1. [***C22***] Unsupervised Feature Learning via Non-Parametric Instance Discrimination, *Zhirong Wu, Yuanjun Xiong, Stella X. Yu, Dahua Lin*
2. [***D3***] Multi-Task Adversarial Network for Disentangled Feature Learning, *Yang Liu, Zhaowen Wang, Hailin Jin, Ian Wassell*
3. [***D6***] Learning From Synthetic Data: Addressing Domain Shift for Semantic Segmentation, *Swami Sankaranarayanan, Yogesh Balaji, Arpit Jain, Ser Nam Lim, Rama Chellappa*
4. [***D9***] Empirical Study of the Topology and Geometry of Deep Networks, *Alhussein Fawzi, Seyed-Mohsen Moosavi-Dezfooli, Pascal Frossard, Stefano Soatto*
5. [***D12***] Boosting Domain Adaptation by Discovering Latent Domains, *Massimiliano Mancini, Lorenzo Porzi, Samuel Rota Bulò, Barbara Caputo, Elisa Ricci*
6. [***D15***] Shape From Shading Through Shape Evolution, *Dawei Yang, Jia Deng*
7. [***D18***] Weakly Supervised Instance Segmentation Using Class Peak Response, *Yanzhao Zhou, Yi Zhu, Qixiang Ye, Qiang Qiu, Jianbin Jiao*
8. [***D21***] Collaborative and Adversarial Network for Unsupervised Domain Adaptation, *Weichen Zhang, Wanli Ouyang, Wen Li, Dong Xu*
9. [***E2***] Environment Upgrade Reinforcement Learning for Non-Differentiable Multi-Stage Pipelines, *Shuqin Xie, Zitian Chen, Chao Xu, Cewu Lu*
10. [***E5***] Teaching Categories to Human Learners With Visual Explanations, *Oisin Mac Aodha, Shihan Su, Yuxin Chen, Pietro Perona, Yisong Yue*

0830-1010 Session 2-1C: 3D Vision III

(Room 255)

Papers in this session are also in Poster Session P2-1. Poster tag in square brackets (e.g., [***E8***])

Chairs: Yaser Sheikh *(Carnegie Mellon Univ.)*
Yasutaka Furukawa *(Washington Univ. in St. Louis)*

0830 Orals (O2-1C)

Format (12 min. for presentation + 2 min. for questions)

1. [***E8***] Density Adaptive Point Set Registration, *Felix Järemo Lawin, Martin Danelljan, Fahad Shahbaz Khan, Per-Erik Forssén, Michael Felsberg*
2. [***E11***] Left-Right Comparative Recurrent Model for Stereo Matching, *Zequn Jie, Pengfei Wang, Yonggen Ling, Bo Zhao, Yunchao Wei, Jiashi Feng, Wei Liu*
3. [***E14***] Im2Pano3D: Extrapolating 360° Structure and Semantics Beyond the Field of View, *Shuran Song, Andy Zeng, Angel X. Chang, Manolis Savva, Silvio Savarese, Thomas Funkhouser*
4. [***E17***] Polarimetric Dense Monocular SLAM, *Luwei Yang, Feitong Tan, Ao Li, Zhaopeng Cui, Yasutaka Furukawa, Ping Tan*

0928 Spotlights (S2-1C)

Format (4 min. for presentation; no questions)

1. [***E20***] A Unifying Contrast Maximization Framework for Event Cameras, With Applications to Motion, Depth, and Optical Flow Estimation, *Guillermo Gallego, Henri Rebecq, Davide Scaramuzza*
2. [***F1***] Modeling Facial Geometry Using Compositional VAEs, *Timur Bagautdinov, Chenglei Wu, Jason Saragih, Pascal Fua, Yaser Sheikh*
3. [***F4***] Tangent Convolutions for Dense Prediction in 3D, *Maxim Tatarchenko, Jaesik Park, Vladlen Koltun, Qian-Yi Zhou*
4. [***F7***] RayNet: Learning Volumetric 3D Reconstruction With Ray Potentials, *Despoina Paschalidou, Osman Ulusoy, Carolin Schmitt, Luc Van Gool, Andreas Geiger*
5. [***F10***] Neural 3D Mesh Renderer, *Hiroharu Kato, Yoshitaka Ushiku, Tatsuya Harada*
6. [***F13***] Structured Attention Guided Convolutional Neural Fields for Monocular Depth Estimation, *Dan Xu, Wei Wang, Hao Tang, Hong Liu, Nicu Sebe, Elisa Ricci*

7. [***F16***] Automatic 3D Indoor Scene Modeling From Single Panorama, *Yang Yang, Shi Jin, Ruiyang Liu, Sing Bing Kang, Jingyi Yu*
8. [***F19***] Extreme 3D Face Reconstruction: Seeing Through Occlusions, *Anh Tuấn Trần, Tal Hassner, Iacopo Masi, Eran Paz, Yuval Nirkin, Gérard Medioni*
9. [***F22***] Beyond Gröbner Bases: Basis Selection for Minimal Solvers, *Viktor Larsson, Magnus Oskarsson, Kalle Åström, Alge Wallis, Zuzana Kukelova, Tomas Pajdla*
10. [***G3***] Lions and Tigers and Bears: Capturing Non-Rigid, 3D, Articulated Shape From Images, *Silvia Zuffi, Angjoo Kanazawa, Michael J. Black*

1000–1045 Morning Break (Halls A-C)

1000-1830 Exhibits (Halls D-E)

- See Exhibits map for list of exhibitors.

1010-1230 Demos (Hall C)

- Sim4CV: A Photo-Realistic Simulator for Computer Vision Applications, *Matthias Mueller, Silvio Giancola, Bernard Ghanem (KAUST)*
- Real-Time Face Anti-Spoofing System, *Yaojie Liu, Amin Jourabloo, Xiaoming Liu (Michigan State Univ.)*
- Efficient Annotation of Segmentation Datasets With Polygon-RNN++, *David Acuna, Huan Ling, Amlan Kar, Sanja Fidler (Univ. of Toronto)*
- Confocal Non-Line-Of-Sight Imaging, *Matthew O'Toole, David B. Lindell, Gordon Wetzstein (Stanford)*
- Deep Understanding of Live Soccer Matches, *Xu ShiKun, Gao Siyi, Zhu Yandong, Li Gen, Xu Zhenqi, Li Lei, Wang Changhu (Bytendance.inc)*
- Zillow 3D Home, *Qi Shan, Alex Colburn, Li Guan, Ivaylo Boyadzhiev, Pierre Moulon, Andrew Mullen (Zillow Group)*
- EPIC-KITCHENS 2018, *Dima Damen, Hazel Doughty, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Giovanni Maria Farinella, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, Michael Wray (Univ. of Bristol, Univ. of Toronto, Univ. of Catania)*
- Rolling Shutter Imaging on the Electric Grid, *Mark Sheinin, Yoav Y. Schechner, Kiriakos N. Kutulakos (Technion - Israel Inst. of Technology, U. of Toronto)*

1010-1230 Poster Session P2-1 (Halls C-E)

Poster tag in square brackets (e.g., [***G6***])

Object Recognition & Scene Understanding

1. [***G6***] Deep Cocktail Network: Multi-Source Unsupervised Domain Adaptation With Category Shift, *Ruijia Xu, Ziliang Chen, Wangmeng Zuo, Junjie Yan, Liang Lin*
2. [***G9***] DOTA: A Large-Scale Dataset for Object Detection in Aerial Images, *Gui-Song Xia, Xiang Bai, Jian Ding, Zhen Zhu, Serge Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, Liangpei Zhang*
3. [***G12***] Finding Beans in Burgers: Deep Semantic-Visual Embedding With Localization, *Martin Engilberge, Louis Chevallier, Patrick Pérez, Matthieu Cord*
4. [***G15***] Feature Super-Resolution: Make Machine See More Clearly, *Weimin Tan, Bo Yan, Bahetiyaer Bare*
5. [***G18***] ClusterNet: Detecting Small Objects in Large Scenes by Exploiting Spatio-Temporal Information, *Rodney LaLonde, Dong Zhang, Mubarak Shah*
6. [***G21***] MaskLab: Instance Segmentation by Refining Object Detection With Semantic and Direction Features, *Liang-Chieh Chen, Alexander Hermans, George Papandreou, Florian Schroff, Peng Wang, Hartwig Adam*
7. [***H2***] Hashing as Tie-Aware Learning to Rank, *Kun He, Fatih Cakir, Sarah Adel Bargal, Stan Sclaroff*
8. [***H5***] Classification-Driven Dynamic Image Enhancement, *Vivek Sharma, Ali Diba, Davy Neven, Michael S. Brown, Luc Van Gool, Rainer Stiefelhagen*
9. [***H8***] Knowledge Aided Consistency for Weakly Supervised Phrase Grounding, *Kan Chen, Jiyang Gao, Ram Nevatia*
10. [***H11***] Who Let the Dogs Out? Modeling Dog Behavior From Visual Data, *Kiana Ehsani, Hessam Bagherinezhad, Joseph Redmon, Roozbeh Mottaghi, Ali Farhadi*
11. [***H14***] Pseudo Mask Augmented Object Detection, *Xiangyun Zhao, Shuang Liang, Yichen Wei*
12. [***H17***] Dual Skipping Networks, *Changmao Cheng, Yanwei Fu, Yu-Gang Jiang, Wei Liu, Wenlian Lu, Jianfeng Feng, Xiangyang Xue*
13. [***H20***] Memory Matching Networks for One-Shot Image Recognition, *Qi Cai, Yingwei Pan, Ting Yao, Chenggang Yan, Tao Mei*

14. [*I1*] IQA: Visual Question Answering in Interactive Environments, *Daniel Gordon, Aniruddha Kembhavi, Mohammad Rastegari, Joseph Redmon, Dieter Fox, Ali Farhadi*
15. [*I4*] Pose Transferrable Person Re-Identification, *Jinxian Liu, Bingbing Ni, Yichao Yan, Peng Zhou, Shuo Cheng, Jianguo Hu*
16. [*I7*] Large Scale Fine-Grained Categorization and Domain-Specific Transfer Learning, *Yin Cui, Yang Song, Chen Sun, Andrew Howard, Serge Belongie*
17. [*I10*] Data Distillation: Towards Omni-Supervised Learning, *Ilija Radosavovic, Piotr Dollár, Ross Girshick, Georgia Gkioxari, Kaiming He*
18. [*I13*] Object Referring in Videos With Language and Human Gaze, *Arun Balajee Vasudevan, Dengxin Dai, Luc Van Gool*
19. [*I16*] Feature Selective Networks for Object Detection, *Yao Zhai, Jingjing Fu, Yan Lu, Houqiang Li*
20. [*I19*] Learning a Discriminative Filter Bank Within a CNN for Fine-Grained Recognition, *Yaming Wang, Vlad I. Morariu, Larry S. Davis*
21. [*I22*] Grounding Referring Expressions in Images by Variational Context, *Hanwang Zhang, Yulei Niu, Shih-Fu Chang*
22. [*J3*] Dynamic Graph Generation Network: Generating Relational Knowledge From Diagrams, *Daesik Kim, YoungJoon Yoo, Jee-Soo Kim, SangKuk Lee, Nojun Kwak*
23. [*J6*] A Network Architecture for Point Cloud Classification via Automatic Depth Images Generation, *Riccardo Roveri, Lukas Rahmann, Cengiz Oztireli, Markus Gross*
24. [*J9*] Towards Dense Object Tracking in a 2D Honeybee Hive, *Katarzyna Bozek, Laetitia Hebert, Alexander S. Mikheyev, Greg J. Stephens*
25. [*J12*] Long-Term On-Board Prediction of People in Traffic Scenes Under Uncertainty, *Apratim Bhattacharyya, Mario Fritz, Bernt Schiele*
26. [*J15*] Single-Shot Refinement Neural Network for Object Detection, *Shifeng Zhang, Longyin Wen, Xiao Bian, Zhen Lei, Stan Z. Li*
27. [*J18*] Video Captioning via Hierarchical Reinforcement Learning, *Xin Wang, Wenhu Chen, Jiawei Wu, Yuan-Fang Wang, William Yang Wang*
28. [*J21*] Tips and Tricks for Visual Question Answering: Learnings From the 2017 Challenge, *Damien Teney, Peter Anderson, Xiaodong He, Anton van den Hengel*
29. [*K2*] Learning to Segment Every Thing, *Ronghang Hu, Piotr Dollár, Kaiming He, Trevor Darrell, Ross Girshick*
30. [*K5*] Self-Supervised Adversarial Hashing Networks for Cross-Modal Retrieval, *Chao Li, Cheng Deng, Ning Li, Wei Liu, Xinbo Gao, Dacheng Tao*
31. [*K8*] Parallel Attention: A Unified Framework for Visual Object Discovery Through Dialogs and Queries, *Bohan Zhuang, Qi Wu, Chunhua Shen, Ian Reid, Anton van den Hengel*
32. [*K11*] Zigzag Learning for Weakly Supervised Object Detection, *Xiaopeng Zhang, Jiashi Feng, Hongkai Xiong, Qi Tian*
33. [*K14*] Attentive Fashion Grammar Network for Fashion Landmark Detection and Clothing Category Classification, *Wenguan Wang, Yuanlu Xu, Jianbing Shen, Song-Chun Zhu*

Machine Learning for Computer Vision

34. [*K17*] Generalized Zero-Shot Learning via Synthesized Examples, *Vinay Kumar Verma, Gundeep Arora, Ashish Mishra, Piyush Rai*
35. [*K20*] Partially Shared Multi-Task Convolutional Neural Network With Local Constraint for Face Attribute Learning, *Jiajiong Cao, Yingming Li, Zhongfei Zhang*
36. [*L1*] SYQ: Learning Symmetric Quantization for Efficient Deep Neural Networks, *Julian Faraone, Nicholas Fraser, Michaela Blott, Philip H.W. Leong*
37. [*L4*] DS*: Tighter Lifting-Free Convex Relaxations for Quadratic Matching Problems, *Florian Bernard, Christian Theobalt, Michael Moeller*
38. [*L7*] Deep Mutual Learning, *Ying Zhang, Tao Xiang, Timothy M. Hospedales, Huchuan Lu*
39. [*L10*] Coupled End-to-End Transfer Learning With Generalized Fisher Information, *Shixing Chen, Caojin Zhang, Ming Dong*
40. [*L13*] Residual Parameter Transfer for Deep Domain Adaptation, *Artem Rozantsev, Mathieu Salzmann, Pascal Fua*
41. [*L16*] High-Order Tensor Regularization With Application to Attribute Ranking, *Kwang In Kim, Juhyun Park, James Tompkin*
42. [*L19*] Learning to Localize Sound Source in Visual Scenes, *Arda Senocak, Tae-Hyun Oh, Junsik Kim, Ming-Hsuan Yang, In So Kweon*

CVPR 2018 At-a-Glance (Main Conference)

0730 0830 0845 0900 0915 0930 0945 1000 1015 1030 1045 1100 1115 1130 1145 1200 1215 1230 1245 1300 1315 1330 1345 1400 1415 1430 1445 1500 1515 1530 1545 1600 1615 1630 1645 1700 1715 1730 1745 1800 1830 1900 2000 2100

Registration (South Lobby): Same all 3 days

Poster Setup (Halls C-E): Same all 3 days

Exhibits (Halls D-E): Same all 3 days; see exhibit map handout for exhibitor list & locations

Social Event (The Gallivan Center): See map on pg. 21

Tuesday, June 19

Tutorial (Room 151 A-C & G): Building Deep Learning Applications on Big Data Platforms, (see other booklet)

Tutorial (Room 151 A-C & G): New from HoloLens – Research Mode, (see other booklet)

Breakfast (Hall A; Halls 1-4)

Oral Session 1-1A: Object Recognition & Scene Understanding I (Ballroom), pg. 6

Break (Halls A-C)

Lunch (Hall A; Halls 1-4)

Break (Halls A-C)

Welcome & Paper Awards (Ballroom)

Special Session: Workshop Competitions (Ballroom), pg. 6

Oral & Spotlight Session 1-1B: Analyzing Humans in Images I (Room 155), pg. 6

Oral & Spotlight Session 1-1C: 3D Vision I (Room 255), pg. 7

Demos (Hall C): pg. 7

Poster Session P1-1 (Halls C-E): pgs. 7-11

Poster Session P1-2 (Halls C-E): pgs. 12-16

Oral & Spotlight Session 1-2A: Machine Learning for Computer Vision I (Ballroom), pgs. 16-17

Oral & Spotlight Session 1-2B: 3D Vision II (Room 155), pg. 17

Oral & Spotlight Session 1-2C: Machine Learning for Computer Vision II (Room 255), pgs. 17-18

Demos (Hall C): pg. 18

Poster Session P1-3 (Halls C-E): pgs. 18-21

Wednesday, June 20

Breakfast

Break (Halls A-C)

Lunch (Hall A; Halls 1-4)

Tutorial (Room 151 A-C & G): Using Intel Deep Learning Deployment Tools for Algorithm Development and Productization, (see other booklet)

Break

Oral & Spotlight Session 2-1A: Object Recognition & Scene Understanding II (Ballroom), pg. 22

Oral & Spotlight Session 2-1B: Machine Learning for Computer Vision III (Room 155), pgs. 22-23

Oral & Spotlight Session 2-1C: 3D Vision III (Room 255), pgs. 23-24

Demos (Hall C): pg. 24

Poster Session P2-1 (Halls C-E): pgs. 24-27

Poster Session P2-2 (Halls C-E): pgs. 28-32

Oral & Spotlight Session 2-2A: Video Analytics (Ballroom), pg. 32

Oral & Spotlight Session 2-2B: Object Recognition & Scene Understanding III (Room 155), pg. 33

Oral & Spotlight Session 2-2C: Computational Photography (Room 255), pgs. 33-34

Demos (Hall C): pg. 34

Poster Session P2-3 (Halls C-E): pgs. 34-37

Thursday, June 21

Breakfast (Hall A; Halls 1-4)

Break (Halls 1-4)

Lunch (Halls 1-4)

Doctoral Consortium: pg. 43

Break (Halls 1-4)

Oral & Spotlight Session 3-1A: Object Recognition & Scene Understanding IV (Ballroom), pg. 38

Oral & Spotlight Session 3-1B: Analyzing Humans (Room 155), pgs. 38-39

Oral & Spotlight Session 3-1C: Applications (Room 255), pgs. 39-40

Poster Session P3-1 (Halls D-E): pgs. 40-42

Oral & Spotlight Session 3-2A: Analyzing Humans in Images II (Ballroom), pg. 44

Oral & Spotlight Session 3-2B: Machine Learning for Computer Vision IV (Room 155), pgs. 44-45

Oral & Spotlight Session 3-2C: Object Recognition & Scene Understanding V (Room 255), pgs. 45-46

Oral & Spotlight Session 3-3A: Machine Learning for Computer Vision V (Ballroom), pg. 46

Oral & Spotlight Session 3-3B: Image Motion & Tracking (Room 155), pg. 47

Oral & Spotlight Session 3-3C: Machine Learning for Computer Vision VI (Room 255), pgs. 47-48

Poster Session P3-2 (Halls D-E): pgs. 48-49

43. [***L22***] Dynamic Few-Shot Visual Learning Without Forgetting, *Spyros Gidaris, Nikos Komodakis*
44. [***M3***] Two-Step Quantization for Low-Bit Neural Networks, *Peisong Wang, Qinghao Hu, Yifan Zhang, Chunjie Zhang, Yang Liu, Jian Cheng*
45. [***M6***] Improved Lossy Image Compression With Priming and Spatially Adaptive Bit Rates for Recurrent Networks, *Nick Johnston, Damien Vincent, David Minnen, Michele Covell, Saurabh Singh, Troy Chinen, Sung Jin Hwang, Joel Shor, George Toderici*
46. [***M9***] Conditional Probability Models for Deep Image Compression, *Fabian Mentzer, Eirikur Agustsson, Michael Tschannen, Radu Timofte, Luc Van Gool*
47. [***M12***] Deep Diffeomorphic Transformer Networks, *Nicki Skafte Detlefsen, Oren Freifeld, Søren Hauberg*
48. [***M15***] The LováSz-Softmax Loss: A Tractable Surrogate for the Optimization of the Intersection-Over-Union Measure in Neural Networks, *Maxim Berman, Amal Rannen Triki, Matthew B. Blaschko*
49. [***M18***] Generative Adversarial Perturbations, *Omid Poursaeed, Isay Katsman, Bicheng Gao, Serge Belongie*
50. [***M21***] Learning Strict Identity Mappings in Deep Residual Networks, *Xin Yu, Zhiding Yu, Srikumar Ramalingam*
51. [***N2***] Geometric Robustness of Deep Networks: Analysis and Improvement, *Can Kanbak, Seyed-Mohsen Moosavi-Dezfooli, Pascal Frossard*
52. [***N5***] View Extrapolation of Human Body From a Single Image, *Hao Zhu, Hao Su, Peng Wang, Xun Cao, Ruigang Yang*
53. [***N8***] Geometry Aware Constrained Optimization Techniques for Deep Learning, *Soumava Kumar Roy, Zakaria Mhammedi, Mehrtash Harandi*
54. [***N11***] PointNetVLAD: Deep Point Cloud Based Retrieval for Large-Scale Place Recognition, *Mikaela Angelina Uy, Gim Hee Lee*
55. [***N14***] An Efficient and Provable Approach for Mixture Proportion Estimation Using Linear Independence Assumption, *Xiyu Yu, Tongliang Liu, Mingming Gong, Kayhan Batmanghelich, Dacheng Tao*
56. [***N17***] VoxelNet: End-to-End Learning for Point Cloud Based 3D Object Detection, *Yin Zhou, Oncel Tuzel*
57. [***N20***] Image to Image Translation for Domain Adaptation, *Zak Murez, Soheil Kolouri, David Kriegman, Ravi Ramamoorthi, Kyungnam Kim*
58. [***O1***] MobileNetV2: Inverted Residuals and Linear Bottlenecks, *Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, Liang-Chieh Chen*

3D Vision

59. [***O4***] Im2Struct: Recovering 3D Shape Structure From a Single RGB Image, *Chengjie Niu, Jun Li, Kai Xu*
60. [***O7***] Trust Your Model: Light Field Depth Estimation With Inline Occlusion Handling, *Hendrik Schilling, Maximilian Diebold, Carsten Rother, Bernd Jähne*
61. [***O10***] Baseline Desensitizing in Translation Averaging, *Bingbing Zhuang, Loong-Fah Cheong, Gim Hee Lee*
62. [***O13***] Mining Point Cloud Local Structures by Kernel Correlation and Graph Pooling, *Yiru Shen, Chen Feng, Yaoqing Yang, Dong Tian*
63. [***O16***] Large-Scale Point Cloud Semantic Segmentation With Superpoint Graphs, *Loic Landrieu, Martin Simonovsky*
64. [***O19***] Very Large-Scale Global SfM by Distributed Motion Averaging, *Siyu Zhu, Runze Zhang, Lei Zhou, Tianwei Shen, Tian Fang, Ping Tan, Long Quan*
65. [***O22***] ScanComplete: Large-Scale Scene Completion and Semantic Segmentation for 3D Scans, *Angela Dai, Daniel Ritchie, Martin Bokeloh, Scott Reed, Jürgen Sturm, Matthias Nießner*
66. [***P3***] Solving the Perspective-2-Point Problem for Flying-Camera Photo Composition, *Ziquan Lan, David Hsu, Gim Hee Lee*
67. [***P6***] Reflection Removal for Large-Scale 3D Point Clouds, *Jae-Seong Yun, Jae-Young Sim*
68. [***P9***] Attentional ShapeContextNet for Point Cloud Recognition, *Saining Xie, Sainan Liu, Zeyu Chen, Zhuowen Tu*
69. [***P12***] Geometry-Aware Deep Network for Single-Image Novel View Synthesis, *Miaomiao Liu, Xuming He, Mathieu Salzmann*
70. [***P15***] InverseFaceNet: Deep Monocular Inverse Face Rendering, *Hyeongwoo Kim, Michael Zollhöfer, Ayush Tewari, Justus Thies, Christian Richardt, Christian Theobalt*
71. [***P18***] Sparse Photometric 3D Face Reconstruction Guided by Morphable Models, *Xuan Cao, Zhang Chen, Anpei Chen, Xin Chen, Shiying Li, Jingyi Yu*
72. [***P21***] Texture Mapping for 3D Reconstruction With RGB-D Sensor, *Yanping Fu, Qingan Yan, Long Yang, Jie Liao, Chunxia Xiao*

73. [**Q2**] Learning Less Is More - 6D Camera Localization via 3D Surface Regression, *Eric Brachmann, Carsten Rother*
74. [**Q5**] Feature Mapping for Learning Fast and Accurate 3D Pose Inference From Synthetic Images, *Mahdi Rad, Markus Oberweger, Vincent Lepetit*
75. [**Q8**] Indoor RGB-D Compass From a Single Line and Plane, *Pyojin Kim, Brian Coltin, H. Jin Kim*
76. [**Q11**] Geometry-Aware Network for Non-Rigid Shape Prediction From a Single View, *Albert Pumarola, Antonio Agudo, Lorenzo Porzi, Alberto Sanfeliu, Vincent Lepetit, Francesc Moreno-Noguer*

Applications

77. [**Q14**] Sim2Real Viewpoint Invariant Visual Servoing by Recurrent Control, *Fereshteh Sadeghi, Alexander Toshev, Eric Jang, Sergey Levine*
78. [**Q17**] DocUNet: Document Image Unwarping via a Stacked U-Net, *Ke Ma, Zhixin Shu, Xue Bai, Jue Wang, Dimitris Samaras*
79. [**Q20**] Analysis of Hand Segmentation in the Wild, *Aisha Urooj, Ali Borji*
80. [**R1**] RoadTracer: Automatic Extraction of Road Networks From Aerial Images, *Favyen Bastani, Songtao He, Sofiane Abbar, Mohammad Alizadeh, Hari Balakrishnan, Sanjay Chawla, Sam Madden, David DeWitt*
81. [**R4**] Alternating-Stereo VINS: Observability Analysis and Performance Evaluation, *Mrinal K. Paul, Stergios I. Roumeliotis*
82. [**R7**] Soccer on Your Tabletop, *Konstantinos Rematas, Ira Kemelmacher-Shlizerman, Brian Curless, Steve Seitz*

1130–1330 Lunch (Hall A; Halls 1-4)

Notes:

1230-1450 Poster Session P2-2 (Halls C-E)

Poster tag in square brackets (e.g., [***A2***])

Computational Photography

1. [***A2***] EPINET: A Fully-Convolutional Neural Network Using Epipolar Geometry for Depth From Light Field Images, *Changha Shin, Hae-Gon Jeon, Youngjin Yoon, In So Kweon, Seon Joo Kim*
2. [***A5***] A Hybrid l_1-l_0 Layer Decomposition Model for Tone Mapping, *Zhetong Liang, Jun Xu, David Zhang, Zisheng Cao, Lei Zhang*
3. [***A8***] Deeply Learned Filter Response Functions for Hyperspectral Reconstruction, *Shijie Nie, Lin Gu, Yinqiang Zheng, Antony Lam, Nobutaka Ono, Imari Sato*
4. [***A22***] CRRN: Multi-Scale Guided Concurrent Reflection Removal Network, *Renjie Wan, Boxin Shi, Ling-Yu Duan, Ah-Hwee Tan, Alex C. Kot*
5. [***B3***] Single Image Reflection Separation With Perceptual Losses, *Xuaner Zhang, Ren Ng, Qifeng Chen*
6. [***B6***] A Robust Method for Strong Rolling Shutter Effects Correction Using Lines With Automatic Feature Selection, *Yizhen Lao, Omar Ait-Aider*
7. [***B9***] Time-Resolved Light Transport Decomposition for Thermal Photometric Stereo, *Kenichiro Tanaka, Nobuhiro Ikeya, Tsuyoshi Takatani, Hiroyuki Kubo, Takuya Funatomi, Yasuhiro Mukaigawa*

Image Motion & Tracking

8. [***B12***] Efficient Diverse Ensemble for Discriminative Co-Tracking, *Kourosh Meshgi, Shigeyuki Oba, Shin Ishii*
9. [***B15***] Rolling Shutter and Radial Distortion Are Features for High Frame Rate Multi-Camera Tracking, *Akash Bapat, True Price, Jan-Michael Frahm*
10. [***B18***] A Twofold Siamese Network for Real-Time Object Tracking, *Anfeng He, Chong Luo, Xinmei Tian, Wenjun Zeng*
11. [***B21***] Multi-Cue Correlation Filters for Robust Visual Tracking, *Ning Wang, Wengang Zhou, Qi Tian, Richang Hong, Meng Wang, Houqiang Li*
12. [***C2***] Learning Attentions: Residual Attentional Siamese Network for High Performance Online Visual Tracking, *Qiang Wang, Zhu Teng, Junliang Xing, Jin Gao, Weiming Hu, Stephen Maybank*
13. [***C5***] SINT++: Robust Visual Tracking via Adversarial Positive Instance Generation, *Xiao Wang, Chenglong Li, Bin Luo, Jin Tang*
14. [***C8***] High-Speed Tracking With Multi-Kernel Correlation Filters, *Ming Tang, Bin Yu, Fan Zhang, Jinqiao Wang*
15. [***C11***] Occlusion Aware Unsupervised Learning of Optical Flow, *Yang Wang, Yi Yang, Zhenheng Yang, Liang Zhao, Peng Wang, Wei Xu*
16. [***C14***] Revisiting Video Saliency: A Large-Scale Benchmark and a New Model, *Wenguan Wang, Jianbing Shen, Fang Guo, Ming-Ming Cheng, Ali Borji*
17. [***C17***] Learning Spatial-Temporal Regularized Correlation Filters for Visual Tracking, *Feng Li, Cheng Tian, Wangmeng Zuo, Lei Zhang, Ming-Hsuan Yang*

Object Recognition & Scene Understanding

18. [***C20***] Multimodal Visual Concept Learning With Weakly Supervised Techniques, *Giorgos Bouritsas, Petros Koutras, Athanasia Zlatintsi, Petros Maragos*
19. [***D1***] Efficient Large-Scale Approximate Nearest Neighbor Search on OpenCL FPGA, *Jialiang Zhang, Soroosh Khoram, Jing Li*
20. [***D4***] Learning a Complete Image Indexing Pipeline, *Himalaya Jain, Joaquin Zepeda, Patrick Pérez, Rémi Gribonval*
21. [***D7***] Transparency by Design: Closing the Gap Between Performance and Interpretability in Visual Reasoning, *David Mascharka, Philip Tran, Ryan Soklaski, Arjun Majumdar*
22. [***D10***] Fooling Vision and Language Models Despite Localization and Attention Mechanism, *Xiaojun Xu, Xinyun Chen, Chang Liu, Anna Rohrbach, Trevor Darrell, Dawn Song*
23. [***D13***] Categorizing Concepts With Basic Level for Vision-to-Language, *Hanzhang Wang, Hanli Wang, Kaisheng Xu*
24. [***D16***] Don't Just Assume; Look and Answer: Overcoming Priors for Visual Question Answering, *Aishwarya Agrawal, Dhruv Batra, Devi Parikh, Aniruddha Kembhavi*
25. [***D19***] Learning Pixel-Level Semantic Affinity With Image-Level Supervision for Weakly Supervised Semantic Segmentation, *Jiwoon Ahn, Suha Kwak*
26. [***D22***] From Lifestyle Vlogs to Everyday Interactions, *David F. Fouhey, Wei-cheng Kuo, Alexei A. Efros, Jitendra Malik*
27. [***E3***] Cross-Domain Weakly-Supervised Object Detection Through Progressive Domain Adaptation, *Naoto Inoue, Ryosuke Furuta, Toshihiko Yamasaki, Kiyoharu Aizawa*

28. [***E6***] RotationNet: Joint Object Categorization and Pose Estimation Using Multiviews From Unsupervised Viewpoints, *Asako Kanezaki, Yasuyuki Matsushita, Yoshifumi Nishida*
29. [***E9***] An End-to-End TextSpotter With Explicit Alignment and Attention, *Tong He, Zhi Tian, Weilin Huang, Chunhua Shen, Yu Qiao, Changming Sun*

Analyzing Humans in Images

30. [***E12***] WILDTRACK: A Multi-Camera HD Dataset for Dense Unscripted Pedestrian Detection, *Tatjana Chavdarova, Pierre Baqué, Stéphane Bouquet, Andrii Maksai, Cijo Jose, Timur Bagautdinov, Louis Lettry, Pascal Fua, Luc Van Gool, François Fleuret*
31. [***E15***] Direct Shape Regression Networks for End-to-End Face Alignment, *Xin Miao, Xiantong Zhen, Xianglong Liu, Cheng Deng, Vassilis Athitsos, Heng Huang*
32. [***E18***] Natural and Effective Obfuscation by Head Inpainting, *Qianru Sun, Liqian Ma, Seong Joon Oh, Luc Van Gool, Bernt Schiele, Mario Fritz*
33. [***E21***] 3D Semantic Trajectory Reconstruction From 3D Pixel Continuum, *Jae Shin Yoon, Ziwei Li, Hyun Soo Park*
34. [***F2***] Optimizing Filter Size in Convolutional Neural Networks for Facial Action Unit Recognition, *Shizhong Han, Zibo Meng, Zhiyuan Li, James O'Reilly, Jie Cai, Xiaofeng Wang, Yan Tong*
35. [***F5***] V2V-PoseNet: Voxel-to-Voxel Prediction Network for Accurate 3D Hand and Human Pose Estimation From a Single Depth Map, *Gyeongsik Moon, Ju Yong Chang, Kyoung Mu Lee*
36. [***F8***] Ring Loss: Convex Feature Normalization for Face Recognition, *Yutong Zheng, Dipan K. Pal, Marios Savvides*
37. [***F11***] Adversarially Occluded Samples for Person Re-Identification, *Houjing Huang, Dangwei Li, Zhang Zhang, Xiaotang Chen, Kaiqi Huang*
38. [***F14***] Classifier Learning With Prior Probabilities for Facial Action Unit Recognition, *Yong Zhang, Weiming Dong, Bao-Gang Hu, Qiang Ji*
39. [***F17***] 4DFAB: A Large Scale 4D Database for Facial Expression Analysis and Biometric Applications, *Shiyang Cheng, Irene Kotsia, Maja Pantic, Stefanos Zafeiriou*
40. [***F20***] Seeing Small Faces From Robust Anchor's Perspective, *Chenchen Zhu, Ran Tao, Khoa Luu, Marios Savvides*
41. [***G1***] 2D/3D Pose Estimation and Action Recognition Using Multitask Deep Learning, *Diogo C. Luvizon, David Picard, Hedi Tabia*
42. [***G4***] Dense 3D Regression for Hand Pose Estimation, *Chengde Wan, Thomas Probst, Luc Van Gool, Angela Yao*
43. [***G7***] Camera Style Adaptation for Person Re-Identification, *Zhun Zhong, Liang Zheng, Zhedong Zheng, Shaozi Li, Yi Yang*
44. [***G10***] PoseTrack: A Benchmark for Human Pose Estimation and Tracking, *Mykhaylo Andriluka, Umar Iqbal, Eldar Insafutdinov, Leonid Pishchulin, Anton Milan, Juergen Gall, Bernt Schiele*
45. [***G13***] Exploit the Unknown Gradually: One-Shot Video-Based Person Re-Identification by Stepwise Learning, *Yu Wu, Yutian Lin, Xuanyi Dong, Yan Yan, Wanli Ouyang, Yi Yang*
46. [***G16***] Pose-Robust Face Recognition via Deep Residual Equivariant Mapping, *Kaidi Cao, Yu Rong, Cheng Li, Xiaoou Tang, Chen Change Loy*
47. [***G19***] DecideNet: Counting Varying Density Crowds Through Attention Guided Detection and Density Estimation, *Jiang Liu, Chenqiang Gao, Deyu Meng, Alexander G. Hauptmann*
48. [***G22***] LSTM Pose Machines, *Yue Luo, Jimmy Ren, Zhouxia Wang, Wenxiu Sun, Jinshan Pan, Jianbo Liu, Jiahao Pang, Liang Lin*
49. [***H3***] Disentangling Features in 3D Face Shapes for Joint Face Reconstruction and Recognition, *Feng Liu, Ronghang Zhu, Dan Zeng, Qijun Zhao, Xiaoming Liu*
50. [***H6***] Convolutional Sequence to Sequence Model for Human Dynamics, *Chen Li, Zhen Zhang, Wee Sun Lee, Gim Hee Lee*
51. [***H9***] Gesture Recognition: Focus on the Hands, *Pradyumna Narayana, Ross Beveridge, Bruce A. Draper*
52. [***H12***] Crowd Counting via Adversarial Cross-Scale Consistency Pursuit, *Zan Shen, Yi Xu, Bingbing Ni, Minsi Wang, Jianguo Hu, Xiaokang Yang*
53. [***H15***] 3D Human Pose Estimation in the Wild by Adversarial Learning, *Wei Yang, Wanli Ouyang, Xiaolong Wang, Jimmy Ren, Hongsheng Li, Xiaogang Wang*
54. [***H18***] CosFace: Large Margin Cosine Loss for Deep Face Recognition, *Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, Wei Liu*
55. [***H21***] Encoding Crowd Interaction With Deep Neural Network for Pedestrian Trajectory Prediction, *Yanyu Xu, Zhixin Piao, Shenghua Gao*

56. [***I2***] Mean-Variance Loss for Deep Age Estimation From a Face, *Hongyu Pan, Hu Han, Shiguang Shan, Xilin Chen*
57. [***I5***] Probabilistic Joint Face-Skull Modelling for Facial Reconstruction, *Dennis Madsen, Marcel Lüthi, Andreas Schneider, Thomas Vetter*
58. [***I8***] Learning Latent Super-Events to Detect Multiple Activities in Videos, *AJ Piergiovanni, Michael S. Ryoo*
59. [***I11***] Temporal Hallucinating for Action Recognition With Few Still Images, *Yali Wang, Lei Zhou, Yu Qiao*
60. [***I14***] Deep Progressive Reinforcement Learning for Skeleton-Based Action Recognition, *Yansong Tang, Yi Tian, Jiwen Lu, Peiyang Li, Jie Zhou*
61. [***I17***] Gaze Prediction in Dynamic 360° Immersive Videos, *Yanyu Xu, Yanbing Dong, Junru Wu, Zhengzhong Sun, Zhiru Shi, Jingyi Yu, Shenghua Gao*
62. [***I20***] When Will You Do What? - Anticipating Temporal Occurrences of Activities, *Yazan Abu Farha, Alexander Richard, Juergen Gall*
63. [***J1***] Fusing Crowd Density Maps and Visual Object Trackers for People Tracking in Crowd Scenes, *Weihong Ren, Di Kang, Yandong Tang, Antoni B. Chan*
64. [***J4***] Dual Attention Matching Network for Context-Aware Feature Sequence Based Person Re-Identification, *Jianlou Si, Honggang Zhang, Chun-Guang Li, Jason Kuen, Xiangfei Kong, Alex C. Kot, Gang Wang*
65. [***J7***] Easy Identification From Better Constraints: Multi-Shot Person Re-Identification From Reference Constraints, *Jiahuan Zhou, Bing Su, Ying Wu*
66. [***J10***] Crowd Counting With Deep Negative Correlation Learning, *Zenglin Shi, Le Zhang, Yun Liu, Xiaofeng Cao, Yangdong Ye, Ming-Ming Cheng, Guoyan Zheng*
67. [***J13***] Human Appearance Transfer, *Mihai Zanfir, Alin-Ionut Popa, Andrei Zanfir, Cristian Sminchisescu*

Machine Learning for Computer Vision

68. [***J16***] Domain Generalization With Adversarial Feature Learning, *Haoliang Li, Sinno Jialin Pan, Shiqi Wang, Alex C. Kot*
69. [***J19***] Pyramid Stereo Matching Network, *Jia-Ren Chang, Yong-Sheng Chen*
70. [***J22***] Event-Based Vision Meets Deep Learning on Steering Prediction for Self-Driving Cars, *Ana I. Maqueda, Antonio Loquercio, Guillermo Gallego, Narciso García, Davide Scaramuzza*
71. [***K3***] Learning Answer Embeddings for Visual Question Answering, *Hexiang Hu, Wei-Lun Chao, Fei Sha*
72. [***K6***] Good View Hunting: Learning Photo Composition From Dense View Pairs, *Zijun Wei, Jianming Zhang, Xiaohui Shen, Zhe Lin, Radomír Mech, Minh Hoai, Dimitris Samaras*
73. [***K9***] CleanNet: Transfer Learning for Scalable Image Classifier Training With Label Noise, *Kuang-Huei Lee, Xiaodong He, Lei Zhang, Linjun Yang*
74. [***K12***] Independently Recurrent Neural Network (IndRNN): Building a Longer and Deeper RNN, *Shuai Li, Wanqing Li, Chris Cook, Ce Zhu, Yanbo Gao*
75. [***K15***] Mix and Match Networks: Encoder-Decoder Alignment for Zero-Pair Image Translation, *Yaxing Wang, Joost van de Weijer, Luis Herranz*
76. [***K18***] Structured Uncertainty Prediction Networks, *Garoe Dorta, Sara Vicente, Lourdes Agapito, Neill D. F. Campbell, Ivor Simpson*
77. [***K21***] Between-Class Learning for Image Classification, *Yuji Tokozume, Yoshitaka Ushiku, Tatsuya Harada*
78. [***L2***] Adversarial Feature Augmentation for Unsupervised Domain Adaptation, *Riccardo Volpi, Pietro Morerio, Silvio Savarese, Vittorio Murino*
79. [***L5***] Generative Image Inpainting With Contextual Attention, *Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, Thomas S. Huang*
80. [***L8***] CSGNet: Neural Shape Parser for Constructive Solid Geometry, *Gopal Sharma, Rishabh Goyal, Difan Liu, Evangelos Kalogerakis, Subhransu Maji*
81. [***L11***] Conditional Image-to-Image Translation, *Jianxin Lin, Yingce Xia, Tao Qin, Zhibo Chen, Tie-Yan Liu*
82. [***L14***] Continuous Relaxation of MAP Inference: A Nonconvex Perspective, *D. Khuê Lê-Huu, Nikos Paragios*
83. [***L17***] Feature Generating Networks for Zero-Shot Learning, *Yongqin Xian, Tobias Lorenz, Bernt Schiele, Zeynep Akata*
84. [***L20***] Joint Optimization Framework for Learning With Noisy Labels, *Daiki Tanaka, Daiki Ikami, Toshihiko Yamasaki, Kiyoharu Aizawa*
85. [***M1***] Convolutional Image Captioning, *Jyoti Aneja, Aditya Deshpande, Alexander G. Schwing*
86. [***M4***] AON: Towards Arbitrarily-Oriented Text Recognition, *Zhanzhan Cheng, Yangliu Xu, Fan Bai, Yi Niu, Shiliang Pu, Shuigeng Zhou*

87. [***M7***] Wrapped Gaussian Process Regression on Riemannian Manifolds, *Anton Mallasto, Aasa Feragen*
88. [***M10***] Geometry Guided Convolutional Neural Networks for Self-Supervised Video Representation Learning, *Chuang Gan, Boqing Gong, Kun Liu, Hao Su, Leonidas J. Guibas*
89. [***M13***] DiverseNet: When One Right Answer Is Not Enough, *Michael Firman, Neill D. F. Campbell, Lourdes Agapito, Gabriel J. Brostow*
90. [***M16***] Deep Face Detector Adaptation Without Negative Transfer or Catastrophic Forgetting, *Muhammad Abdullah Jamal, Haoxiang Li, Boqing Gong*
91. [***M19***] Analyzing Filters Toward Efficient ConvNet, *Takumi Kobayashi*
92. [***M22***] Regularizing Deep Networks by Modeling and Predicting Label Structure, *Mohammadreza Mostajabi, Michael Maire, Gregory Shakhnarovich*
93. [***N3***] In-Place Activated BatchNorm for Memory-Optimized Training of DNNs, *Samuel Rota Bulò, Lorenzo Porzi, Peter Kontschieder*
94. [***N6***] DVQA: Understanding Data Visualizations via Question Answering, *Kushal Kafle, Brian Price, Scott Cohen, Christopher Kanan*
95. [***N9***] DA-GAN: Instance-Level Image Translation by Deep Attention Generative Adversarial Networks, *Shuang Ma, Jianlong Fu, Chang Wen Chen, Tao Mei*
96. [***N12***] Unsupervised Learning of Depth and Ego-Motion From Monocular Video Using 3D Geometric Constraints, *Reza Mahjourian, Martin Wicke, Anelia Angelova*

Object Recognition & Scene Understanding

97. [**N15**] FOTS: Fast Oriented Text Spotting With a Unified Network, Xuebo Liu, Ding Liang, Shi Yan, Dagui Chen, Yu Qiao, Junjie Yan
98. [**N18**] Mobile Video Object Detection With Temporally-Aware Feature Maps, Mason Liu, Menglong Zhu
99. [**N21**] Weakly Supervised Phrase Localization With Multi-Scale Anchored Transformer Network, Fang Zhao, Jianshu Li, Jian Zhao, Jiashi Feng
100. [**O2**] Revisiting Oxford and Paris: Large-Scale Image Retrieval Benchmarking, Filip Radenović, Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, Ondřej Chum
101. [**O5**] Cross-Dataset Adaptation for Visual Question Answering, Wei-Lun Chao, Hexiang Hu, Fei Sha
102. [**O8**] Globally Optimal Inlier Set Maximization for Atlanta Frame Estimation, Kyungdon Joo, Tae-Hyun Oh, In So Kweon, Jean-Charles Bazin
103. [**O11**] End-to-End Convolutional Semantic Embeddings, Quanzeng You, Zhengyou Zhang, Jiebo Luo
104. [**O14**] Referring Image Segmentation via Recurrent Refinement Networks, Ruiyu Li, Kaican Li, Yi-Chun Kuo, Michelle Shu, Xiaojuan Qi, Xiaoyong Shen, Jiaya Jia
105. [**O17**] Two Can Play This Game: Visual Dialog With Discriminative Question Generation and Answering, Unnat Jain, Svetlana Lazebnik, Alexander G. Schwing
106. [**O20**] Generative Adversarial Learning Towards Fast Weakly Supervised Detection, Yunhan Shen, Rongrong Ji, Shengchuan Zhang, Wangmeng Zuo, Yan Wang
107. [**P1**] A Deeper Look at Power Normalizations, Piotr Koniusz, Hongguang Zhang, Fatih Porikli
108. [**P4**] Dimensionality's Blessing: Clustering Images by Underlying Distribution, Wen-Yan Lin, Siying Liu, Jian-Huang Lai, Yasuyuki Matsushita
109. [**P7**] Eliminating Background-Bias for Robust Person Re-Identification, Maoqing Tian, Shuai Yi, Hongsheng Li, Shihua Li, Xuesen Zhang, Jianping Shi, Junjie Yan, Xiaogang Wang
110. [**P10**] Learning to Evaluate Image Captioning, Yin Cui, Guandao Yang, Andreas Veit, Xun Huang, Serge Belongie
111. [**P13**] Single-Shot Object Detection With Enriched Semantics, Zhishuai Zhang, Siyuan Qiao, Cihang Xie, Wei Shen, Bo Wang, Alan L. Yuille
112. [**P16**] Low-Shot Learning With Imprinted Weights, Hang Qi, Matthew Brown, David G. Lowe
113. [**P19**] Neural Motifs: Scene Graph Parsing With Global Context, Rowan Zellers, Mark Yatskar, Sam Thomson, Yejin Choi

Applications

114. [***P22***] Variational Autoencoders for Deforming 3D Mesh Models, *Qingyang Tan, Lin Gao, Yu-Kun Lai, Shihong Xia*
115. [***Q3***] Fast Monte-Carlo Localization on Aerial Vehicles Using Approximate Continuous Belief Representations, *Aditya Dhawale, Kumar Shaurya Shankar, Nathan Michael*
116. [***Q6***] DeLS-3D: Deep Localization and Segmentation With a 3D Semantic Map, *Peng Wang, Ruigang Yang, Binbin Cao, Wei Xu, Yuanqing Lin*

117. [***Q9***] LiDAR-Video Driving Dataset: Learning Driving Policies Effectively, *Yiping Chen, Jingkang Wang, Jonathan Li, Cewu Lu, Zhipeng Luo, Han Xue, Cheng Wang*
118. [***Q12***] Logo Synthesis and Manipulation With Clustered Generative Adversarial Networks, *Alexander Sage, Eirikur Agustsson, Radu Timofte, Luc Van Gool*
119. [***Q15***] Egocentric Basketball Motion Planning From a Single First-Person Image, *Gedas Bertasius, Aaron Chan, Jianbo Shi*
120. [***Q18***] Human-Centric Indoor Scene Synthesis Using Stochastic Grammar, *Siyuan Qi, Yixin Zhu, Siyuan Huang, Chenfanfu Jiang, Song-Chun Zhu*
121. [***Q21***] Rotation-Sensitive Regression for Oriented Scene Text Detection, *Minghui Liao, Zhen Zhu, Baoguang Shi, Gui-song Xia, Xiang Bai*
122. [***R2***] Separating Self-Expression and Visual Content in Hashtag Supervision, *Andreas Veit, Maximilian Nickel, Serge Belongie, Laurens van der Maaten*
123. [***R5***] Distort-and-Recover: Color Enhancement Using Deep Reinforcement Learning, *Jongchan Park, Joon-Young Lee, Donggeun Yoo, In So Kweon*

1420–1450 Afternoon Break (Halls A-C)

1450-1630 Session 2-2A: Video Analytics (Ballroom)

Papers in this session are also in Poster Session P2-3. Poster tag in square brackets (e.g., [***A3***])

Chairs: Karteek Alahari *(Inria Grenoble)*
Carl Vondrick *(Columbia Univ.)*

1450 Orals (O2-2A)

Format (12 min. for presentation + 2 min. for questions)

1. [***A3***] Im2Flow: Motion Hallucination From Static Images for Action Recognition, *Ruohan Gao, Bo Xiong, Kristen Grauman*
2. [***A6***] Finding "It": Weakly-Supervised Reference-Aware Visual Grounding in Instructional Videos, *De-An Huang, Shyamal Buch, Lucio Dery, Animesh Garg, Li Fei-Fei, Juan Carlos Niebles*
3. [***A9***] Actor and Action Video Segmentation From a Sentence, *Kirill Gavrilyuk, Amir Ghodrati, Zhenyang Li, Cees G. M. Snoek*
4. [***B1***] Egocentric Activity Recognition on a Budget, *Rafael Possas, Sheila Pinto Caceres, Fabio Ramos*

1548 Spotlights (S2-2A)

Format (4 min. for presentation; no questions)

1. [***B4***] CNN in MRF: Video Object Segmentation via Inference in a CNN-Based Higher-Order Spatio-Temporal MRF, *Linchao Bao, Baoyuan Wu, Wei Liu*
2. [***B7***] Action Sets: Weakly Supervised Action Segmentation Without Ordering Constraints, *Alexander Richard, Hilde Kuehne, Juergen Gall*
3. [***B10***] Low-Latency Video Semantic Segmentation, *Yule Li, Jianping Shi, Dahua Lin*
4. [***B13***] Fine-Grained Video Captioning for Sports Narrative, *Huanyu Yu, Shuo Cheng, Bingbing Ni, Minsi Wang, Jian Zhang, Xiaokang Yang*
5. [***B16***] End-to-End Learning of Motion Representation for Video Understanding, *Lijie Fan, Wenbing Huang, Chuang Gan, Stefano Ermon, Boqing Gong, Junzhou Huang*
6. [***B19***] Compressed Video Action Recognition, *Chao-Yuan Wu, Manzil Zaheer, Hexiang Hu, R. Manmatha, Alexander J. Smola, Philipp Krähenbühl*
7. [***B22***] Features for Multi-Target Multi-Camera Tracking and Re-Identification, *Ergys Ristani, Carlo Tomasi*
8. [***C3***] AVA: A Video Dataset of Spatio-Temporally Localized Atomic Visual Actions, *Chunhui Gu, Chen Sun, David A. Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, Cordelia Schmid, Jitendra Malik*
9. [***C6***] Who's Better? Who's Best? Pairwise Deep Ranking for Skill Determination, *Hazel Doughty, Dima Damen, Walterio Mayol-Cuevas*
10. [***C9***] MX-LSTM: Mixing Tracklets and Vislets to Jointly Forecast Trajectories and Head Poses, *Irtiza Hasan, Francesco Setti, Theodore Tsesmelis, Alessio Del Bue, Fabio Galasso, Marco Cristani*

1450-1630 Session 2-2B: Object Recognition & Scene Understanding III (Room 155)

Papers in this session are also in Poster Session P2-3. Poster tag in square brackets (e.g., [***C12***])

Chairs: Abhinav Shrivastava *(Carnegie Mellon Univ.)*
Iasonas Kokkinos *(Facebook AI Research)*

1450 Orals (O2-2B)

Format (12 min. for presentation + 2 min. for questions)

1. [***C12***] Bottom-Up and Top-Down Attention for Image Captioning and Visual Question Answering, *Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, Lei Zhang*
2. [***C15***] Improved Fusion of Visual and Language Representations by Dense Symmetric Co-Attention for Visual Question Answering, *Duy-Kien Nguyen, Takayuki Okatani*
3. [***C18***] FLIPDIAL: A Generative Model for Two-Way Visual Dialogue, *Daniela Massiceti, N. Siddharth, Puneet K. Dokania, Philip H.S. Torr*
4. [***C21***] Are You Talking to Me? Reasoned Visual Dialog Generation Through Adversarial Learning, *Qi Wu, Peng Wang, Chunhua Shen, Ian Reid, Anton van den Hengel*

1548 Spotlights (S2-2B)

Format (4 min. for presentation; no questions)

1. [***D2***] Visual Question Generation as Dual Task of Visual Question Answering, *Yikang Li, Nan Duan, Bolei Zhou, Xiao Chu, Wanli Ouyang, Xiaogang Wang, Ming Zhou*
2. [***D5***] Unsupervised Textual Grounding: Linking Words to Image Concepts, *Raymond A. Yeh, Minh N. Do, Alexander G. Schwing*
3. [***D8***] Focal Visual-Text Attention for Visual Question Answering, *Junwei Liang, Lu Jiang, Liangliang Cao, Li-Jia Li, Alexander G. Hauptmann*
4. [***D11***] SeGAN: Segmenting and Generating the Invisible, *Kiana Ehsani, Roozbeh Mottaghi, Ali Farhadi*
5. [***D14***] Cascade R-CNN: Delving Into High Quality Object Detection, *Zhaowei Cai, Nuno Vasconcelos*
6. [***D17***] Learning Semantic Concepts and Order for Image and Sentence Matching, *Yan Huang, Qi Wu, Chunfeng Song, Liang Wang*
7. [***D20***] Functional Map of the World, *Gordon Christie, Neil Fendley, James Wilson, Ryan Mukherjee*
8. [***E1***] MegDet: A Large Mini-Batch Object Detector, *Chao Peng, Tete Xiao, Zeming Li, Yuning Jiang, Xiangyu Zhang, Kai Jia, Gang Yu, Jian Sun*
9. [***E4***] Learning Globally Optimized Object Detector via Policy Gradient, *Yongming Rao, Dahua Lin, Jiwen Lu, Jie Zhou*
10. [***E7***] Photographic Text-to-Image Synthesis With a Hierarchically-Nested Adversarial Network, *Zizhao Zhang, Yuanpu Xie, Lin Yang*

1450-1630 Session 2-2C: Computational Photography (Room 255)

Papers in this session are also in Poster Session P2-3. Poster tag in square brackets (e.g., [***E10***])

Chairs: Ayan Chakrabarti *(TTI Chicago)*
Aswin Sankaranarayanan *(Carnegie Mellon Univ.)*

1450 Orals (O2-2C)

Format (12 min. for presentation + 2 min. for questions)

1. [***E10***] Illuminant Spectra-Based Source Separation Using Flash Photography, *Zhuo Hui, Kalyan Sunkavalli, Sunil Hadap, Aswin C. Sankaranarayanan*
2. [***E13***] Trapping Light for Time of Flight, *Ruilin Xu, Mohit Gupta, Shree K. Nayar*
3. [***E16***] The Perception-Distortion Tradeoff, *Yochai Blau, Tomer Michaeli*

1534 Spotlights (S2-2C)

Format (4 min. for presentation; no questions)

1. [***E19***] Label Denoising Adversarial Network (LDAN) for Inverse Lighting of Faces, *Hao Zhou, Jin Sun, Yaser Yacoob, David W. Jacobs*
2. [***E22***] Optimal Structured Light à La Carte, *Parsa Mirdehghan, Wenzheng Chen, Kiriakos N. Kutulakos*
3. [***F3***] Tracking Multiple Objects Outside the Line of Sight Using Speckle Imaging, *Brandon M. Smith, Matthew O'Toole, Mohit Gupta*
4. [***F6***] Inferring Light Fields From Shadows, *Manel Baradad, Vickie Ye, Adam B. Yedidia, Frédo Durand, William T. Freeman, Gregory W. Wornell, Antonio Torralba*
5. [***F9***] Modifying Non-Local Variations Across Multiple Views, *Tal Tlusty, Tomer Michaeli, Tali Dekel, Lihi Zelnik-Manor*

6. [***F12***] Robust Video Content Alignment and Compensation for Rain Removal in a CNN Framework, *Jie Chen, Cheen-Hau Tan, Junhui Hou, Lap-Pui Chau, He Li*
7. [***F15***] SfSNet: Learning Shape, Reflectance and Illuminance of Faces `in the Wild', *Soumyadip Sengupta, Angjoo Kanazawa, Carlos D. Castillo, David W. Jacobs*
8. [***F18***] Deep Photo Enhancer: Unpaired Learning for Image Enhancement From Photographs With GANs, *Yu-Sheng Chen, Yu-Ching Wang, Man-Hsin Kao, Yung-Yu Chuang*
9. [***F21***] LIME: Live Intrinsic Material Estimation, *Abhimitra Meka, Maxim Maximov, Michael Zollhöfer, Avishek Chatterjee, Hans-Peter Seidel, Christian Richardt, Christian Theobalt*
10. [***G2***] Learning to Detect Features in Texture Images, *Linguang Zhang, Szymon Rusinkiewicz*
11. [***G5***] Learning to Extract a Video Sequence From a Single Motion-Blurred Image, *Meiguang Jin, Givi Meishvili, Paolo Favaro*
12. [***G8***] Lose the Views: Limited Angle CT Reconstruction via Implicit Sinogram Completion, *Rushil Anirudh, Hyojin Kim, Jayaraman J. Thiagarajan, K. Aditya Mohan, Kyle Champley, Timo Bremer*
13. [***G11***] A Common Framework for Interactive Texture Transfer, *Yifang Men, Zhouhui Lian, Yingmin Tang, Jianguo Xiao*

1630-1830 Demos (Hall C)

- Sim4CV: A Photo-Realistic Simulator for Computer Vision Applications, *Matthias Mueller, Silvio Giancola, Bernard Ghanem (KAUST)*
- Real-Time Face Anti-Spoofing System, *Yaojie Liu, Amin Jourabloo, Xiaoming Liu (Michigan State Univ.)*
- Efficient Annotation of Segmentation Datasets With Polygon-RNN++, *David Acuna, Huan Ling, Amlan Kar, Sanja Fidler (Univ. of Toronto)*
- Confocal Non-Line-Of-Sight Imaging, *Matthew O'Toole, David B. Lindell, Gordon Wetzstein (Stanford)*
- Semi-Dense, Event-Based Visual SLAM, *Guillermo Gallego, Henri Rebecq, Davide Scaramuzza (Univ. of Zurich, ETH Zurich)*
- Ultimate SLAM? Combining Events, Frames and IMU for Robust Visual SLAM, *Henri Rebecq, Guillermo Gallego, Davide Scaramuzza (Univ. of Zurich, ETH Zurich)*
- Real-Time Visual SLAM Using a Jointly Optimized, Compact Dense Code, *Jan Czarnowski, Michael Bloesch, Ronald Clark, Robert Lukierski, Stefan Leutenegger, Andrew J. Davison (Imperial College London)*
- Cytomine — An Open-Source Web Application for Collaborative Analysis of Multi-Gigapixel Images, *Romain Mormont, Ulysse Rubens, Renaud Hoyoux, Rémy Vandaele, Pierre Geurts, Raphaël Marée (Univ. of Liège)*

1630-1830 Poster Session P2-3 (Halls C-E)

Poster tag in square brackets (e.g., [***G14***])

Computational Photography

1. [**G14**] AMNet: Memorability Estimation With Attention, Jiri Fajtl, Vasileios Argyriou, Dorothy Monekosso, Paolo Remagnino
2. [**G17**] Blind Predicting Similar Quality Map for Image Quality Assessment, Da Pan, Ping Shi, Ming Hou, Zefeng Ying, Sizhe Fu, Yuan Zhang
3. [**G20**] Deep End-to-End Time-of-Flight Imaging, Shuochen Su, Felix Heide, Gordon Wetzstein, Wolfgang Heidrich
4. [**H1**] Aperture Supervision for Monocular Depth Estimation, Pratul P. Srinivasan, Rahul Garg, Neal Wadhwa, Ren Ng, Jonathan T. Barron
5. [**H4**] Seeing Temporal Modulation of Lights From Standard Cameras, Naoki Sakakibara, Fumihiko Sakaue, Jun Sato
6. [**H7**] Statistical Tomography of Microscopic Life, Aviad Levis, Yoav Y. Schechner, Ronen Talmon
7. [**H10**] Divide and Conquer for Full-Resolution Light Field Deblurring, M. R. Mahesh Mohan, A. N. Rajagopalan
8. [**H13**] Multispectral Image Intrinsic Decomposition via Subspace Constraint, Qian Huang, Weixin Zhu, Yang Zhao, Linsen Chen, Yao Wang, Tao Yue, Xun Cao
9. [**H16**] Improving Color Reproduction Accuracy on Cameras, Hakki Can Karaimer, Michael S. Brown

Video Analytics

10. [***H19***] A Closer Look at Spatiotemporal Convolutions for Action Recognition, *Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, Manohar Paluri*
11. [***H22***] Inferring Shared Attention in Social Scene Videos, *Lifeng Fan, Yixin Chen, Ping Wei, Wenguan Wang, Song-Chun Zhu*
12. [***I3***] Making Convolutional Networks Recurrent for Visual Sequence Learning, *Xiaodong Yang, Pavlo Molchanov, Jan Kautz*
13. [***I6***] Real-World Anomaly Detection in Surveillance Videos, *Waqas Sultani, Chen Chen, Mubarak Shah*

14. [**I9**] Viewpoint-Aware Attentive Multi-View Inference for Vehicle Re-Identification, *Yi Zhou, Ling Shao*
15. [**I12**] Efficient Video Object Segmentation via Network Modulation, *Linjie Yang, Yanran Wang, Xuehan Xiong, Jianchao Yang, Aggelos K. Katsaggelos*
16. [**I15**] Weakly-Supervised Action Segmentation With Iterative Soft Boundary Assignment, *Li Ding, Chenliang Xu*
17. [**I18**] Depth-Aware Stereo Video Retargeting, *Bing Li, Chia-Wen Lin, Boxin Shi, Tiejun Huang, Wen Gao, C.-C. Jay Kuo*
18. [**I21**] Instance Embedding Transfer to Unsupervised Video Object Segmentation, *Siyang Li, Bryan Seybold, Alexey Vorobyov, Alireza Fathi, Qin Huang, C.-C. Jay Kuo*
19. [**J2**] Future Frame Prediction for Anomaly Detection – A New Baseline, *Wen Liu, Weixin Luo, Dongze Lian, Shenghua Gao*
20. [**J5**] Can Spatiotemporal 3D CNNs Retrace the History of 2D CNNs and ImageNet?, *Kensho Hara, Hirokatsu Kataoka, Yutaka Satoh*
21. [**J8**] Dynamic Video Segmentation Network, *Yu-Syuan Xu, Tsu-Jui Fu, Hsuan-Kung Yang, Chun-Yi Lee*
22. [**J11**] Recognize Actions by Disentangling Components of Dynamics, *Yue Zhao, Yuanjun Xiong, Dahua Lin*
23. [**J14**] Motion-Appearance Co-Memory Networks for Video Question Answering, *Jiyang Gao, Runzhou Ge, Kan Chen, Ram Nevatia*

Low-level & Mid-level Vision

24. [**J17**] Learning to Understand Image Blur, *Shanghang Zhang, Xiaohui Shen, Zhe Lin, Radomír Měch, João P. Costeira, José M. F. Moura*
25. [**J20**] Dense Decoder Shortcut Connections for Single-Pass Semantic Segmentation, *Piotr Bilinski, Victor Prisacariu*
26. [**K1**] Generative Adversarial Image Synthesis With Decision Tree Latent Controller, *Takuhiro Kaneko, Kaoru Hiramatsu, Kunio Kashino*
27. [**K4**] Learning a Discriminative Prior for Blind Image Deblurring, *Lerenhan Li, Jinshan Pan, Wei-Sheng Lai, Changxin Gao, Nong Sang, Ming-Hsuan Yang*
28. [**K7**] Frame-Recurrent Video Super-Resolution, *Mehdi S. M. Sajjadi, Raviteja Vemulapalli, Matthew Brown*
29. [**K10**] Discovering Point Lights With Intensity Distance Fields, *Edward Zhang, Michael F. Cohen, Brian Curless*
30. [**K13**] Video Rain Streak Removal by Multiscale Convolutional Sparse Coding, *Minghan Li, Qi Xie, Qian Zhao, Wei Wei, Shuhang Gu, Jing Tao, Deyu Meng*
31. [**K16**] Stereoscopic Neural Style Transfer, *Dongdong Chen, Lu Yuan, Jing Liao, Nenghai Yu, Gang Hua*
32. [**K19**] Multi-Frame Quality Enhancement for Compressed Video, *Ren Yang, Mai Xu, Zulin Wang, Tianyi Li*
33. [**K22**] CNN Based Learning Using Reflection and Retinex Models for Intrinsic Image Decomposition, *Anil S. Baslamisli, Hoang-An Le, Theo Gevers*
34. [**L3**] Image Restoration by Estimating Frequency Distribution of Local Patches, *Jaeyoung Yoo, Sang-ho Lee, Nojun Kwak*
35. [**L6**] Latent RANSAC, *Simon Korman, Roee Litman*
36. [**L9**] Two-Stream Convolutional Networks for Dynamic Texture Synthesis, *Matthew Tesfaldet, Marcus A. Brubaker, Konstantinos G. Derpanis*
37. [**L12**] Towards Open-Set Identity Preserving Face Synthesis, *Jianmin Bao, Dong Chen, Fang Wen, Houqiang Li, Gang Hua*
38. [**L15**] A Revised Underwater Image Formation Model, *Derya Akkaynak, Tali Treibitz*
39. [**L18**] Graph-Cut RANSAC, *Daniel Barath, Jiří Matas*

Video Analytics

40. [**L21**] Temporal Deformable Residual Networks for Action Segmentation in Videos, *Peng Lei, Sinisa Todorovic*
41. [**M2**] Weakly Supervised Action Localization by Sparse Temporal Pooling Network, *Phuc Nguyen, Ting Liu, Gautam Prasad, Bohyung Han*
42. [**M5**] PoseFlow: A Deep Motion Representation for Understanding Human Behaviors in Videos, *Dingwen Zhang, Guangyu Guo, Dong Huang, Junwei Han*
43. [**M8**] FFNet: Video Fast-Forwarding via Reinforcement Learning, *Shuyue Lan, Rameswar Panda, Qi Zhu, Amit K. Roy-Chowdhury*
44. [**M11**] Multi-Shot Pedestrian Re-Identification via Sequential Decision Making, *Jianfu Zhang, Naiyan Wang, Liqing Zhang*
45. [**M14**] Attend and Interact: Higher-Order Object Interactions for Video Understanding, *Chih-Yao Ma, Asim Kadav, Iain Melvin, Zsolt Kira, Ghassan AlRegib, Hans Peter Graf*

46. [***M17***] Where and Why Are They Looking? Jointly Inferring Human Attention and Intentions in Complex Tasks, *Ping Wei, Yang Liu, Tianmin Shu, Nanning Zheng, Song-Chun Zhu*
47. [***M20***] Fully Convolutional Adaptation Networks for Semantic Segmentation, *Yiheng Zhang, Zhaofan Qiu, Ting Yao, Dong Liu, Tao Mei*
48. [***N1***] Semantic Video Segmentation by Gated Recurrent Flow Propagation, *David Nilsson, Cristian Sminchisescu*
49. [***N4***] Interpretable Video Captioning via Trajectory Structured Localization, *Xian Wu, Guanbin Li, Qingxing Cao, Qingge Ji, Liang Lin*

Object Recognition & Scene Understanding

50. [***N7***] Deep Hashing via Discrepancy Minimization, *Zhixiang Chen, Xin Yuan, Jiwen Lu, Qi Tian, Jie Zhou*
51. [***N10***] ShuffleNet: An Extremely Efficient Convolutional Neural Network for Mobile Devices, *Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, Jian Sun*
52. [***N13***] Zero-Shot Recognition via Semantic Embeddings and Knowledge Graphs, *Xiaolong Wang, Yufei Ye, Abhinav Gupta*
53. [***N16***] Referring Relationships, *Ranjay Krishna, Ines Chami, Michael Bernstein, Li Fei-Fei*
54. [***N19***] Improving Object Localization With Fitness NMS and Bounded IoU Loss, *Lachlan Tychsen-Smith, Lars Petersson*
55. [***N22***] End-to-End Deep Kronecker-Product Matching for Person Re-Identification, *Yantao Shen, Tong Xiao, Hongsheng Li, Shuai Yi, Xiaogang Wang*
56. [***O3***] Semantic Visual Localization, *Johannes L. Schönberger, Marc Pollefeys, Andreas Geiger, Torsten Sattler*
57. [***O6***] Objects as Context for Detecting Their Semantic Parts, *Abel Gonzalez-Garcia, Davide Modolo, Vittorio Ferrari*
58. [***O9***] End-to-End Weakly-Supervised Semantic Alignment, *Ignacio Rocco, Relja Arandjelović, Josef Sivic*
59. [***O12***] Dynamic Zoom-In Network for Fast Object Detection in Large Images, *Mingfei Gao, Ruichi Yu, Ang Li, Vlad I. Morariu, Larry S. Davis*
60. [***O15***] Learning Markov Clustering Networks for Scene Text Detection, *Zichuan Liu, Guosheng Lin, Sheng Yang, Jiashi Feng, Weisi Lin, Wang Ling Goh*
61. [***O18***] Deep Reinforcement Learning of Region Proposal Networks for Object Detection, *Aleksis Pirinen, Cristian Sminchisescu*
62. [***O21***] Beyond Holistic Object Recognition: Enriching Image Understanding With Part States, *Cewu Lu, Hao Su, Yonglu Li, Yongyi Lu, Li Yi, Chi-Keung Tang, Leonidas J. Guibas*
63. [***P2***] Discriminability Objective for Training Descriptive Captions, *Ruotian Luo, Brian Price, Scott Cohen, Gregory Shakhnarovich*
64. [***P5***] Visual Question Answering With Memory-Augmented Networks, *Chao Ma, Chunhua Shen, Anthony Dick, Qi Wu, Peng Wang, Anton van den Hengel, Ian Reid*
65. [***P8***] Structure Inference Net: Object Detection Using Scene-Level Context and Instance-Level Relationships, *Yong Liu, Ruiping Wang, Shiguang Shan, Xilin Chen*
66. [***P11***] Occluded Pedestrian Detection Through Guided Attention in CNNs, *Shanshan Zhang, Jian Yang, Bernt Schiele*
67. [***P14***] Reward Learning From Narrated Demonstrations, *Hsiao-Yu Tung, Adam W. Harley, Liang-Kang Huang, Katerina Fragkiadaki*
68. [***P17***] Weakly-Supervised Semantic Segmentation Network With Deep Seeded Region Growing, *Zilong Huang, Xinggang Wang, Jiasi Wang, Wenyu Liu, Jingdong Wang*

Analyzing Humans in Images

69. [***P20***] PoTion: Pose MoTion Representation for Action Recognition, *Vasileios Choutas, Philippe Weinzaepfel, Jérôme Revaud, Cordelia Schmid*
70. [***Q1***] Bilateral Ordinal Relevance Multi-Instance Regression for Facial Action Unit Intensity Estimation, *Yong Zhang, Rui Zhao, Weiming Dong, Bao-Gang Hu, Qiang Ji*
71. [***Q4***] Pulling Actions out of Context: Explicit Separation for Effective Combination, *Yang Wang, Minh Hoai*
72. [***Q7***] Dynamic Feature Learning for Partial Face Recognition, *Lingxiao He, Haiqing Li, Qi Zhang, Zhenan Sun*
73. [***Q10***] Exploiting Transitivity for Learning Person Re-Identification Models on a Budget, *Sourya Roy, Sujoy Paul, Neal E. Young, Amit K. Roy-Chowdhury*
74. [***Q13***] Deep Spatial Feature Reconstruction for Partial Person Re-Identification: Alignment-Free Approach, *Lingxiao He, Jian Liang, Haiqing Li, Zhenan Sun*
75. [***Q16***] Every Smile Is Unique: Landmark-Guided Diverse Smile Generation, *Wei Wang, Xavier Alameda-Pineda, Dan Xu, Pascal Fua, Elisa Ricci, Nicu Sebe*

76. [***Q19***] UV-GAN: Adversarial Facial UV Map Completion for Pose-Invariant Face Recognition, *Jiankang Deng, Shiyang Cheng, Niannan Xue, Yuxiang Zhou, Stefanos Zafeiriou*
77. [***Q22***] Cascaded Pyramid Network for Multi-Person Pose Estimation, *Yilun Chen, Zhicheng Wang, Yuxiang Peng, Zhiqiang Zhang, Gang Yu, Jian Sun*
78. [***R3***] A Face-to-Face Neural Conversation Model, *Hang Chu, Daiqing Li, Sanja Fidler*
79. [***R6***] End-to-End Recovery of Human Shape and Pose, *Angjoo Kanazawa, Michael J. Black, David W. Jacobs, Jitendra Malik*

Notes:

Thursday, June 21

0730–1830 Registration (South Lobby)

0730–0900 Breakfast (Halls 1-4)

0800–1000 Poster Setup (Halls D-E)

0830-1010 Session 3-1A: Object Recognition & Scene Understanding IV (Ballroom)

Papers in this session are also in Poster Session P3-1. Poster tag in square brackets (e.g., [***A1***])

Chairs: Georgia Gkioxari *(Facebook AI Research)*
Adriana Kovashka *(Univ. of Pittsburgh)*

0830 Orals (O3-1A)

Format (12 min. for presentation + 2 min. for questions)

1. [***A1***] Squeeze-and-Excitation Networks, *Jie Hu, Li Shen, Gang Sun*
2. [***A3***] Revisiting Salient Object Detection: Simultaneous Detection, Ranking, and Subitizing of Multiple Salient Objects, *Md Amirul Islam, Mahmoud Kalash, Neil D. B. Bruce*
3. [***A5***] Context Encoding for Semantic Segmentation, *Hang Zhang, Kristin Dana, Jianping Shi, Zhongyue Zhang, Xiaogang Wang, Ambrish Tyagi, Amit Agrawal*

0914 Spotlights (S3-1A)

Format (4 min. for presentation; no questions)

1. [***A7***] Creating Capsule Wardrobes From Fashion Images, *Wei-Lin Hsiao, Kristen Grauman*
2. [***A9***] Webly Supervised Learning Meets Zero-Shot Learning: A Hybrid Approach for Fine-Grained Classification, *Li Niu, Ashok Veeraraghavan, Ashutosh Sabharwal*
3. [***A22***] Look, Imagine and Match: Improving Textual-Visual Cross-Modal Retrieval With Generative Models, *Jiuxiang Gu, Jianfei Cai, Shafiq R. Joty, Li Niu, Gang Wang*
4. [***B2***] Bidirectional Attentive Fusion With Context Gating for Dense Video Captioning, *Jingwen Wang, Wenhao Jiang, Lin Ma, Wei Liu, Yong Xu*
5. [***B4***] InLoc: Indoor Visual Localization With Dense Matching and View Synthesis, *Hajime Taira, Masatoshi Okutomi, Torsten Sattler, Mircea Cimpoi, Marc Pollefeys, Josef Sivic, Tomas Pajdla, Akihiko Torii*
6. [***B6***] Towards High Performance Video Object Detection, *Xizhou Zhu, Jifeng Dai, Lu Yuan, Yichen Wei*
7. [***B8***] Neural Baby Talk, *Jiasen Lu, Jianwei Yang, Dhruv Batra, Devi Parikh*
8. [***B10***] Few-Shot Image Recognition by Predicting Parameters From Activations, *Siyuan Qiao, Chenxi Liu, Wei Shen, Alan L. Yuille*
9. [***B12***] Iterative Visual Reasoning Beyond Convolutions, *Xinlei Chen, Li-Jia Li, Li Fei-Fei, Abhinav Gupta*
10. [***B14***] Visual Question Reasoning on General Dependency Tree, *Qingxing Cao, Xiaodan Liang, Bailing Li, Guanbin Li, Liang Lin*
11. [***B16***] CVM-Net: Cross-View Matching Network for Image-Based Ground-to-Aerial Geo-Localization, *Sixing Hu, Mengdan Feng, Rang M. H. Nguyen, Gim Hee Lee*
12. [***B18***] Revisiting Dilated Convolution: A Simple Approach for Weakly- and Semi-Supervised Semantic Segmentation, *Yunchao Wei, Huaxin Xiao, Honghui Shi, Zequn Jie, Jiashi Feng, Thomas S. Huang*
13. [***B20***] Low-Shot Learning From Imaginary Data, *Yu-Xiong Wang, Ross Girshick, Martial Hebert, Bharath Hariharan*

0830-1010 Session 3-1B: Analyzing Humans (Room 155)

Papers in this session are also in Poster Session P3-1. Poster tag in square brackets (e.g., [***B22***])

Chairs: Greg Rogez *(INRIA Rhones-Alpes)*
Hamed Pirsiavash *(U. of Maryland, Baltimore County)*

0830 Orals (O3-1B)

Format (12 min. for presentation + 2 min. for questions)

1. [***B22***] DoubleFusion: Real-Time Capture of Human Performances With Inner Body Shapes From a Single Depth Sensor, *Tao Yu, Zerong Zheng, Kaiwen Guo, Jianhui Zhao, Qionghai Dai, Hao Li, Gerard Pons-Moll, Yebin Liu*
2. [***C2***] DensePose: Dense Human Pose Estimation in the Wild, *Rıza Alp Güler, Natalia Neverova, Iasonas Kokkinos*
3. [***C4***] Ordinal Depth Supervision for 3D Human Pose Estimation, *Georgios Pavlakos, Xiaowei Zhou, Kostas Daniilidis*

0914 Spotlights (S3-1B)

Format (4 min. for presentation; no questions)

1. [***C6***] Consensus Maximization for Semantic Region Correspondences, *Pablo Speciale, Danda P. Paudel, Martin R. Oswald, Hayko Riemenschneider, Luc Van Gool, Marc Pollefeys*
2. [***C8***] Robust Hough Transform Based 3D Reconstruction From Circular Light Fields, *Alessandro Vianello, Jens Ackermann, Maximilian Diebold, Bernd Jähne*
3. [***C10***] Alive Caricature From 2D to 3D, *Qianyi Wu, Juyong Zhang, Yu-Kun Lai, Jianmin Zheng, Jianfei Cai*
4. [***C12***] Nonlinear 3D Face Morphable Model, *Luan Tran, Xiaoming Liu*
5. [***C14***] Through-Wall Human Pose Estimation Using Radio Signals, *Mingmin Zhao, Tianhong Li, Mohammad Abu Alsheikh, Yonglong Tian, Hang Zhao, Antonio Torralba, Dina Katabi*
6. [***C16***] What Makes a Video a Video: Analyzing Temporal Information in Video Understanding Models and Datasets, *De-An Huang, Vignesh Ramanathan, Dhruv Mahajan, Lorenzo Torresani, Manohar Paluri, Li Fei-Fei, Juan Carlos Niebles*
7. [***C18***] Fast Video Object Segmentation by Reference-Guided Mask Propagation, *Seoung Wug Oh, Joon-Young Lee, Kalyan Sunkavalli, Seon Joo Kim*
8. [***C20***] NeuralNetwork-Viterbi: A Framework for Weakly Supervised Video Learning, *Alexander Richard, Hilde Kuehne, Ahsan Iqbal, Juergen Gall*
9. [***C22***] Actor and Observer: Joint Modeling of First and Third-Person Videos, *Gunnar A. Sigurdsson, Abhinav Gupta, Cordelia Schmid, Ali Farhadi, Karteek Alahari*
10. [***D2***] HSA-RNN: Hierarchical Structure-Adaptive RNN for Video Summarization, *Bin Zhao, Xuelong Li, Xiaoqiang Lu*
11. [***D4***] Fast and Accurate Online Video Object Segmentation via Tracking Parts, *Jingchun Cheng, Yi-Hsuan Tsai, Wei-Chih Hung, Shengjin Wang, Ming-Hsuan Yang*
12. [***D6***] Now You Shake Me: Towards Automatic 4D Cinema, *Yuhao Zhou, Makarand Tapaswi, Sanja Fidler*
13. [***D8***] Viewpoint-Aware Video Summarization, *Atsushi Kanehira, Luc Van Gool, Yoshitaka Ushiku, Tatsuya Harada*

0830-1010 Session 3-1C: Applications
(Room 255)

Papers in this session are also in Poster Session P3-1. Poster tag in square brackets (e.g., [***D10***])

Chairs: Devi Parikh *(Georgia Tech)*
Kristen Grauman *(Univ. of Texas at Austin)*

0830 Orals (O3-1C)

Format (12 min. for presentation + 2 min. for questions)

1. [***D10***] Photometric Stereo in Participating Media Considering Shape-Dependent Forward Scatter, *Yuki Fujimura, Masaaki Iiyama, Atsushi Hashimoto, Michihiko Minoh*
2. [***D12***] Direction-Aware Spatial Context Features for Shadow Detection, *Xiaowei Hu, Lei Zhu, Chi-Wing Fu, Jing Qin, Pheng-Ann Heng*
3. [***D14***] Discriminative Learning of Latent Features for Zero-Shot Recognition, *Yan Li, Junge Zhang, Jianguo Zhang, Kaiqi Huang*

0914 Spotlights (S3-1C)

Format (4 min. for presentation; no questions)

1. [***D16***] Learning to Adapt Structured Output Space for Semantic Segmentation, *Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schulter, Kihyuk Sohn, Ming-Hsuan Yang, Manmohan Chandraker*
2. [***D18***] Multi-Task Learning Using Uncertainty to Weigh Losses for Scene Geometry and Semantics, *Alex Kendall, Yarin Gal, Roberto Cipolla*
3. [***D20***] Jointly Localizing and Describing Events for Dense Video Captioning, *Yehao Li, Ting Yao, Yingwei Pan, Hongyang Chao, Tao Mei*
4. [***D22***] Going From Image to Video Saliency: Augmenting Image Salience With Dynamic Attentional Push, *Siavash Gorji, James J. Clark*
5. [***E2***] M^3: Multimodal Memory Modelling for Video Captioning, *Junbo Wang, Wei Wang, Yan Huang, Liang Wang, Tieniu Tan*
6. [***E4***] Emotional Attention: A Study of Image Sentiment and Visual Attention, *Shaojing Fan, Zhiqi Shen, Ming Jiang, Bryan L. Koenig, Juan Xu, Mohan S. Kankanhalli, Qi Zhao*
7. [***E6***] A Low Power, High Throughput, Fully Event-Based Stereo System, *Alexander Andreopoulos, Hirak J. Kashyap, Tapan K. Nayak, Arnon Amir, Myron D. Flickner*

8. [***E8***] VITON: An Image-Based Virtual Try-On Network, *Xintong Han, Zuxuan Wu, Zhe Wu, Ruichi Yu, Larry S. Davis*
9. [***E10***] Multi-Oriented Scene Text Detection via Corner Localization and Region Segmentation, *Pengyuan Lyu, Cong Yao, Wenhao Wu, Shuicheng Yan, Xiang Bai*
10. [***E12***] Multi-Content GAN for Few-Shot Font Style Transfer, *Samaneh Azadi, Matthew Fisher, Vladimir G. Kim, Zhaowen Wang, Eli Shechtman, Trevor Darrell*
11. [***E14***] Audio to Body Dynamics, *Eli Shlizerman, Lucio Dery, Hayden Schoen, Ira Kemelmacher-Shlizerman*
12. [***E16***] Weakly Supervised Coupled Networks for Visual Sentiment Analysis, *Jufeng Yang, Dongyu She, Yu-Kun Lai, Paul L. Rosin, Ming-Hsuan Yang*
13. [***E18***] Future Person Localization in First-Person Videos, *Takuma Yagi, Karttikeya Mangalam, Ryo Yonetani, Yoichi Sato*

1000–1045 Morning Break (Halls 1-4)

1000-1830 Exhibits (Halls D-E)

- See Exhibits map for list of exhibitors.

1010-1230 Poster Session P3-1 (Halls D-E)

Poster tag in square brackets (e.g., [***E20***])

Object Recognition & Scene Understanding

1. [***E20***] Preserving Semantic Relations for Zero-Shot Learning, *Yashas Annadani, Soma Biswas*
2. [***E22***] Show Me a Story: Towards Coherent Neural Story Illustration, *Hareesh Ravi, Lezi Wang, Carlos Muniz, Leonid Sigal, Dimitris Metaxas, Mubbasir Kapadia*
3. [***F2***] Reconstruction Network for Video Captioning, *Bairui Wang, Lin Ma, Wei Zhang, Wei Liu*
4. [***F4***] Fast Spectral Ranking for Similarity Search, *Ahmet Iscen, Yannis Avrithis, Giorgos Tolias, Teddy Furon, Ondřej Chum*
5. [***F6***] Mining on Manifolds: Metric Learning Without Labels, *Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, Ondřej Chum*
6. [***F8***] PIXOR: Real-Time 3D Object Detection From Point Clouds, *Bin Yang, Wenjie Luo, Raquel Urtasun*
7. [***F10***] Leveraging Unlabeled Data for Crowd Counting by Learning to Rank, *Xialei Liu, Joost van de Weijer, Andrew D. Bagdanov*
8. [***F12***] Zero-Shot Kernel Learning, *Hongguang Zhang, Piotr Koniusz*
9. [***F14***] Differential Attention for Visual Question Answering, *Badri Patro, Vinay P. Namboodiri*
10. [***F16***] Learning From Noisy Web Data With Category-Level Supervision, *Li Niu, Qingtao Tang, Ashok Veeraraghavan, Ashutosh Sabharwal*
11. [***F18***] Toward Driving Scene Understanding: A Dataset for Learning Driver Behavior and Causal Reasoning, *Vasili Ramanishka, Yi-Ting Chen, Teruhisa Misu, Kate Saenko*
12. [***F20***] Learning Attribute Representations With Localization for Flexible Fashion Search, *Kenan E. Ak, Ashraf A. Kassim, Joo Hwee Lim, Jo Yew Tham*
13. [***F22***] Bidirectional Retrieval Made Simple, *Jônatas Wehrmann, Rodrigo C. Barros*
14. [***G2***] Learning Multi-Instance Enriched Image Representations via Non-Greedy Ratio Maximization of the l_1-Norm Distances, *Kai Liu, Hua Wang, Feiping Nie, Hao Zhang*
15. [***G4***] Learning Visual Knowledge Memory Networks for Visual Question Answering, *Zhou Su, Chen Zhu, Yinpeng Dong, Dongqi Cai, Yurong Chen, Jianguo Li*
16. [***G6***] Visual Grounding via Accumulated Attention, *Chaorui Deng, Qi Wu, Qingyao Wu, Fuyuan Hu, Fan Lyu, Mingkui Tan*
17. [***G8***] Beyond Trade-Off: Accelerate FCN-Based Face Detector With Higher Accuracy, *Guanglu Song, Yu Liu, Ming Jiang, Yujie Wang, Junjie Yan, Biao Leng*
18. [***G10***] PackNet: Adding Multiple Tasks to a Single Network by Iterative Pruning, *Arun Mallya, Svetlana Lazebnik*
19. [***G12***] Repulsion Loss: Detecting Pedestrians in a Crowd, *Xinlong Wang, Tete Xiao, Yuning Jiang, Shuai Shao, Jian Sun, Chunhua Shen*
20. [***G14***] Neural Sign Language Translation, *Necati Cihan Camgoz, Simon Hadfield, Oscar Koller, Hermann Ney, Richard Bowden*

Video Analytics

21. [***G16***] Non-Local Neural Networks, *Xiaolong Wang, Ross Girshick, Abhinav Gupta, Kaiming He*
22. [***G18***] LAMV: Learning to Align and Match Videos With Kernelized Temporal Layers, *Lorenzo Baraldi, Matthijs Douze, Rita Cucchiara, Hervé Jégou*
23. [***G20***] Optimizing Video Object Detection via a Scale-Time Lattice, *Kai Chen, Jiaqi Wang, Shuo Yang, Xingcheng Zhang, Yuanjun Xiong, Chen Change Loy, Dahua Lin*
24. [***G22***] Learning Compressible 360° Video Isomers, *Yu-Chuan Su, Kristen Grauman*

25. [***H2***] Attention Clusters: Purely Attention Based Local Feature Integration for Video Classification, *Xiang Long, Chuang Gan, Gerard de Melo, Jiajun Wu, Xiao Liu, Shilei Wen*
26. [***H4***] What Have We Learned From Deep Representations for Action Recognition?, *Christoph Feichtenhofer, Axel Pinz, Richard P. Wildes, Andrew Zisserman*
27. [***H6***] Controllable Video Generation With Sparse Trajectories, *Zekun Hao, Xun Huang, Serge Belongie*

Machine Learning for Computer Vision

28. [***H8***] Representing and Learning High Dimensional Data With the Optimal Transport Map From a Probabilistic Viewpoint, *Serim Park, Matthew Thorpe*
29. [***H10***] CLIP-Q: Deep Network Compression Learning by In-Parallel Pruning-Quantization, *Frederick Tung, Greg Mori*
30. [***H12***] Inference in Higher Order MRF-MAP Problems With Small and Large Cliques, *Ishant Shanu, Chetan Arora, S.N. Maheshwari*
31. [***H14***] ROAD: Reality Oriented Adaptation for Semantic Segmentation of Urban Scenes, *Yuhua Chen, Wen Li, Luc Van Gool*
32. [***H16***] Eye In-Painting With Exemplar Generative Adversarial Networks, *Brian Dolhansky, Cristian Canton Ferrer*
33. [***H18***] ClcNet: Improving the Efficiency of Convolutional Neural Network Using Channel Local Convolutions, *Dong-Qing Zhang*
34. [***H20***] Towards Effective Low-Bitwidth Convolutional Neural Networks, *Bohan Zhuang, Chunhua Shen, Mingkui Tan, Lingqiao Liu, Ian Reid*
35. [***H22***] Stochastic Downsampling for Cost-Adjustable Inference and Improved Regularization in Convolutional Networks, *Jason Kuen, Xiangfei Kong, Zhe Lin, Gang Wang, Jianxiong Yin, Simon See, Yap-Peng Tan*
36. [***I2***] Face Aging With Identity-Preserved Conditional Generative Adversarial Networks, *Zongwei Wang, Xu Tang, Weixin Luo, Shenghua Gao*
37. [***I4***] Unsupervised Cross-Dataset Person Re-Identification by Transfer Learning of Spatial-Temporal Patterns, *Jianming Lv, Weihang Chen, Qing Li, Can Yang*
38. [***I6***] Feature Quantization for Defending Against Distortion of Images, *Zhun Sun, Mete Ozay, Yan Zhang, Xing Liu, Takayuki Okatani*
39. [***I8***] Tagging Like Humans: Diverse and Distinct Image Annotation, *Baoyuan Wu, Weidong Chen, Peng Sun, Wei Liu, Bernard Ghanem, Siwei Lyu*
40. [***I10***] Re-Weighted Adversarial Adaptation Network for Unsupervised Domain Adaptation, *Qingchao Chen, Yang Liu, Zhaowen Wang, Ian Wassell, Kevin Chetty*
41. [***I12***] Inferring Semantic Layout for Hierarchical Text-to-Image Synthesis, *Seunghoon Hong, Dingdong Yang, Jongwook Choi, Honglak Lee*
42. [***I14***] Regularizing RNNs for Caption Generation by Reconstructing the Past With the Present, *Xinpeng Chen, Lin Ma, Wenhao Jiang, Jian Yao, Wei Liu*
43. [***I16***] Unsupervised Domain Adaptation With Similarity Learning, *Pedro O. Pinheiro*
44. [***I18***] Learning Deep Sketch Abstraction, *Umar Riaz Muhammad, Yongxin Yang, Yi-Zhe Song, Tao Xiang, Timothy M. Hospedales*
45. [***I20***] Matching Adversarial Networks, *Gellért Máttyus, Raquel Urtasun*
46. [***I22***] SoS-RSC: A Sum-of-Squares Polynomial Approach to Robustifying Subspace Clustering Algorithms, *Mario Sznaier, Octavia Camps*
47. [***J2***] Resource Aware Person Re-Identification Across Multiple Resolutions, *Yan Wang, Lequn Wang, Yurong You, Xu Zou, Vincent Chen, Serena Li, Gao Huang, Bharath Hariharan, Kilian Q. Weinberger*
48. [***J4***] Learning and Using the Arrow of Time, *Donglai Wei, Joseph J. Lim, Andrew Zisserman, William T. Freeman*
49. [***J6***] Neural Style Transfer via Meta Networks, *Falong Shen, Shuicheng Yan, Gang Zeng*
50. [***J8***] People, Penguins and Petri Dishes: Adapting Object Counting Models to New Visual Domains and Object Types Without Forgetting, *Mark Marsden, Kevin McGuinness, Suzanne Little, Ciara E. Keogh, Noel E. O'Connor*
51. [***J10***] HydraNets: Specialized Dynamic Architectures for Efficient Inference, *Ravi Teja Mullapudi, William R. Mark, Noam Shazeer, Kayvon Fatahalian*
52. [***J12***] SketchMate: Deep Hashing for Million-Scale Human Sketch Retrieval, *Peng Xu, Yongye Huang, Tongtong Yuan, Kaiyue Pang, Yi-Zhe Song, Tao Xiang, Timothy M. Hospedales, Zhanyu Ma, Jun Guo*

53. [**J14**] From Source to Target and Back: Symmetric Bi-Directional Adaptive GAN, *Paolo Russo, Fabio M. Carlucci, Tatiana Tommasi, Barbara Caputo*
54. [**J16**] OLÉ: Orthogonal Low-Rank Embedding - A Plug and Play Geometric Loss for Deep Learning, *José Lezama, Qiang Qiu, Pablo Musé, Guillermo Sapiro*
55. [**J18**] Efficient Parametrization of Multi-Domain Deep Neural Networks, *Sylvestre-Alvise Rebuffi, Hakan Bilen, Andrea Vedaldi*
56. [**J20**] Deep Density Clustering of Unconstrained Faces, *Wei-An Lin, Jun-Cheng Chen, Carlos D. Castillo, Rama Chellappa*
57. [**J22**] Geometric Multi-Model Fitting With a Convex Relaxation Algorithm, *Paul Amayo, Pedro Piniés, Lina M. Paz, Paul Newman*
58. [**K2**] Fast and Robust Estimation for Unit-Norm Constrained Linear Fitting Problems, *Daiki Ikami, Toshihiko Yamasaki, Kiyoharu Aizawa*
59. [**K4**] Importance Weighted Adversarial Nets for Partial Domain Adaptation, *Jing Zhang, Zewei Ding, Wanqing Li, Philip Ogunbona*

Low-level & Mid-level Vision

60. [**K6**] Efficient Subpixel Refinement With Symbolic Linear Predictors, *Vincent Lui, Jonathon Geeves, Winston Yii, Tom Drummond*
61. [**K8**] Scale-Recurrent Network for Deep Image Deblurring, *Xin Tao, Hongyun Gao, Xiaoyong Shen, Jue Wang, Jiaya Jia*
62. [**K10**] DeblurGAN: Blind Motion Deblurring Using Conditional Adversarial Networks, *Orest Kupyn, Volodymyr Budzan, Mykola Mykhailych, Dmytro Mishkin, Jiří Matas*
63. [**K12**] A2-RL: Aesthetics Aware Reinforcement Learning for Image Cropping, *Debang Li, Huikai Wu, Junge Zhang, Kaiqi Huang*
64. [**K14**] Single Image Dehazing via Conditional Generative Adversarial Network, *Runde Li, Jinshan Pan, Zechao Li, Jinhui Tang*
65. [**K16**] On the Duality Between Retinex and Image Dehazing, *Adrian Galdran, Aitor Alvarez-Gila, Alessandro Bria, Javier Vazquez-Corral, Marcelo Bertalmío*
66. [**K18**] Arbitrary Style Transfer With Deep Feature Reshuffle, *Shuyang Gu, Congliang Chen, Jing Liao, Lu Yuan*
67. [**K20**] Nonlocal Low-Rank Tensor Factor Analysis for Image Restoration, *Xinyuan Zhang, Xin Yuan, Lawrence Carin*
68. [**K22**] Avatar-Net: Multi-Scale Zero-Shot Style Transfer by Feature Decoration, *Lu Sheng, Ziyi Lin, Jing Shao, Xiaogang Wang*
69. [**L2**] Missing Slice Recovery for Tensors Using a Low-Rank Model in Embedded Space, *Tatsuya Yokota, Burak Erem, Seyhmus Guler, Simon K. Warfield, Hidekata Hontani*
70. [**L4**] Deep Semantic Face Deblurring, *Ziyi Shen, Wei-Sheng Lai, Tingfa Xu, Jan Kautz, Ming-Hsuan Yang*
71. [**L6**] GraphBit: Bitwise Interaction Mining via Deep Reinforcement Learning, *Yueqi Duan, Ziwei Wang, Jiwen Lu, Xudong Lin, Jie Zhou*

Biomedical Image

72. [**L8**] Recurrent Saliency Transformation Network: Incorporating Multi-Stage Visual Cues for Small Organ Segmentation, *Qihang Yu, Lingxi Xie, Yan Wang, Yuyin Zhou, Elliot K. Fishman, Alan L. Yuille*
73. [**L10**] Thoracic Disease Identification and Localization With Limited Supervision, *Zhe Li, Chong Wang, Mei Han, Yuan Xue, Wei Wei, Li-Jia Li, Li Fei-Fei*
74. [**L12**] Quantization of Fully Convolutional Networks for Accurate Biomedical Image Segmentation, *Xiaowei Xu, Qing Lu, Lin Yang, Sharon Hu, Danny Chen, Yu Hu, Yiyu Shi*
75. [**L14**] Visual Feature Attribution Using Wasserstein GANs, *Christian F. Baumgartner, Lisa M. Koch, Kerem Can Tezcan, Jia Xi Ang, Ender Konukoglu*

1130–1330 Lunch (Halls 1-4)

Notes:

1230–1430 Doctoral Consortium (Room 254)
(by invitation only)

Supported by: intel

- Anurag Arnab (Univ. of Oxford)
- Tolga Birdal (TU Munich)
- Wei-Lun Chao (USC)
- Shixing Chen (Wayne State Univ.)
- Xinlei Chen (CMU/Facebook)
- Donghyeon Cho (KAIST)
- Jongwon Choi (Seoul National Univ.)
- Arthur D. Costea (Technical Univ. of Cluj-Napoca)
- Kun He (Boston Univ.)
- Sarfaraz Hussein (Univ. of Central Florida)
- Hae-Gon Jeon ((KAIST)
- Meiguang Jin (Univ. of Bern)
- Hanbyul Joo (CMU)
- Amin Jourabloo (Michigan State Univ.)
- Heechul Jung (KAIST)
- Pyojin Kim (Seoul National Univ.)
- Ksenia Konyushova (EPFL)
- Jason Kuen (Nanyang Technological Univ.)
- Arun Kumar (Univ. of Georgia)
- Yixuan Li (Facebook Research)
- Ding Liu (UIUC)
- Seong Joon Oh (MPI-INF)
- Rameswar Panda (UC Riverside)
- Xiaojuan Qi (CUHK)
- Zhile Ren (Brown Univ.)
- Soumyadip Sengupta (Univ. of Maryland, College Park)
- Subeesh Vasu (IIT Madras)
- Jingya Wang (Queen Mary Univ. of London)
- Saining Xie (UCSD)
- Dingwen Zhang (Northwestern Polytechnical Univ.)
- Li Zhang (Queen Mary Univ. of London)
- Yong Zhang (Inst. of Automation, CAS)
- Qianggong Zhang (Univ. of Adelaide)
- He Zhang (Rutgers Univ.)
- Shanghang Zhang (CMU)
- Zizhao Zhang (Univ. of Florida)
- Rui Zhao (Rensselaer Polytechnic Inst.)
- Bolei Zhou (MIT)

Notes:

1250-1430 Session 3-2A: Analyzing Humans in Images II (Ballroom)

Papers in this session are also in Poster Session P3-2.
Poster tag in square brackets (e.g., [***A2***])

Chairs: Alexandre Alahi *(EPFL)*
Juan Carlos Niebles *(Stanford Univ.)*

1250 Orals (O3-2A)

Format (12 min. for presentation + 2 min. for questions)

1. [***A2***] Total Capture: A 3D Deformation Model for Tracking Faces, Hands, and Bodies, *Hanbyul Joo, Tomas Simon, Yaser Sheikh*
2. [***A4***] Augmented Skeleton Space Transfer for Depth-Based Hand Pose Estimation, *Seungryul Baek, Kwang In Kim, Tae-Kyun Kim*
3. [***A6***] Synthesizing Images of Humans in Unseen Poses, *Guha Balakrishnan, Amy Zhao, Adrian V. Dalca, Frédo Durand, John Guttag*

1334 Spotlights (S3-2A)

Format (4 min. for presentation; no questions)

1. [***A8***] SSNet: Scale Selection Network for Online 3D Action Prediction, *Jun Liu, Amir Shahroudy, Gang Wang, Ling-Yu Duan, Alex C. Kot*
2. [***A10***] Detecting and Recognizing Human-Object Interactions, *Georgia Gkioxari, Ross Girshick, Piotr Dollár, Kaiming He*
3. [***B1***] Unsupervised Learning and Segmentation of Complex Activities From Video, *Fadime Sener, Angela Yao*
4. [***B3***] Unsupervised Training for 3D Morphable Model Regression, *Kyle Genova, Forrester Cole, Aaron Maschinot, Aaron Sarna, Daniel Vlasic, William T. Freeman*
5. [***B5***] Video Based Reconstruction of 3D People Models, *Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, Gerard Pons-Moll*
6. [***B7***] Pose-Guided Photorealistic Face Rotation, *Yibo Hu, Xiang Wu, Bing Yu, Ran He, Zhenan Sun*
7. [***B9***] Mesoscopic Facial Geometry Inference Using Deep Neural Networks, *Loc Huynh, Weikai Chen, Shunsuke Saito, Jun Xing, Koki Nagano, Andrew Jones, Paul Debevec, Hao Li*
8. [***B11***] Hand PointNet: 3D Hand Pose Estimation Using Point Sets, *Liuhao Ge, Yujun Cai, Junwu Weng, Junsong Yuan*
9. [***B13***] Seeing Voices and Hearing Faces: Cross-Modal Biometric Matching, *Arsha Nagrani, Samuel Albanie, Andrew Zisserman*
10. [***B15***] Learning Monocular 3D Human Pose Estimation From Multi-View Images, *Helge Rhodin, Jörg Spörri, Isinsu Katircioglu, Victor Constantin, Frédéric Meyer, Erich Müller, Mathieu Salzmann, Pascal Fua*
11. [***B17***] Separating Style and Content for Generalized Style Transfer, *Yexun Zhang, Ya Zhang, Wenbin Cai*
12. [***B19***] TextureGAN: Controlling Deep Image Synthesis With Texture Patches, *Wenqi Xian, Patsorn Sangkloy, Varun Agrawal, Amit Raj, Jingwan Lu, Chen Fang, Fisher Yu, James Hays*
13. [***B21***] Connecting Pixels to Privacy and Utility: Automatic Redaction of Private Information in Images, *Tribhuvanesh Orekondy, Mario Fritz, Bernt Schiele*

1250-1430 Session 3-2B: Machine Learning for Computer Vision IV (Room 155)

Papers in this session are also in Poster Session P3-2.
Poster tag in square brackets (e.g., [***C1***])

Chairs: Olga Russakovsky *(Princeton Univ.)*
Philipp Krahenbuhl *(Univ. of Texas at Austin)*

1250 Orals (O3-2B)

Format (12 min. for presentation + 2 min. for questions)

1. [***C1***] MapNet: An Allocentric Spatial Memory for Mapping Environments, *João F. Henriques, Andrea Vedaldi*
2. [***C3***] Accurate and Diverse Sampling of Sequences Based on a "Best of Many" Sample Objective, *Apratim Bhattacharyya, Bernt Schiele, Mario Fritz*
3. [***C5***] VirtualHome: Simulating Household Activities via Programs, *Xavier Puig, Kevin Ra, Marko Boben, Jiaman Li, Tingwu Wang, Sanja Fidler, Antonio Torralba*

1334 Spotlights (S3-2B)

Format (4 min. for presentation; no questions)

1. [***C7***] Generate to Adapt: Aligning Domains Using Generative Adversarial Networks, *Swami Sankaranarayanan, Yogesh Balaji, Carlos D. Castillo, Rama Chellappa*
2. [***C9***] Multi-Agent Diverse Generative Adversarial Networks, *Arnab Ghosh, Viveka Kulharia, Vinay P. Namboodiri, Philip H.S. Torr, Puneet K. Dokania*

3. [***C11***] A PID Controller Approach for Stochastic Optimization of Deep Networks, *Wangpeng An, Haoqian Wang, Qingyun Sun, Jun Xu, Qionghai Dai, Lei Zhang*
4. [***C13***] "Learning-Compression" Algorithms for Neural Net Pruning, *Miguel Á. Carreira-Perpiñán, Yerlan Idelbayev*
5. [***C15***] Large-Scale Distance Metric Learning With Uncertainty, *Qi Qian, Jiasheng Tang, Hao Li, Shenghuo Zhu, Rong Jin*
6. [***C17***] Guide Me: Interacting With Deep Networks, *Christian Rupprecht, Iro Laina, Nassir Navab, Gregory D. Hager, Federico Tombari*
7. [***C19***] Art of Singular Vectors and Universal Adversarial Perturbations, *Valentin Khrulkov, Ivan Oseledets*
8. [***C21***] Deflecting Adversarial Attacks With Pixel Deflection, *Aaditya Prakash, Nick Moran, Solomon Garber, Antonella DiLillo, James Storer*
9. [***D1***] MovieGraphs: Towards Understanding Human-Centric Situations From Videos, *Paul Vicol, Makarand Tapaswi, Lluís Castrejón, Sanja Fidler*
10. [***D3***] SemStyle: Learning to Generate Stylised Image Captions Using Unaligned Text, *Alexander Mathews, Lexing Xie, Xuming He*
11. [***D5***] Benchmarking 6DOF Outdoor Visual Localization in Changing Conditions, *Torsten Sattler, Will Maddern, Carl Toft, Akihiko Torii, Lars Hammarstrand, Erik Stenborg, Daniel Safari, Masatoshi Okutomi, Marc Pollefeys, Josef Sivic, Fredrik Kahl, Tomas Pajdla*
12. [***D7***] IVQA: Inverse Visual Question Answering, *Feng Liu, Tao Xiang, Timothy M. Hospedales, Wankou Yang, Changyin Sun*
13. [***D9***] Unsupervised Person Image Synthesis in Arbitrary Poses, *Albert Pumarola, Antonio Agudo, Alberto Sanfeliu, Francesc Moreno-Noguer*

1250-1430 Session 3-2C: Object Recognition & Scene Understanding V (Room 255)

Papers in this session are also in Poster Session P3-2. Poster tag in square brackets (e.g., [***D11***])

Chairs: Greg Shakhnarovich *(TTI Chicago)*
Bernard Ghanem *(KAUST)*

1250 Orals (O3-2C)

Format (12 min. for presentation + 2 min. for questions)

1. [***D11***] Learning Descriptor Networks for 3D Shape Synthesis and Analysis, *Jianwen Xie, Zilong Zheng, Ruiqi Gao, Wenguan Wang, Song-Chun Zhu, Ying Nian Wu*
2. [***D13***] Neural Kinematic Networks for Unsupervised Motion Retargetting, *Ruben Villegas, Jimei Yang, Duygu Ceylan, Honglak Lee*
3. [***D15***] Group Consistent Similarity Learning via Deep CRF for Person Re-Identification, *Dapeng Chen, Dan Xu, Hongsheng Li, Nicu Sebe, Xiaogang Wang*

1334 Spotlights (S3-2C)

Format (4 min. for presentation; no questions)

1. [***D17***] Learning Compositional Visual Concepts With Mutual Consistency, *Yunye Gong, Srikrishna Karanam, Ziyan Wu, Kuan-Chuan Peng, Jan Ernst, Peter C. Doerschuk*
2. [***D19***] NestedNet: Learning Nested Sparse Structures in Deep Neural Networks, *Eunwoo Kim, Chanho Ahn, Songhwai Oh*
3. [***D21***] Context Embedding Networks, *Kun Ho Kim, Oisin Mac Aodha, Pietro Perona*
4. [***E1***] Iterative Learning With Open-Set Noisy Labels, *Yisen Wang, Weiyang Liu, Xingjun Ma, James Bailey, Hongyuan Zha, Le Song, Shu-Tao Xia*
5. [***E3***] Learning Transferable Architectures for Scalable Image Recognition, *Barret Zoph, Vijay Vasudevan, Jonathon Shlens, Quoc V. Le*
6. [***E5***] SBNet: Sparse Blocks Network for Fast Inference, *Mengye Ren, Andrei Pokrovsky, Bin Yang, Raquel Urtasun*
7. [***E7***] Language-Based Image Editing With Recurrent Attentive Models, *Jianbo Chen, Yelong Shen, Jianfeng Gao, Jingjing Liu, Xiaodong Liu*
8. [***E9***] Net2Vec: Quantifying and Explaining How Concepts Are Encoded by Filters in Deep Neural Networks, *Ruth Fong, Andrea Vedaldi*

9. [***E11***] End-to-End Dense Video Captioning With Masked Transformer, *Luowei Zhou, Yingbo Zhou, Jason J. Corso, Richard Socher, Caiming Xiong*
10. [***E13***] A Neural Multi-Sequence Alignment TeCHnique (NeuMATCH), *Pelin Dogan, Boyang Li, Leonid Sigal, Markus Gross*
11. [***E15***] Path Aggregation Network for Instance Segmentation, *Shu Liu, Lu Qi, Haifang Qin, Jianping Shi, Jiaya Jia*
12. [***E17***] The INaturalist Species Classification and Detection Dataset, *Grant Van Horn, Oisin Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, Serge Belongie*
13. [***E19***] Multimodal Explanations: Justifying Decisions and Pointing to the Evidence, *Dong Huk Park, Lisa Anne Hendricks, Zeynep Akata, Anna Rohrbach, Bernt Schiele, Trevor Darrell, Marcus Rohrbach*

1430–1450 Afternoon Break (Halls 1-4)

1450-1630 Session 3-3A: Machine Learning for Computer Vision V (Ballroom)

Papers in this session are also in Poster Session P3-2. Poster tag in square brackets (e.g., [***E21***])

Chairs: Jon Barron *(Google Research)*
Vladlen Koltun *(Intel)*

1450 Orals (O3-3A)

Format (12 min. for presentation + 2 min. for questions)

1. [***E21***] StarGAN: Unified Generative Adversarial Networks for Multi-Domain Image-to-Image Translation, *Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, Jaegul Choo*
2. [***F1***] High-Resolution Image Synthesis and Semantic Manipulation With Conditional GANs, *Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, Bryan Catanzaro*
3. [***F3***] Semi-Parametric Image Synthesis, *Xiaojuan Qi, Qifeng Chen, Jiaya Jia, Vladlen Koltun*

1534 Spotlights (S3-3A)

Format (4 min. for presentation; no questions)

1. [***F5***] BlockDrop: Dynamic Inference Paths in Residual Networks, *Zuxuan Wu, Tushar Nagarajan, Abhishek Kumar, Steven Rennie, Larry S. Davis, Kristen Grauman, Rogerio Feris*
2. [***F7***] Interpretable Convolutional Neural Networks, *Quanshi Zhang, Ying Nian Wu, Song-Chun Zhu*
3. [***F9***] Deep Cross-Media Knowledge Transfer, *Xin Huang, Yuxin Peng*
4. [***F11***] Interleaved Structured Sparse Convolutional Neural Networks, *Guotian Xie, Jingdong Wang, Ting Zhang, Jianhuang Lai, Richang Hong, Guo-Jun Qi*
5. [***F13***] A Variational U-Net for Conditional Appearance and Shape Generation, *Patrick Esser, Ekaterina Sutter, Björn Ommer*
6. [***F15***] Detach and Adapt: Learning Cross-Domain Disentangled Deep Representation, *Yen-Cheng Liu, Yu-Ying Yeh, Tzu-Chien Fu, Sheng-De Wang, Wei-Chen Chiu, Yu-Chiang Frank Wang*
7. [***F17***] Learning Deep Structured Active Contours End-to-End, *Diego Marcos, Devis Tuia, Benjamin Kellenberger, Lisa Zhang, Min Bai, Renjie Liao, Raquel Urtasun*
8. [***F19***] Deep Learning Under Privileged Information Using Heteroscedastic Dropout, *John Lambert, Ozan Sener, Silvio Savarese*
9. [***F21***] Smooth Neighbors on Teacher Graphs for Semi-Supervised Learning, *Yucen Luo, Jun Zhu, Mengxi Li, Yong Ren, Bo Zhang*
10. [***G1***] Interpret Neural Networks by Identifying Critical Data Routing Paths, *Yulong Wang, Hang Su, Bo Zhang, Xiaolin Hu*
11. [***G3***] Deep Spatio-Temporal Random Fields for Efficient Video Segmentation, *Siddhartha Chandra, Camille Couprie, Iasonas Kokkinos*
12. [***G5***] Customized Image Narrative Generation via Interactive Visual Question Generation and Answering, *Andrew Shin, Yoshitaka Ushiku, Tatsuya Harada*

1450-1630 Session 3-3B: Image Motion & Tracking (Room 155)

Papers in this session are also in Poster Session P3-2. Poster tag in square brackets (e.g., [***G7***])

Chairs: Katerina Fragkiadaki *(Carnegie Mellon Univ.)*
Nathan Jacobs *(Univ. of Kentucky)*

1450 Orals (O3-3B)

Format (12 min. for presentation + 2 min. for questions)

1. [***G7***] PWC-Net: CNNs for Optical Flow Using Pyramid, Warping, and Cost Volume, *Deqing Sun, Xiaodong Yang, Ming-Yu Liu, Jan Kautz*
2. [***G9***] Revisiting Deep Intrinsic Image Decompositions, *Qingnan Fan, Jiaolong Yang, Gang Hua, Baoquan Chen, David Wipf*
3. [***G11***] Multi-Cell Detection and Classification Using a Generative Convolutional Model, *Florence Yellin, Benjamin D. Haeffele, Sophie Roth, René Vidal*

1534 Spotlights (S3-3B)

Format (4 min. for presentation; no questions)

1. [***G13***] Learning Spatial-Aware Regressions for Visual Tracking, *Chong Sun, Dong Wang, Huchuan Lu, Ming-Hsuan Yang*
2. [***G15***] High Performance Visual Tracking With Siamese Region Proposal Network, *Bo Li, Junjie Yan, Wei Wu, Zheng Zhu, Xiaolin Hu*
3. [***G17***] LiteFlowNet: A Lightweight Convolutional Neural Network for Optical Flow Estimation, *Tak-Wai Hui, Xiaoou Tang, Chen Change Loy*
4. [***G19***] VITAL: VIsual Tracking via Adversarial Learning, *Yibing Song, Chao Ma, Xiaohe Wu, Lijun Gong, Linchao Bao, Wangmeng Zuo, Chunhua Shen, Rynson W.H. Lau, Ming-Hsuan Yang*
5. [***G21***] Super SloMo: High Quality Estimation of Multiple Intermediate Frames for Video Interpolation, *Huaizu Jiang, Deqing Sun, Varun Jampani, Ming-Hsuan Yang, Erik Learned-Miller, Jan Kautz*
6. [***H1***] Real-World Repetition Estimation by Div, Grad and Curl, *Tom F. H. Runia, Cees G. M. Snoek, Arnold W. M. Smeulders*
7. [***H3***] Recurrent Pixel Embedding for Instance Grouping, *Shu Kong, Charless C. Fowlkes*
8. [***H5***] Deep Unsupervised Saliency Detection: A Multiple Noisy Labeling Perspective, *Jing Zhang, Tong Zhang, Yuchao Dai, Mehrtash Harandi, Richard Hartley*
9. [***H7***] Learning Intrinsic Image Decomposition From Watching the World, *Zhengqi Li, Noah Snavely*
10. [***H9***] TieNet: Text-Image Embedding Network for Common Thorax Disease Classification and Reporting in Chest X-Rays, *Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Ronald M. Summers*
11. [***H11***] Generating Synthetic X-Ray Images of a Person From the Surface Geometry, *Brian Teixeira, Vivek Singh, Terrence Chen, Kai Ma, Birgi Tamersoy, Yifan Wu, Elena Balashova, Dorin Comaniciu*
12. [***H13***] Gibson Env: Real-World Perception for Embodied Agents, *Fei Xia, Amir R. Zamir, Zhiyang He, Alexander Sax, Jitendra Malik, Silvio Savarese*
13. [***H15***] Reinforcement Cutting-Agent Learning for Video Object Segmentation, *Junwei Han, Le Yang, Dingwen Zhang, Xiaojun Chang, Xiaodan Liang*

1450-1630 Session 3-3C: Machine Learning for Computer Vision VI (Room 255)

Papers in this session are also in Poster Session P3-2. Poster tag in square brackets (e.g., [***H17***])

Chairs: Kaiming He *(Facebook AI Research)*
Victor Lempitsky *(Skolkovo Inst. of Science & Tech.)*

1450 Orals (O3-3C)

Format (12 min. for presentation + 2 min. for questions)

1. [***H17***] Feature Space Transfer for Data Augmentation, *Bo Liu, Xudong Wang, Mandar Dixit, Roland Kwitt, Nuno Vasconcelos*
2. [***H19***] Analytic Expressions for Probabilistic Moments of PL-DNN With Gaussian Input, *Adel Bibi, Modar Alfadly, Bernard Ghanem*
3. [***H21***] Detail-Preserving Pooling in Deep Networks, *Faraz Saeedan, Nicolas Weber, Michael Goesele, Stefan Roth*

1534 Spotlights (S3-3C)

Format (4 min. for presentation; no questions)

1. [***I1***] Rethinking Feature Distribution for Loss Functions in Image Classification, *Weitao Wan, Yuanyi Zhong, Tianpeng Li, Jiansheng Chen*

2. [*I3*] Shift: A Zero FLOP, Zero Parameter Alternative to Spatial Convolutions, *Bichen Wu, Alvin Wan, Xiangyu Yue, Peter Jin, Sicheng Zhao, Noah Golmant, Amir Gholaminejad, Joseph Gonzalez, Kurt Keutzer*
3. [*I5*] Sketch-a-Classifier: Sketch-Based Photo Classifier Generation, *Conghui Hu, Da Li, Yi-Zhe Song, Tao Xiang, Timothy M. Hospedales*
4. [*I7*] Light Field Intrinsics With a Deep Encoder-Decoder Network, *Anna Alperovich, Ole Johannsen, Michael Strecke, Bastian Goldluecke*
5. [*I9*] Learning Generative ConvNets via Multi-Grid Modeling and Sampling, *Ruiqi Gao, Yang Lu, Junpei Zhou, Song-Chun Zhu, Ying Nian Wu*
6. [*I11*] Manifold Learning in Quotient Spaces, *Éloi Mehr, André Lieutier, Fernando Sanchez Bermudez, Vincent Guitteny, Nicolas Thome, Matthieu Cord*
7. [*I13*] Learning Intelligent Dialogs for Bounding Box Annotation, *Ksenia Konyushkova, Jasper Uijlings, Christoph H. Lampert, Vittorio Ferrari*
8. [*I15*] Boosting Adversarial Attacks With Momentum, *Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, Jianguo Li*
9. [*I17*] NISP: Pruning Networks Using Neuron Importance Score Propagation, *Ruichi Yu, Ang Li, Chun-Fu Chen, Jui-Hsin Lai, Vlad I. Morariu, Xintong Han, Mingfei Gao, Ching-Yung Lin, Larry S. Davis*
10. [*I19*] PointGrid: A Deep Network for 3D Shape Understanding, *Truc Le, Ye Duan*
11. [*I21*] Tell Me Where to Look: Guided Attention Inference Network, *Kunpeng Li, Ziyan Wu, Kuan-Chuan Peng, Jan Ernst, Yun Fu*
12. [*J1*] 3D Semantic Segmentation With Submanifold Sparse Convolutional Networks, *Benjamin Graham, Martin Engelcke, Laurens van der Maaten*
13. [*J3*] TOM-Net: Learning Transparent Object Matting From a Single Image, *Guanying Chen, Kai Han, Kwan-Yee K. Wong*

1630-1830 Poster Session P3-2 (Halls D-E)

Poster tag in square brackets (e.g., [*J5*])

Biomedical Image

1. [*J5*] Translating and Segmenting Multimodal Medical Volumes With Cycle- and Shape-Consistency Generative Adversarial Network, *Zizhao Zhang, Lin Yang, Yefeng Zheng*
2. [*J7*] An Unsupervised Learning Model for Deformable Medical Image Registration, *Guha Balakrishnan, Amy Zhao, Mert R. Sabuncu, John Guttag, Adrian V. Dalca*
3. [*J9*] Deep Lesion Graphs in the Wild: Relationship Learning and Organization of Significant Radiology Image Findings in a Diverse Large-Scale Lesion Database, *Ke Yan, Xiaosong Wang, Le Lu, Ling Zhang, Adam P. Harrison, Mohammadhadi Bagheri, Ronald M. Summers*
4. [*J11*] Learning Distributions of Shape Trajectories From Longitudinal Datasets: A Hierarchical Model on a Manifold of Diffeomorphisms, *Alexandre Bône, Olivier Colliot, Stanley Durrleman*
5. [*J13*] CNN Driven Sparse Multi-Level B-Spline Image Registration, *Pingge Jiang, James A. Shackleford*
6. [*J15*] Anatomical Priors in Convolutional Networks for Unsupervised Biomedical Segmentation, *Adrian V. Dalca, John Guttag, Mert R. Sabuncu*
7. [*J17*] 3D Registration of Curves and Surfaces Using Local Differential Information, *Carolina Raposo, João P. Barreto*
8. [*J19*] Weakly Supervised Learning of Single-Cell Feature Embeddings, *Juan C. Caicedo, Claire McQuin, Allen Goodman, Shantanu Singh, Anne E. Carpenter*
9. [*J21*] Guided Proofreading of Automatic Segmentations for Connectomics, *Daniel Haehn, Verena Kaynig, James Tompkin, Jeff W. Lichtman, Hanspeter Pfister*

Machine Learning for Computer Vision

10. [*K1*] Wide Compression: Tensor Ring Nets, *Wenqi Wang, Yifan Sun, Brian Eriksson, Wenlin Wang, Vaneet Aggarwal*
11. [*K3*] Improvements to Context Based Self-Supervised Learning, *T. Nathan Mundhenk, Daniel Ho, Barry Y. Chen*
12. [*K5*] Learning Structure and Strength of CNN Filters for Small Sample Size Training, *Rohit Keshari, Mayank Vatsa, Richa Singh, Afzel Noore*
13. [*K7*] Boosting Self-Supervised Learning via Knowledge Transfer, *Mehdi Noroozi, Ananth Vinjimoor, Paolo Favaro, Hamed Pirsiavash*

14. [***K9***] The Power of Ensembles for Active Learning in Image Classification, *William H. Beluch, Tim Genewein, Andreas Nürnberger, Jan M. Köhler*
15. [***K11***] Learning Compact Recurrent Neural Networks With Block-Term Tensor Decomposition, *Jinmian Ye, Linnan Wang, Guangxi Li, Di Chen, Shandian Zhe, Xinqi Chu, Zenglin Xu*
16. [***K13***] Spatially-Adaptive Filter Units for Deep Neural Networks, *Domen Tabernik, Matej Kristan, Aleš Leonardis*
17. [***K15***] SO-Net: Self-Organizing Network for Point Cloud Analysis, *Jiaxin Li, Ben M. Chen, Gim Hee Lee*
18. [***K17***] SGAN: An Alternative Training of Generative Adversarial Networks, *Tatjana Chavdarova, François Fleuret*
19. [***K19***] SketchyGAN: Towards Diverse and Realistic Sketch to Image Synthesis, *Wengling Chen, James Hays*
20. [***K21***] Explicit Loss-Error-Aware Quantization for Low-Bit Deep Neural Networks, *Aojun Zhou, Anbang Yao, Kuan Wang, Yurong Chen*
21. [***L1***] Towards Universal Representation for Unseen Action Recognition, *Yi Zhu, Yang Long, Yu Guan, Shawn Newsam, Ling Shao*
22. [***L3***] Deep Image Prior, *Dmitry Ulyanov, Andrea Vedaldi, Victor Lempitsky*

Applications

23. [***L5***] ST-GAN: Spatial Transformer Generative Adversarial Networks for Image Compositing, *Chen-Hsuan Lin, Ersin Yumer, Oliver Wang, Eli Shechtman, Simon Lucey*
24. [***L7***] CartoonGAN: Generative Adversarial Networks for Photo Cartoonization, *Yang Chen, Yu-Kun Lai, Yong-Jin Liu*

END OF MAIN CVPR 2018 CONFERENCE

Notes:

GOLD SPONSORS

SILVER SPONSORS

DIAMOND SPONSORS

PLATINUM SPONSORS