

## Message from the General and Program Chairs

Welcome to Columbus, Ohio and the 27th IEEE Conference on Computer Vision and Pattern Recognition (CVPR). In addition to the main four-day program of oral and poster presentations, keynote talks, demos, exhibitions, and social functions, CVPR 2014 has a number of co-located events, including 19 workshops and 15 tutorials.

This year, we received 1807 valid submissions to the main conference, of which 1778 were fully reviewed. (The others were either administratively rejected for technical or ethical reasons or withdrawn before review.) To select papers from these submissions, we invited 55 highly regarded researchers to act as Area Chairs (ACs). ACs were selected to include a broad range of geographical locations while providing balance in both gender and seniority. No more than one AC per institution was selected. The numbers of senior and junior ACs were approximately equal to allow the pairing of more experienced ACs with less experienced ones during paper discussions.

We recruited an expert team of 981 experienced reviewers from the broader CV&PR community. Our criteria to select reviewers were a proven publication record in the top computer vision conferences and journals and a completed (or close to completed) doctoral degree. The original list of reviewers was slightly increased by additional reviewers recommended by the ACs to add expertise for papers where appropriate reviewers were not available.

As in previous years, we used the CMT conference management service provided by Microsoft Research to manage the submission and selection of papers from beginning to end. Also, for the first time, we used the iThenticate software to detect potential cases of plagiarism.

After the submission deadline, the Program Chairs distributed the papers to the ACs with help from the automated Toronto Paper Matching System (TPMS), developed by Charlin and Zemel [ICML 2013]. TPMS suggests matches between papers and ACs based on the PDF files of submitted manuscripts and representative publications by each AC; for CVPR 2014 we had two Technology Chairs who were in

charge of the interface with TPMS. The ACs, in turn, used TPMS to help them determine the potential reviewers for each of their assigned papers. ACs suggested 5-10 reviewers per paper from which CMT automatically selected three non-conflicted reviewers per paper. Finally, manual adjustments were made by the ACs and Program Chairs to achieve better matches between the papers and reviewers. As a result, each reviewer was assigned a maximum of 9 papers, with a median load of 7 papers.

Two key innovations to the review form for CVPR 2014 were to include an 'oral/poster' rating and to eliminate the borderline rating. Reviewers were given six weeks to complete their reviews, at which time the ACs stepped back in to vet the reviews for quality (initiating discussions, where necessary) and write an initial consolidation report of the reviews before releasing them to the authors, who were given one week to write rebuttals. After collecting the rebuttals, the ACs finished their pre-meeting work, i.e., consolidating the reviews and author rebuttals, initiating discussions among the reviewers for clarification, and making recommendations for decisions on papers. Reviewers were asked to carefully read and respond to the authors' rebuttals. The Program Chairs and the ACs strove to ensure that every paper eligible for full review received at least three high quality reviews and a thoughtful final consolidation report that summarized the reviewer discussion

Every paper, review, and author rebuttal was evaluated by a primary AC, who made a preliminary recommendation on the paper. All papers that were recommended for rejection were then reviewed by a secondary AC. Papers that were not considered acceptable by both ACs were rejected before the AC meeting, and the AC wrote a consolidation report explaining the decision. The 984 papers that passed this review were discussed at the AC meeting held in February at the University of Maryland.

At the AC meeting, the ACs were divided into 18 panels of three members each, with no conflicts between the ACs and papers associated with each

1

## Message from the General and Program Chairs

panel. The Program Chairs monitored the panel discussions and worked hard to maintain consistency among the panels. Decisions to accept or reject a paper were made by the three ACs working together and, as needed, by input from additional non-conflicted ACs. By the end of the discussion, the ACs were asked to produce detailed consolidation reports to justify all their decisions. The panels selected 540 papers for presentation at the main conference (a 29.88% acceptance rate). Papers were selected on merit, not to fulfil any quotas. Panels of at least six non-conflicting ACs were assembled to evaluate accepted papers. The panels selected 104 of these papers (5.75%) to be presented as orals at the conference.

The Program Chairs and General Chairs did not submit any papers to CVPR 2014, allowing them to work without any direct conflicts throughout the review process. (The General Chairs were also not present at the AC Meeting.) Additionally, the respective ACs were excluded from any decisions associated with papers from their research groups, affiliated institutions or collaborators. The double-blind nature of the CVPR review process was thus strictly maintained throughout.

The overall acceptance rate was comparable to those of previous CVPRs. However, the percentage of oral presentations was significantly increased from recent years, from 3.3% last year to 5.75% this year. This was a deliberate decision based on extensive consultation with members of the CVPR community. The main conference at CVPR was scheduled as a three-day event many years ago, when the number of submissions was fewer than half of that seen in recent years. As the number of submissions has grown, the percentage of oral presentations has decreased. At first, this increased the quality of oral presentations, but it eventually became frustrating for reviewers, ACs, and authors that many excellent papers could not be selected for oral presentation.

The main conference this year has been expanded from three to four days, at the expense of losing a workshop/tutorial day (to maintain the same overall

conference length). Poster spotlights have been moved to video spotlights (available for oral presenters as well), which are available online before, during and after the conference. This has allowed for a significant increase in oral presentations to provide more authors an opportunity to present their research to a broader audience.

Another innovation of CVPR 2014 is the creation of a PAMI/IJCV journal session. The goal of this session is to expose the CVPR community to high-impact journal publications that have appeared in recent years and were not previously presented at a major computer vision conference. Six recent high-impact journal papers were selected by a committee consisting of two former Editors-in-Chief of PAMI. These papers were selected from a pool of 20+ papers nominated by Associate Editors of PAMI and IJCV.

The main meeting also includes two invited speakers, Dr. Doris Tsao and Dr. Stéphane Mallat.

The proceedings of CVPR 2014 are being provided on USB drives at the conference. All papers in the main conference and associated workshops will be made available through the IEEE Computer Society Digital Library and through IEEE Xplore.

We wish to thank all members of the Organizing Committee, the Area Chairs, reviewers, authors, and the CMT and TPMS teams for the immense amount of hard work and professionalism that has gone into making CVPR 2014 a first-rate conference. Our thanks also go to the organizers of previous CVPRs for their helpful advice and support. We are grateful to the sponsors as well, and we are happy to report that CVPR 2014 has seen a significant increase in industrial support, which is further evidence of the relevance and importance of this community.

Finally, we wish all the attendees a highly stimulating, informative, and enjoyable conference.

Sven Dickinson, Dimitri Metaxas, and Matthew Turk, General Co-Chairs

Ronen Basri, Cornelia Fermuller, Aleix Martinez, and René Vidal, Program Co-Chairs

## Organizing Committee, Area Chairs, & Outstanding Reviewers

CVPR 2014 Organizing Committee

 General Chairs:
 Sven Dickinson
 Publications Chairs:
 Eric Mortensen

 Dimitri Metaxas
 Sanja Fidler

**Program Chairs:** 

Matthew Turk Demos/Exhibitions Chair: Michael Brown
Ronen Basri Corporate Relations Chairs: Kari Pulli
Cornelia Fermuller

Workshops Chairs: Jim Davis Doctoral Consortium Chairs: Philippos Mordohai Raquel Urtasun

Margrit Betke Publicity Chair: Kristin Dana
Tutorials Chairs: Kristen Grauman

Raquel Urtasun

Student Activities Chair: Jason Corso

Video Proceedings Chair: Joannis Kakadiaris

Student Activities Chair: Jason Corso

Finance Chairs: Terry Boult Octavia Camps Local Arrangements Chair: Jason Coiso

Website Chairs: Ryan Farrell Technology Chairs: Yaser Sheikh

G<sup>´</sup>ary Huang Jingyi Yu Yaser Yacoob **Logistics Advisor:** Ginger Boult

CVPR 2014 Area Chairs

Frik Learned-Miller Leo Grady Raquel Urtasun Lourdes Agapito James Rehg Tal Arbel Edwin Hancock Zhouchen Lin Arun Ross Manik Varma Kobus Barnard Tal Hacener Haibin Ling Sudeep Sarkar Nuno Vasconcelos David Jacobs Yoichi Sato Olga Veksler Tamara Berg Simon Lucey Hui li Konrad Schindler Horst Bischof Yi Ma Baha Vemuri Daniel Cremers Frederic Jurie Jiri Matas Thomas Serre Daphna Weinshall Rita Cucchiara Fredrik Kahl Gerard Medioni Rahul Sukthankar Allen Yang Kostas Daniilidis Ioannis Kakadiaris Sinisa Todorovic Greg Mori Ming-Hsuan Yang Piotr Dollar Jana Kosecka Nikos Paragios Carlo Tomasi Laurent Younes Ahmed Elgammal Christoph Lampert Robert Pless Fernando de la Torre Lihi Zelnik-Manor Paolo Favaro Svetlana Lazebnik Petia Radeva 7huowen Tu

CVPR 2014 Outstanding Reviewers

We are pleased to recognize the following researchers as "Outstanding Reviewers for CVPR 2014". These reviewers were selected from almost 1000 reviewers for their hard work Area Chairs, who found their reviews of high quality.

Elli Angelopoulou Eric Mortensen **David Suter** Jason Corso Anthony Hoogs Joao Barreto James Crowlev Omar Javed Ko Nishino Yuichi Taguchi Ohad Ben-Shahar Alessio Del Bue Bjorn Ommer Hao Jiang Robby Tan Moshe Ben-ezra Roland Kwitt Devi Parikh Roberto Tron Konstantinos Derpanis Margrit Betke Jan-Michael Frahm Jean-Francois Lalonde Nikhil Rasiwasia Stefan Walk Fabio Galasso Stefan Roth Ross Beveridge Ivan Laptev Jingdong Wang Mohammad Saberian Thomas Brox Peter Geheler Subhransu Maji Lior Wolf Octavia Camps Christopher Gever Tim Marks Albert Salah John Wright Antoni Chan Michal Havlena Yasuvuki Matsushita Radim Sara Larry Zitnick Rainer Stiefelhagen Tsung-han Chan Xuming He Peyman Milanfar **Bob Collins** Adrian Hilton Francesc Moreno-Noquer Yusuke Sugano

## Monday, June 23

**0700–1700 Registration** (Exhibit Hall C Lobby)

0730-0830 Breakfast (Exhibit Hall C)

**1200–1330 Lunch** (Exhibit Hall C)

### **Perceptual Organization**

Organizers: Michael Maire

Stella Yu

**Location:** C112-113 **Schedule:** Full Day

o830 Welcome

### S1: Segmentation & Video (0830-1010)

- o830 Moving vs. Static Objects in Video Segmentation, Thomas Brox (Univ. of Freiburg)
- o850 Composite Statistical Learning and Inference in Semantic and Video Segmentation, Fuxin Li (Georgia Institute of Technology)
- og1o Interactive Learning for Point-Cloud Motion Segmentation, *Tal Hassner (Open Univ. of Israel)*
- og3o Learning from YouTube Videos, Rahul Sukthankar (Google)
- 0950 Spotlight Presentations

### 1000 Morning Break

### S2: Representation (1030-1200)

- 1030 From Edges to Objects, *Piotr Dollar (Microsoft Research)*
- 1050 Perceiving Crowds: Bypassing the Bottleneck of Conscious Vision, David Whitney (UC Berkeley)
- 1130 Hierarchy, Reasoning, and Representation Learning, Yann LeCun (New York Univ.)

#### 1200 Lunch Break (Exhibit Hall C)

### S3: Activity (1330-1500)

- 1330 Live Counting of Repetitive Actions via a Convolutional Deep Neural Network Trained on Unrealistic Synthetic Data, Lior Wolf (Tel Aviv Univ.)
- 1350 Weakly-Supervised Learning of Actions, *Ivan Laptev*(INRIA)
- 1410 Recognizing Human Activities: Interactions, Groups, and Context, *Greg Mori (Simon Fraser Univ.)*
- 1430 Video Scene Segmentation and Recognition by Location-Independent Activity Classes, Anthony Hoogs (Kitware)
- 1450 Spotlight Presentations
- 1500 Afternoon Break

#### S4: Motion, Shape, & Parts (1520-1700)

- 1520 From Shallow Hollywood Motions to Deep Academic Gestures, *Christoph Bregler (New York Univ.)*
- 1540 Perceptual Organization of Motion: Event Recognition and the Perception of 3D Structure from Motion, James Todd (Ohio State Univ.)
- 1620 Data-Driven Perceptual Interpretation of Shape,
  William Freeman (MIT)
- 1640 Do Mid-Level Parts Still Matter in the Age of CNNs? *Lubomir Bourdev (Facebook)*

### **Mobile Vision**

Organizers: Zhengyou Zhang

Marc Pollefeys Gang Hua Matthew Turk Kari Pulli Raja Bala

**Location:** Grand Ballroom 1

Schedule: Full Day

0825 Opening Remarks, Gang Hua

#### S1: Mobile Visual Recognition and Search (0830-0930)

- o83o Fast Target Recognition on Mobile Devices: Revisiting Gaussian Elimination for the the Estimation of Planar Homographies, Olexa Bilaniuk, Hamid Bazargani, Robert Laganière
- 0850 Cascade of Box (CABOX) Filters for Optimal Scale Space Approximation, Victor Fragoso, Gaurav Srivastava, Abhishek Nagar, Zhu Li, Kyungmo Park, Matthew Turk
- og1o Real-time Mobile Facial Expression Recognition
  System A Case Study, Myunghoon Suk, Balakrishnan
  Prabhakaran
- og3o Invited Talk: Towards Ubiquitous Embedded 3D Visual Sensing in Mobile Devices, Achintya Bhowmik (Intel Corporation)

### 1015 Morning Break

## S2: Mobile Computational Photography and Multiview Analysis (1045-1145)

- 1045 Dense View Interpolation on Mobile Devices using Focal Stacks, Parikshit Sakurikar, P. J. Narayanan
- 1105 Dynamic Image Stacks, *David E. Jacobs, Orazio Gallo, Kari A. Pulli*
- 1125 Robust Three-view Triangulation Done Fast, Johan Hedborg, Andreas Robinson, Michael Felsberg

#### 1200 Lunch Break (Exhibit Hall C)

1345 Invited Talk: Mobile Imaging: The Future Of The Image, Edward J. Delp (Purdue Univ.)

## S3: Mobile 3D Modeling and Other Mobile Related Applications (1425-1525)

- 1425 3D Hallway Modeling Using A Single Image, *Greg Olmschenk, Zhigang Zhu*
- 1445 Estimating Gaze Direction of Vehicle Drivers using a Smartphone Camera, Meng-Che Chuang, Raja Bala, Edgar A. Bernal, Peter Paul, Aaron Burry
- 1505 GPS Refinement and Camera Orientation Estimation from a Single Image and a 2D Map, *Hang Chu, Andrew Gallagher, Tsuhan Chen*

### 1525 Afternoon Break

#### S4: Other Mobile Related Applications (1555-1635)

- 1555 Fast and Robust Object Detection Using Visual Subcategories, Eshed Ohn-Bar, Mohan M. Trivedi
- 1615 Vision on Wheels: Looking at Driver, Vehicle, and Surround for On-Road Maneuver Analysis, Eshed Ohn-Bar, Ashish Tawari, Sujitha Martin, Mohan M. Trivedi

#### S5: Demos (1640-1730)

- 1640 Space-Variant Image Deblurring on Smartphones using Inertial Sensors, Ondřej Šindelář, Filip Šroubek, Peyman Milanfar
- 1650 Offline 1000-Class Classification on a Smartphone, Yoshiyuki Kawano, Keiji Yanai
- 1700 A Compact 3D Camera Suited for Mobile and Embedded Vision Applications, Stefano Mattoccia, Ilario Marchio, Marco Casadio
- 1710 Fast and Robust Perspective Rectification of Document Images on a Smartphone, Williem, Christian Simon, Sungdae Cho, In Kyu Park
- 1720 Correcting Photometric Distortion of Document Images on a Smartphone, Christian Simon, Williem, Jihwan Choe, Il Dong Yun, In Kyu Park
- 1730 Best Paper Award Announcement (Sponsored by Microsoft Research and Nvidia Research)

### **Scene Understanding**

**Organizers:** James Hays

Derek Hoiem Aditya Khosla Jianxiong Xiao

Location: Exhibit Hall C

Schedule: Full Day

o830 Welcome by Organizers

o835 **Invited Talk:** Irving Biederman (Univ. of Southern California)

0905 **Invited Talk:** Jitendra Malik (Univ. of California, Berkeley)

og35 **Sponsor Invited Talk:** Chang Huang (Baidu Inst. of Deep Learning)

0945 Sponsor Invited Talk: Yuanging Lin (NEC Labs)

1000 Morning Break

1030 **Invited Talk:** Alexei A Efros (Univ. of California, Berkeley)

1100 Invited Talk: Ashutosh Saxena (Cornell University)

1130 Poster Spotlights

1200 Lunch Break (Exhibit Hall C)

1330 Poster Session

1500 Afternoon Break

1530 Invited Talk: David Forsyth (Univ. of Illinois, Urbana-Champaign)

1600 Invited Talk: Rob Fergus (New York Univ.)

1630 Sponsor Invited Talk: Rahul Sukthankar (Google Research)

1640 Invited Talk: Antonio Torralba (Massachusetts Inst. of Technology)

1710 Invited Talk: Larry Zitnick (Microsoft Research)

1720 **Invited Talk:** Martial Hebert (Carnegie Mellon University)

1750 Sponsor Invited Talk: Marc'Aurelio Ranzato (Facebook Al Research)

1800 Award Ceremony, Workshop Chairs

### **Vision Meets Cognition**

Organizers: Yibiao Zhao

Lap-Fai Yu Bo Zheng Peter Battaglia Tao Gao

C213-215

**Schedule:** Full Day 0830 Welcome Message

Location:

0840 Invited Talk: Beyond What and Where: Joint Spatial, Temporal and Causal Parsing with Commonsense Reasoning, Song-Chun Zhu (UCLA)

0925 Invited Talk: Learning-from-Observation: From Assembly Robot Through Dancing Humanoid, Katsushi Ikeuchi (Univ. of Tokyo)

og55 Invited Talk: Seeing Time's Arrow / Inferring Properties of Cloth from Watching it Move, Bill Freeman (MIT)

1025 Morning Break

1045 Invited Talk: Cognitive Vision and its Application to Medical Imaging, Visual Surveillance, and Space Robotics, Demetri Terzopoulos (UCLA)

1115 Posters

1200 Lunch Break (Exhibit Hall C)

1300 Posters

1355 Invited Talk: Modeling Human Common-Sense Scene Understanding, Josh Tenenbaum (MIT)

1440 Invited Talk: Physics, Humans, and Intention: Interpretable Machine Learning for 3D Scene Understanding, Ashutosh Saxena (Cornell Univ.)

1510 **Invited Talk:** Vision is for Agents, *Benjamin Kuipers* (*Univ. of Michigan*)

1540 Afternoon Break

1600 Invited Talk: Seeing Into the Future, Larry Zitnick (Microsoft Research)

1630 Panel discussion

## Monday, June 23 Workshops

### Vision Industry & Entrepreneur Workshop

Organizers: Sek Chai

**Boaz Super** 

Himanshu Arora

Terrance Boult

Arnab Dhua

Marshall Tappen

Raja Bala Yu Wang

Location: Grand Ballroom 3

Schedule: Full Day

o835 Welcome: Sek Chai (SRI), Himanshu Arora (A9.com)

#### S1: Mobile Platforms for Computer Vision (0845-1105)

o845 Invited Talk: Mobile Computational Imaging, Kari Pulli (Nvidia)

og25 **Invited Talk**: Computer Vision Applications for Mobile and Beyond, *Mahesh Ramachandran (Qualcomm)* 

#### 1005 Morning Break

1025 Invited Talk: OpenVX: The Computer Vision Hardware Abstraction Layer, Victor Eruhimov (Khronos)

### S2: Distinguished Speaker (1105-1145)

1105 Invited Talk: The Pit and the Pendulum: Academic Research in Industrial Labs, Rahul Sukthankar (Google)

### S3: Industry Session: Demos, Posters, Recruiting (1145— 1500)

1145 Industry Session Spotlights: Moderators: Arnab Dhua (Ag.com), Marshall Tappen (Amazon)

### 1210 Lunch Break (Exhibit Hall C)

### 1315 Demos, Posters, Recruiting

- Computer Vision for Enterprise and Public Safety at Motorola Solutions, Ankur Patel
- Xerox Computer Vision for Roadway Transportation Systems: Robust, High Yield Automated License Plate Recognition, Aaron Burry, Vladimir Kozitsky
- Xerox Computer Vision for Roadway Transportation Systems: Automated Image-Based Detection of Front Seat Passengers in Vehicles, Aaron Burry, Peter Paul, Yusuf Artan, Florent Perronnin

- Video Analytics at United Technologies Research Center, Alan Finn
- Real-time Urban Metering with Pedestrian and Vehicle Recognition, Alexandre Winter, Ignacio Mellado Bataller, Tuan Thi
- Collaborative Computer Vision R&D at Kitware, Arslan Basharat, Sangmin Oh, Matt Leotta, Rusty Blue, Keith Fieldhouse, Matt Turek, Brad Davis, Heather James, Anthony Hoogs
- TeraDeep: Intelligent Vision Systems, *Eugenio Culurciello*
- ViPanix: Panoramic Videos, Francisco Hernandez-Lopez, Mariano Rivera
- Enhancing Confidence in Video Analytics, Gary Rubin, David Berger
- PercepTonic: We See Solutions, Goksel Dedeoglu, Susan Rossbach
- Real-time Image Classification on Mobile Phones, Harro Stokman, Samir Kumar, Daniel Fontijne, Ork de Rooij
- Amazon Fulfillment Technology Computer Vision, Marshall Tappen
- ReKognition API Platform, Meng Wang, Tianqiang Liu, Yushan Chen
- Computer Vision at Eyenuk: Image Analysis for Your Health and Your Photos, Sandeep Bhat, Chaithanya Ramachandra, Malavika Bhaskaranand, Kaushal Solanki
- SRI International: Breakthrough Ideas... Real-World Solutions, *Sek Chai*
- Visual Search Technology for Amazon, Sunil Ramesh, Arnab Dhua, Himanshu Arora

### 1500 Afternoon Break

#### S4: Computer Vision in Services Industry (1525-1645)

- 1525 Invited Talk: An Overview of Retail Video Analytics, Quanfu Fan (IBM)
- 1605 **Invited Talk:** Computer Vision with an Eye on Services, Raja Bala and Peter Paul (Xerox)

#### S5: Panel Session (1645-1730)

- 1645 **Panel:** Computer Vision Industry and Community, *Moderator: Terrance Boult (UCCS)*
- 1720 **Beyond VIEW 2014**: Terrance Boult (UCCS), Sek Chai (SRI)

# Computer Vision & Human Computation

Organizers: Jia Deng

Subhransu Maji Pietro Perona

**Location:** C114-115 **Schedule:** Full Day og2o Opening Remarks

0925 Invited Talk: Building Large Datasets to Represent the World, Larry Zitnick (Microsoft Research)

### 1005 Morning Break

1030 Invited Talk: Beyond Mindless Labeling: \*Really\* Leveraging Humans to Build Intelligent Machines, *Devi Parikh (Viginia Tech)* 

1105 Spotlights and Posters

#### 1200 Lunch Break (Exhibit Hall C)

1330 **Invited Talk:** Quickly Answering General Visual Questions, *Jeffrey Bigham (CMU)* 

1405 **Keynote Talk:** EyeWire, A Game to Map the Brain, Sebastian Seung (Princeton Univ.)

### 1500 Afternoon Break

1530 Invited Talk: Crowd One Shot Learning, James Hays (Brown Univ.)

1605 Invited Talk: Visipedia Tool Ecosystem for Dataset
Curation and Annotation, Serge Belongie (Cornell Tech.)

# Perception Beyond the Visible Spectrum

Organizers: Riad I. Hammoud

Guoliang Fan Firooz Sadjadi

Behzad Kamgar-Parsi

**Location:** C110-111

Schedule: Half Day — Morning

o8oo Welcome Message

0810 Keynote Talk: Multi-frame Data Association with Higher-Order Cost Functions, Erik P. Blasch (Air Force Research Lab)

### S1: Thermal & Infrared Imaging (0845-0930)

o845 A Thermal Infrared Video Benchmark for Visual Analysis, Zheng Wu, Nathan Fuller, Diane Theriault, Margrit Betke

ogoo Low Resolution Person Detection with a Moving Thermal Infrared Camera by Hot Spot Classification, Michael Teutsch, Thomas Müller, Marco Huber, Jürgen Beyerer

og15 Improving Person Tracking Using an Inexpensive
Thermal Infrared Sensor, Suren Kumar, Tim K. Marks,
Michael Iones

### S2: Activity Recognition & Surveillance (0930-1015)

og3o Driver Cell Phone Usage Detection From HOV/HOT NIR Images, Yusuf Artan, Orhan Bulan, Robert P. Loce, Peter Paul

og45 Ground-Based Activity Recognition at Distance and Behind Wall, *Tao Wang, Riad Hammoud, Zhigang Zhu* 

1000 Multi-Source Multi-Modal Activity Recognition in Aerial Video Surveillance, Riad I. Hammoud, Cem S. Sahin, Erik P. Blasch, Bradley J. Rhodes

1015 Morning Break: Poster preparation

### S3: Point Registration, 3D Estimation, and 3D Segmentation (1045-1130)

- 1045 Non-rigid Point Set Registration with Global-Local Topology Preservation, *Song Ge, Guoliang Fan, Meng Ding*
- 1100 3D Scene Estimation with Perturbation-Modulated Light and Distributed Sensors, Quan Wang, Xinchi Zhang, Kim L. Boyer
- 1115 Edge-Weighted Centroid Voronoi Tessellation with Propagation of Consistency Constraint for 3D Grain Segmentation in Microscopic Superalloy Images, Youjie Zhou, Lili Ju, Yu Cao, Jarrell Waggoner, Yuewei Lin, Jeff Simmons, Song Wang

#### S4: Posters (1130-1230)

- Joint Shape and Texture Based X-Ray Cargo Image Classification, Jian Zhang, Li Zhang, Ziran Zhao, Yaohong Liu, Jianping Gu, Qiang Li, Duokun Zhang
- Use of Sparse Representation for Pedestrian Detection in Thermal Images, Bin Qi, Vijay John, Zheng Liu, Seiichi Mita
- A Photon-Mapping Informed Chan-Vese Segmentation Algorithm to Enable Multispectral Sensing and Path-Planning in 3D Virtual Environments, Bruce A. Johnson, Hairong Qi, Jason C. Isaacs
- Superpixel Estimation for Hyperspectral Imagery, Pegah Massoudifar, Anand Rangarajan, Paul Gader
- Automatic Target Recognition in Infrared Imagery
  Using Dense HOG Features and Relevance Grouping of
  Vocabulary, M.N.A. Khan, Guoliang Fan, Douglas R.
  Heisterkamp, Liangjiang Yu
- Ego-Motion Estimation on Range Images using High-Order Polynomial Expansion, Brian Okorn, Josh Harquess

### **Registration of Very Large Images**

Organizers: Ardy Goshtasby

Chang Shu

Akihiro Sugimoto

John Camp

Nathan Netanyahu

Claude Cariou Hector Erives Clark Taylor

Lyubomir Zagorchev

Martin Satter Marcel Jackowski

Simon Warfield

**Location:** C110-111

Schedule: Half Day — Afternoon

- 1300 Invited Talk: Efficient High-Resolution Stereo Matching using Local Plane Sweeps, Sudipta Sinha (Microsoft Research)
- 1345 Automatic Geo-location Correction of Satellite Imagery, Ozge C. Ozcanli, Yi Dong, Joseph L. Mundy, Helen Webb, Riad Hammoud, Tom Victor
- 1410 Efficient Change Detection for Very Large Motion Blurred Images, Vijay Rengarajan, Abhijith Punnappurath, A.N. Rajagopalan, Guna Seetharaman
- 1435 Poster: Non-rigid Registration of 3D Ultrasound Images Using Model-based Segmentation, Babak Matinfar, Lyubomir Zagorchev
- 1445 Poster: Image Registration of Very Large Images via Genetic Programming, Sarit Chicotay, Omid E. David, Nathan S. Netanyahu

### 1455 Afternoon Break

- 1530 Invited Talk: Representing 3D Models with
  Discriminative Visual Elements, Mathieu Aubry (INRIA)
- 1615 Efficient and Automated Multimodal Satellite Data Registration Through MRFs and Linear Programming, Konstantinos Karantzalos, Aristeidis Sotiras, Nikos Paragios

## Monday, June 23 Workshops

- 1640 Poster: Variational Deformation Method for the Computation of the Average Shape of Organs, *Shun* Inagaki, Atsushi Imiya
- 1650 Poster: Adaptive Registration of Very Large Images, Brian P. Jackson, A. Ardeshir Goshtasby

### 1330 Effect of Pupil Dilation and Constriction on the Distribution of Bit Errors within the Iris, Inmaculada Tomeo-Reyes, Vinod Chandran

- 1335 Optimization of Iris Codes for Improved Recognition, Nitin K. Mahadeo, Andrew P. Papliński, Sid Ray
- 1340 Invited Talk: Object Detection with Deep Neural
  Network, Dumitru Erhan (Google)

#### S2: Extended Poster Spotlights (1430-1510)

- 1430 Reliable Posterior Probability Estimation for Streaming Face Recognition, Abhijit Bendale, Terrance Boult
- 1435 Learning Minutiae Neighborhoods: A New Binary Representation for Matching Fingerprints, Akhil Vij, Anoop Namboodiri
- 1440 Performance Improvement of Phase-Based Correspondence Matching for Palmprint Recognition, Vincent Roux, Shoichiro Aoyama, Koichi Ito, Takafumi Aoki
- 1445 A Robust Approach for Singular Point Extraction Based on Complex Polynomial Model, *Jin Qi, Suxing Liu*
- 1450 Secure Fingerprint Matching With Generic Local
  Structures, Matthew Morse, Jesse Hartloff, Thomas
  Effland, Jim Schuler, Jennifer Cordaro, Sergey Tulyakov,
  Atri Rudra, Venu Govindaraju
- 1455 The Value of Multiple Viewpoints in Gesture-Based User Authentication, Jonathan Wu, Janusz Konrad, Prakash Ishwar
- 1500 Context-Aware Active Authentication Using Smartphone Accelerometer Measurements, Abena Primo, Vir V. Phoha, Rajesh Kumar, Abdul Serwadda
- 1505 Can We Use Second Minor Finger Knuckle Patterns to Identify Humans?, *Ajay Kumar, Zhihuan Xu*

#### 1510 Afternoon Break and Poster Session

- 1600 Invited Talk: Face Biometrics under Spoofing Attacks: Vulnerabilities, Countermeasures, Open Issues and Research Directions, Abdenour Hadid (Univ. of Oulu)
- 1650 Valedictory, Awards and Closing Remarks

### **Biometrics**

Organizers: Bir Bhanu

Ross Beveridge Ajay Kumar

**Location**: C220-221

Schedule: Half Day — Afternoon

#### S1: Extended Poster Spotlights (1300-1340)

- 1300 Hallucinating the Full Face from the Periocular Region via Dimensionally Weighted K-SVD, Felix Juefei-Xu, Dipan K. Pal, Marios Savvides
- 1305 Improving 3D Face Details based on Normal Map of Hetero-source Images, Chang Yang, Jiansheng Chen, Nan Su, Guangda Su
- 1310 Globality-Locality Preserving Projections for Biometric Data Dimensionality Reduction, Sheng Huang, Ahmed Elgammal, Luwen Huangfu, Dan Yang, Xiaohong Zhang
- 1315 Robust Low-Rank Regularized Regression for Face Recognition with Occlusion, Jianjun Qian, Jian Yang, Fanglong Zhang, Zhouchen Lin
- 1320 Natural vs Artificial Face Classification using Uniform Local Directional Patterns and Wavelet Uniform Local Directional Patterns, Darryl D'Souza, Roman V. Yampolskiy
- 1325 Landmark Based Facial Component Reconstruction for Recognition Across Pose, Gee-Sern Hsu, Hsiao-Chia Peng, Kai-Hsiang Chang

### **Deep Learning for Computer Vision**

Organizer: Graham Taylor

Marc'Aurelio Ranzato

Honglak Lee

**Time:** 0830-1700 (Full Day) **Location:** Grand Ballroom 2

Description: A central challenge in visual reasoning is that of untangling the many factors of variation that explain an image or video. Photometric and geometric "nuisance" factors are intertwined with the variables of interest, for example, object identity in recognition tasks. To date, the dominant methodology for addressing this challenge has been to engineer a feature extraction pipeline, usually containing multiple stages of processing. An alternative approach is "Representation Learning": relying on the data, instead of feature engineering to learn representations that are invariant to nuisance factors. Techniques that learn multiple layers of representation, which are referred to as "Deep Learning", have demonstrated not only impressive success in recent benchmarks and competitions but applicability to multiple domains. The tutorial will be structured in two parts. In the morning, we will review the foundations of deep learning applied to vision in both the supervised and unsupervised setting. We will also highlight the most frequently used practical development libraries and tools. In the afternoon, we will invite leading experts in the field to discuss the most relevant application areas, including object detection, structured prediction, large-scale classification and hardware acceleration, video, multi-modal and multi-task learning, and regression methods for localization

# Dense Image Correspondences for Computer Vision

Organizer: Michael Rubinstein

Jaechul Kim Zhuowen Tu Ce Liu

**Time:** 0830-1700 (Full Day)

Location: C210

Description: Correspondence, namely how pixels in one image correspond to pixels in another image, is a fundamental problem in computer vision. Although correspondence has been mostly used for analyzing transformations between images from one scene, a new era has started recently when correspondence can be established across different scenes. In this tutorial, we will give an overview of dense correspondence algorithms for aligning images from different scenes. We will survey a variety of representations, including pixels (SIFT flow, Non-Rigid Dense Correspondence), semantic segments (layer flow) and image pyramid (deformable spatial pyramid). These dense alignment algorithms are powerful tools to analyze images and videos. We can not only transform information such as semantic labels, image details and geometry from images and videos in a labeled dataset, but also analyze an entire image database as a whole via information propagation. Recent advances on scene parsing, 2D video to 3D, annotation propagation (image to text), object discovery, cosegmentation, image hallucination, and biomedical image analysis demonstrate that across-scene correspondence can be a fundamental building block for computer vision.

### **BASIS-14: BASes for Images & Surfaces**

Organizer: Alex Bronstein

Michael Bronstein Iasonas Kokkinos George Papandreou

**Time:** 0830-1700 (Full Day)

Location: C212

**Description:** BASIS-14 will be a full day tutorial covering the current state-of-the-art in linear and nonlinear image and surface analysis techniques. Starting with the fundamentals of linear image processing, we will see how the main notions of Fourier transforms can be understood in terms of a change of basis, and will explore the multifold ramifications of this intuition to nonlinear image processing (sparse coding, dictionary learning, exemplar representations, invariant transforms) and surface analysis (heat diffusion on surfaces, spectral decomposition of the Laplace-Beltrami operator, surface descriptors).

We will present applications to both classical problems, such as denoising and deblurring, but also cutting edge-computer vision problems involving image classification, object detection, shape retrieval, and surface registration.

### Open CV 3.0: Solving Problems

Organizer: Gary Bradski

Vadim Pisarevsky Vincent Rabaud Grace Vesom

**Time:** 0830-1700 (Full Day)

Location: C222

**Description:** The third major release of OpenCV is aimed at building solid ground for computer vision development. C++ and Python will be covered here. OpenCV 3.0 enables:

 More algorithms to be integrated (of which we will showcase the latest)

- More supported languages (Matlab, Ruby, Haskell)
- More optimizations (NEON, OpenVX)
- A new and modular way of participating to the core development

While showcasing the aforementioned features, we will focus building end-to-end vision pipelines through several application walk-throughs. Code, instructions, and mobile applications will be available online before the tutorial.

# **Emerging Topics in Human Activity Recognition**

Organizers: Michael Ryoo

Ivan Laptev Greg Mori Sangmin Oh

**Time:** 0830-1230 (Half Day — Morning)

Location: C211

**Description:** In the past 5 years, the field of human activity recognition has grown dramatically, reflecting its importance in many high-impact societal applications including smart surveillance, web-video search and retrieval, quality-of-life devices for elderly people, and human-computer interfaces. Given the initial success of bag-of-words methods for action classification, the field is gradually moving towards more structured interpretation of complex human activities involving multiple people and objects as well as interactions among them in various realistic scenarios. New important research topics and problems are appearing as a consequence, including (i) modeling temporal structure of activities, (ii) learning relations between actions and objects/scenes/social roles, (iii) group activity recognition, and (iv) first-person activity recognition. The objective of this tutorial is to introduce and overview recent progress in these emerging topics, as well as to discuss, motivate and encourage future research in diverse subfields of activity recognition.

### State of the Art 3D Reconstruction Techniques: Very Large Scale 3D Reconstruction and the Role of Priors

Organizers: Noah Snavely

Yasutaka Furukawa

**Time:** 0830-1230 (Half Day — Morning)

**Location:** C123-125

**Description:** This course will cover state-of-the-art 3D reconstruction techniques beyond standard SfM and MVS techniques, focusing on two key areas. The first focus is in the large scale 3D reconstruction. As the core Structure from Motion (SfM) and Multi-View Stereo (MVS) technologies become mature and robust, more and more interests and demands arose for very large-scale SfM and MVS executions, primarily for digital mapping applications. The second focus is in the use of structure priors in 3D reconstruction, such as planarity, orthogonality, symmetry, and repetition, which pose challenges to standard SfM and MVS techniques, but can yield rich structure information about the scene for better 3D modeling.

### Learning Visual Semantics: Models, Massive Computation, & Innovative Applications

Organizer: Shih-Fu Chang

John Smith Rogerio Feris Liangliang Cao

Time: 1300-1700 (Half Day — Afternoon)

Location: C211

Description: The explosion of digital multimedia data including visual content from surveillance cameras, mobile phones, personal photo collections, news footage, or medical images - is creating significant opportunities for automated visual analysis. However, the most interesting content in multimedia files is often unconstrained and complex in nature, reflecting a diversity of human behaviors, scenes, activities, and events, which poses serious challenges for computer vision approaches. In this tutorial, we will present the state-of-the-art on large-scale visual semantic modeling, covering methods for obtaining intuitive mid-level semantic feature representations, while presenting innovative applications. The organizers will share their experience in achieving top performance on several recent competitions, including TRECVID, ImageNet, and ImageCLEF, and developing large-scale data and tool resources.

### **Video Segmentation**

Organizer: Jason Corso

Matthias Grundman

Irfan Essa

**Time:** 1300-1700 (Half Day — Afternoon)

Location: C216

Description: In recent years, segmentation has emerged as a plausible first step in early video processing of unconstrained videos, without needing to make an assumption of a static background as earlier methods have. Video segmentation and over-segmentation, or more commonly supervoxel extraction, is a complementary early video processing step to the more traditional feature extraction, such as STIP and trajectories, and it extends the long history of image segmentation methods. This tutorial will survey and present the important models and algorithms for video segmentation. We will cover direct extensions of image segmentation methods through video-specific spatiotemporal and streaming methods. In addition to core methodological elements, the tutorial will also cover benchmark and evaluation of video segmentation as well as applications of video segmentation. Participants will be introduced to the details of these methods not only through traditional slide presentations but also example implementations through the LIBSVX library.

# Large-Scale Visual Place Recognition and Image-Based Localization

Organizer: Torsten Sattler

Akihiko Torii

Time: 1300-1700 (Half Day — Afternoon)

**Location:** C123-125

Description: The tutorial consists of two parts covering the general problems of visual place recognition and image-based localization. The first part is about visual place recognition and considers an application scenario in which the scene is represented by a set of geo-tagged images. The aim of visual place recognition is to approximate the position of the viewer by identifying the place visible in the guery image using (image) retrieval methods. We discuss several improvements to the standard retrieval pipeline that detect and remove confusing features, exploit the known spatial relations between the images, incorporate priors on the viewer's position, and enable place recognition systems to handle the repetitive structures prevalent in urban environments. The second part of the tutorial is about image-based localization and considers the more specific task of precisely estimating the pose of the query image relative to a 3D model of the scene. Assuming that this 3D model was reconstructed using Structure-from-Motion, we can find correspondences between 2D features in the query image and 3D points in the model using descriptor matching. We first introduce the standard data structures for descriptor matching as well as different approaches to estimate the camera pose from the 2D-3D matches. We then detail the prioritized matching schemes that enable state-of-the-art localization systems to efficiently handle 3D models consisting of millions of 3D points. We also discuss how to exploit existing visibility information between 3D points in the model and the database images and how to reduce the memory requirements by using only a subset of all 3D points without loss of localization performance. Throughout the tutorial, we provide links to publicly available source code for the discussed approaches as well as publicly available datasets.

## Tuesday, June 24 (Morning)

### **Program**

## Tuesday, June 24

**0700–1700 Registration** (Exhibit Hall C Lobby)

0730-0830 Breakfast (Exhibit Hall C)

**0820-0830 Welcome by the General Chairs** (Battelle Grand)

**0830-1200** AM Video Spotlights (C213-215)

## **0830–1000 Oral 1A: Matching & Reconstruction** (Battelle Grand South)

Poster IDs for this session: O-1A-# where # is the paper #.

Chairs: Jana Kosecka (George Mason Univ.) Antonis Argyros (Univ. of Crete)

Format (13 min. for presentation + 2 min. for questions)

- Fast and Accurate Image Matching with Cascade Hashing for 3D Reconstruction, Jian Cheng, Cong Leng, Jiaxiang Wu, Hainan Cui, Hanqing Lu
- 2. Predicting Matchability, Wilfried Hartmann, Michal Havlena, Konrad Schindler
- ${\scriptstyle 3.}\ \ {\rm Trinocular\ Geometry\ Revisited}, \textit{\it Jean\ Ponce,\ Martial\ Hebert}$
- Critical Configurations For Radial Distortion Self-Calibration, Changchang Wu
- Solvers for Relative Pose with a Single Unknown Radial Distortion Minimal, Yubin Kuang, Jan Erik Solem, Fredrik Kahl, Kalle Åström
- Reconstructing PASCAL VOC, Sara Vicente, João Carreira, Lourdes Agapito, Jorge Batista

## **0830–1000 Oral 1B: Segmentation & Grouping**(Battelle Grand North)

Poster IDs for this session: O-1B-# where # is the paper #.

Chairs: Piotr Dollar (Microsoft Research)
Tal Arbel (McGill Univ.)

Format (13 min. for presentation + 2 min. for questions)

- Spectral Graph Reduction for Efficient Image and
   Streaming Video Segmentation, Fabio Galasso, Margret
   Keuper, Thomas Brox, Bernt Schiele
- Weakly Supervised Multiclass Video Segmentation, Xiao Liu, Dacheng Tao, Mingli Song, Ying Ruan, Chun Chen, Jiajun Bu
- Video Motion Segmentation Using New Adaptive Manifold Denoising Model, Dijun Luo, Heng Huang
- Cut, Glue & Cut: A Fast, Approximate Solver for Multicut Partitioning, Thorsten Beier, Thorben Kroeger, Jörg H. Kappes, Ullrich Köthe, Fred A. Hamprecht
- Neural Decision Forests for Semantic Image Labelling, Samuel Rota Bulò, Peter Kontschieder
- 6. Pulling Things out of Perspective, *Lubor Ladický*, *Jianbo Shi, Marc Pollefeys*

### 1000-1030 Break (Grand Ballroom Prefunction)

### 1000-1200 Exhibits (Grand Ballrooms 1-3)

- Microsoft
- Google
- Xerox
- Amazon
- NVIDIA
- A9
- Face++
- Metaio
- Intel
- Curalate
- IBM
- Orbeus
- OMRON
- Qualcomm
- Itseez, Inc

- CoaniVue
- Elsevier
- MathWorks
- Point Grey
- now publishers
- CRC Press/Taylor & Francis
- Morgan & Claypool
- Springer
- KAUST
- Apple, Inc
- 3dMD
- Occam Vision Group
- Eyeris
- Spotscale
- Samsung MPI Lab

### 1000-1200 Demos (C110-115)

- Turning Mobile Phones into 3D Scanners, Petri Tanskanen, Kalin Kolev, Amael Delaunoy, Marc Pollefeys (ETH Zurich)
- Virtual Makeup, Sifei Liu, Jimei Yang, Zhe Hu (UC Merced)
- Learning to be a Depth Camera, Sean Ryan Fanello, Cem Keskin, Shahram Izadi, Pushmeet Kohli, David Kim, David Sweeney, Antonio Criminisi, Jamie Shotton, Sing Bing Kang, Tim Paek (Microsoft Research)
- Filter Forest for Learning Data-Dependent Filters, Sean Ryan Fanello, Cem Keskin, Pushmeet Kohli, Shahram Izadi, Jamie Shotton, Antonio Criminisi, Tim Paek (Microsoft Research)

### 1000-1200 Poster 1A: Recognition, Segmentation, Stereo & SFM (Grand Ballrooms 1-3)

### Poster IDs for this session: P-1A-# where # is the paper #.

- Event Detection using Multi-Level Relevance Labels and Multiple Features, Zhongwen Xu, Ivor W. Tsang, Yi Yang, Zhigang Ma, Alexander G. Hauptmann
- Full-Angle Quaternions for Robustly Matching Vectors of 3D Rotations, Stephan Liwicki, Minh-Tri Pham, Stefanos Zafeiriov, Maja Pantic, Björn Stenger
- Human vs. Computer in Scene and Object Recognition, Ali Borji, Laurent Itti
- 4. Semi-supervised Spectral Clustering for Image Set Classification, *Arif Mahmood, Ajmal Mian, Robyn Owens*
- Look at the Driver, Look at the Road: No Distraction! No Accident!, Mahdi Rezaei, Reinhard Klette
- Measuring Distance Between Unordered Sets of Different Sizes, Andrew Gardner, Jinko Kanno, Christian A. Duncan, Rastko Selmic
- Learning Mid-level Filters for Person Re-identification, Rui Zhao, Wanli Ouyang, Xiaogang Wang
- DeepReID: Deep Filter Pairing Neural Network for Person Re-Identification, Wei Li, Rui Zhao, Tong Xiao, Xiaogang Wang
- Lacunarity Analysis on Image Patterns for Texture Classification, Yuhui Quan, Yong Xu, Yuping Sun, Yu Luo
- 10. Segmentation-aware Deformable Part Models, Eduard Trulls, Stavros Tsogkas, Iasonas Kokkinos, Alberto Sanfeliu, Francesc Moreno-Noquer

- 11. From Categories to Individuals in Real Time A Unified Boosting Approach, David Hall, Pietro Perona
- 12. NMF-KNN: Image Annotation using Weighted Multi-view Non-negative Matrix Factorization, Mahdi M. Kalayeh, Haroon Idrees, Mubarak Shah
- 13. Fine-Grained Visual Comparisons with Local Learning, Aron Yu, Kristen Grauman
- 14. Inferring Analogous Attributes, Chao-Yeh Chen, Kristen Grauman
- 15. Beyond Comparing Image Pairs: Setwise Active Learning for Relative Attributes, *Lucy Liang, Kristen Grauman*
- Visual Persuasion: Inferring Communicative Intents of Images, Jungseock Joo, Weixin Li, Francis F. Steen, Song-Chun Zhu
- 17. Histograms of Pattern Sets for Image Classification and Object Recognition, Winn Voravuthikunchai, Bruno Crémilleux, Frédéric Jurie
- Incorporating Scene Context and Object Layout into Appearance Modeling, Hamid Izadinia, Fereshteh Sadeghi, Ali Farhadi
- Co-Segmentation of Textured 3D Shapes with Sparse Annotations, Mehmet Ersin Yumer, Won Chun, Ameesh Makadia
- 20. How to Evaluate Foreground Maps?, Ran Margolin, Lihi Zelnik-Manor, Ayellet Tal
- 21. MILCut: A Sweeping Line Multiple Instance Learning Paradigm for Interactive Image Segmentation, Jiajun Wu, Yibiao Zhao, Jun-Yan Zhu, Siwei Luo, Zhuowen Tu
- 22. SCAMS: Simultaneous Clustering and Model Selection, Zhuwen Li, Loong-Fah Cheong, Steven Zhiying Zhou
- 23. The Shape-Time Random Field for Semantic Video Labeling, Andrew Kae, Benjamin Marlin, Erik Learned-Miller
- 24. The Secrets of Salient Object Segmentation, Yin Li, Xiaodi Hou, Christof Koch, James M. Rehg, Alan L. Yuille
- 25. Non-rigid Segmentation using Sparse Low Dimensional Manifolds and Deep Belief Networks, Jacinto C. Nascimento, Gustavo Carneiro
- 26. An Exemplar-based CRF for Multi-instance Object Segmentation, *Xuming He, Stephen Gould*
- 27. Object Partitioning using Local Convexity, Simon Christoph Stein, Markus Schoeler, Jeremie Papon, Florentin Wörgötter

- 28. Bayesian Active Contours with Affine-Invariant, Elastic Shape Prior, *Darshan Bryner, Anuj Srivastava*
- 29. Max-Margin Boltzmann Machines for Object Segmentation, Jimei Yang, Simon Safar, Ming-Hsuan Yang
- 30. Multiscale Combinatorial Grouping, Pablo Arbeláez, Jordi Pont-Tuset, Jonathan T. Barron, Ferran Marques, Jitendra Malik
- 31. RIGOR: Reusing Inference in Graph Cuts for Generating Object Regions, Ahmad Humayun, Fuxin Li, James M. Rehg
- 32. Efficient Hierarchical Graph-Based Segmentation of RGBD Videos, Steven Hickson, Stan Birchfield, Irfan Essa, Henrik Christensen
- 33. Point Matching in the Presence of Outliers in Both Point Sets: A Concave Optimization Approach, Wei Lian, Lei Zhang
- 34. Multiple Structured-Instance Learning for Semantic Segmentation with Uncertain Training Data, Feng-Ju Chang, Yen-Yu Lin, Kuang-Jui Hsu
- Joint Motion Segmentation and Background Estimation in Dynamic Scenes, Adeel Mumtaz, Weichen Zhang, Antoni B. Chan
- 36. SeamSeg: Video Object Segmentation using Patch Seams, S. Avinash Ramakanth, R. Venkatesh Babu
- 37. Laplacian Coordinates for Seeded Image Segmentation, Wallace Casaca, Luis Gustavo Nonato, Gabriel Taubin
- 38. Error-tolerant Scribbles Based Interactive Image Segmentation, *Junjie Bai, Xiaodong Wu*
- 39. Iterative Multilevel MRF Leveraging Context and Voxel Information for Brain Tumour Segmentation in MRI, Nagesh Subbanna, Doina Precup, Tal Arbel
- 40. Large Scale Multi-view Stereopsis Evaluation, Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, Henrik Aanæs
- 41. Timing-Based Local Descriptor for Dynamic Surfaces, *Tony Tung, Takashi Matsuyama*
- 42. A Minimal Solution to the Generalized Pose-and-Scale Problem, Jonathan Ventura, Clemens Arth, Gerhard Reitmayr, Dieter Schmalstieg
- 43. A General and Simple Method for Camera Pose and Focal Length Determination, Yinqiang Zheng, Shigeki Sugimoto, Imari Sato, Masatoshi Okutomi

- 44. Partial Symmetry in Polynomial Systems and its Applications in Computer Vision, Yubin Kuang, Yinqiang Zhenq, Kalle Åström
- 45. Efficient Computation of Relative Pose for Multi-Camera Systems, Laurent Kneip, Hongdong Li
- 46. Simultaneous Localization and Calibration: Self-Calibration of Consumer Depth Cameras, Qian-Yi Zhou, Vladlen Koltun
- Minimal Scene Descriptions from Structure from Motion Models, Song Cao, Noah Snavely
- 48. Fast, Approximate Piecewise-Planar Modeling Based on Sparse Structure-from-Motion and Superpixels, András Bódis-Szomorú, Hayko Riemenschneider, Luc Van Gool
- 49. On Projective Reconstruction In Arbitrary Dimensions, Behrooz Nasihatkon, Richard Hartley, Jochen Trumpf
- 50. Stereo under Sequential Optimal Sampling: A Statistical Analysis Framework for Search Space Reduction, Yilin Wang, Ke Wang, Enrique Dunn, Jan-Michael Frahm
- 51. Efficient Pruning LMI Conditions for Branch-and-Prune Rank and Chirality-Constrained Estimation of the Dual Absolute Quadric, Adlane Habed, Danda Pani Paudel, Cédric Demonceaux, David Fofi
- 52. Very Fast Solution to the PnP Problem with Algebraic Outlier Rejection, Luis Ferraz, Xavier Binefa, Francesc Moreno-Noguer
- 53. Finding Vanishing Points via Point Alignments in Image Primal and Dual Domains, José Lezama, Rafael Grompone von Gioi, Gregory Randall, Jean-Michel Morel
- 54. Discriminative Feature-to-Point Matching in Image-Based Localization, *Michael Donoser, Dieter Schmalstieg*
- 55. Two-View Camera Housing Parameters Calibration for Multi-Layer Flat Refractive Interface, Xida Chen, Yee-Hong Yang
- 56. Accurate Localization and Pose Estimation for Large 3D Models, Linus Svärm, Olof Enqvist, Magnus Oskarsson, Fredrik Kahl
- 57. Relative Pose Estimation for a Multi-Camera System with Known Vertical Direction, Gim Hee Lee, Marc Pollefeys, Friedrich Fraundorfer

### 1200-1330 Lunch (Exhibit Hall C)

## **1200–1330 Doctoral Consortium** (C210-212) (by invitation only)

Supported by:



## amazon.com

- Guang Chen (Univ. of Missouri)
- Yu Chen (Northwestern Univ.)
- Bao Chenglong (National Univ. of Singapore)
- R. Gokberk Cinbis (INRIA Rhone-Alpes & Univ. of Grenoble)
- Jifeng Dai (Tsinghua Univ.)
- Jian Dong (National Univ. of Singapore)
- Kun Duan (Indiana Univ.)
- Alexander Fix (Cornell Univ.)
- Victor M. Fragoso (Univ. of California, Santa Barbara)
- Efstratios Gavves (Univ. of Amsterdam)
- Yunchao Gong (Univ. of North Carolina)
- Abner Guzmán-Rivera (Univ. of Illinois, Urbana-Champaign)
- Han Hu (Tsinghua Univ.)
- Zhe Hu (Univ. of California, Merced)
- Zhiwu Huang (Chinese Academy of Sciences)
- Satoshi Ikehata (Univ. of Tokyo)
- Catalin Ionescu (Univ. of Bonn)
- Sadeep Jayasumana (Australian National Univ.)
- Andrew Kae (Univ. of Massachusetts Amherst)
- Le Kang (Univ. of Maryland)
- Vahid Kazemi (KTH)
- Gim Hee Lee (ETH Zürich)
- Wen Li (Nanyang Technological Univ.)
- Stephan Liwicki (Imperial College London)

- Cewu Lu (Hong Konk Univ. of Science & Technology)
- Ping Luo (The Chinese Univ. of Hong Kong)
- Mohammad Norouzi (Univ. of Toronto)
- Iason Oikonomidis (Univ. of Crete)
- Matt O'Toole (Univ. of Toronto)
- Jeremie Papon (Georg-August-Universität Göttingen)
- Hyun Soo Park (Carnegie Mellon Univ.)
- Bryan Poling (Univ. of Minnesota)
- Vittal Premachandran (National Univ. of Singapore)
- Liu Shuaicheng (National Univ. of Singapore)
- Eran Swears (Rensselaer Polytechnic Inst.)
- Danhang Tang (Imperial College London)
- Eduard Trulls (CSIC-UPC)
- Xiaoyang Wang (Rensselaer Polytechnic Inst.)
- Zhaowen Wang (Univ. of Illinois, Urbana-Champaign)
- Chih-Yuan Yang (Univ. of California, Merced)
- Jimei Yang (Univ. of California, Merced)
- Jinwei Ye (Univ. of Delaware)
- Quanshi Zhang (Univ. of Tokyo)
- Yingying Zhu (Univ. of Queensland & CSIRO)
- Li Zhuwen (National Univ. of Singapore)

## Tuesday, June 24 (Afternoon)

### **Program**

1330-1830 PM Video Spotlights (C213-215)

## 1330-1445 Oral 1C: Statistical Methods & Learning I (Battelle Grand South)

Poster IDs for this session: O-1C-# where # is the paper #.

**Chairs :** Raquel Urtasun (Univ. of Toronto) Andrea Vedaldi (Univ. of Oxford)

Format (13 min. for presentation + 2 min. for questions)

- Optimal Decisions from Probabilistic Models: The Intersection-over-Union Case, Sebastian Nowozin
- 2. Covariance Trees for 2D and 3D Processing, *Thierry Guillemot, Andrés Almansa, Tamy Boubekeur*
- Hierarchical Subquery Evaluation for Active Learning on a Graph, Oisin Mac Aodha, Neill D.F. Campbell, Jan Kautz, Gabriel J. Brostow
- Anytime Recognition of Objects and Scenes, Sergey Karayev, Mario Fritz, Trevor Darrell
- Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation, Ross Girshick, Jeff Donahue, Trevor Darrell, Jitendra Malik

## **1330–1445 Oral 1D: Action Recognition** (Battelle Grand North)

Poster IDs for this session: O-1D-# where # is the paper #.

Chairs: Sinisa Todorovic (Oregon State Univ.) Sudeep Sarkar (Univ. of South Florida)

Format (13 min. for presentation + 2 min. for questions)

- Human Action Recognition by Representing 3D Skeletons as Points in a Lie Group, Raviteja Vemulapalli, Felipe Arrate, Rama Chellappa
- 2. Multi-View Super Vector for Action Recognition, *Zhuowei* Cai, Limin Wang, Xiaojiang Peng, Yu Qiao
- Unsupervised Spectral Dual Assignment Clustering of Human Actions in Context, Simon Jones, Ling Shao
- 4. Parsing Videos of Actions with Segmental Grammars, Hamed Pirsiavash, Deva Ramanan
- Rate-Invariant Analysis of Trajectories on Riemannian Manifolds with Application in Visual Speech Recognition, Jingyong Su, Anuj Srivastava, Fillipe D. M. de Souza, Sudeep Sarkar

### 1445-1515 Break (Battelle Grand Prefunction)

## 1515-1630 Special 1: Awards & Plenary Session (Battelle Grand)

Chairs: Aleix Martinez (Ohio State Univ.)
Cornelia Fermüller (Univ. of Maryland)

- Awards Ceremony: Program Chairs
- Plenary Talk: Neural Mechanisms for Face Processing, Doris Tsao (California Inst. of Technology)

Abstract: How the brain distills a representation of meaningful objects from retinal input is one of the central challenges of systems neuroscience. Functional imaging experiments in the macaque reveal that one ecologically important class of objects, faces, is represented by a system of six discrete, strongly interconnected regions. Electrophysiological recordings show that these 'face patches' have unique functional profiles. By studying the distinct visual representations maintained in these six face patches, the sequence of information flow between them, and the role each plays in face perception, we are gaining new insights into hierarchical information processing in the brain.

### 1630-1830 Exhibits (Grand Ballrooms 1-3)

• Same as Tuesday morning Exhibits (see pg. 15)

### **1630–1830 Demos** (C110-115)

- Object Partioning using Local Convexity, Markus Schoeler, Jeremie Papon (Univ. of Göttingen)
- Analysis by Synthesis: 3D Object Recognition by Object Reconstruction, Mohsen Hejrati, Deva Ramanan (UC Irvine)
- Estimating Image Depth Using Shape Collections, Hao Su, Qixing Huang, Niloy Mitra, Yangyan Li, Leonidas Guibas (Stanford, Univ. College London)
- Visipedia Backend: Collaborative Tools for Image Dataset Creation and Management, Grant Van Horn, Steve Branson, Catherine Wah, Pietro Perona, Serge Belongie (UC San Diego, Caltech, Cornell Tech)

### 1630-1830 Poster 1B: 3D Vision, Action Recognition, Recognition, Statistical Methods & Learning (Grand Ballrooms 1-3)

### Poster IDs for this session: P-1B-# where # is the paper #.

- Piecewise Planar and Compact Floorplan Reconstruction from Images, Ricardo Cabral, Yasutaka Furukawa
- 2. Data-driven Flower Petal Modeling with Botany Priors, Chenxi Zhang, Mao Ye, Bo Fu, Ruigang Yang
- User-Specific Hand Modeling from Monocular Depth Sequences, Jonathan Taylor, Richard Stebbing, Varun Ramakrishna, Cem Keskin, Jamie Shotton, Shahram Izadi, Aaron Hertzmann, Andrew Fitzgibbon
- Class Specific 3D Object Shape Priors Using Surface Normals, Christian Häne, Nikolay Savinov, Marc Pollefeys
- Frequency-Based 3D Reconstruction of Transparent and Specular Objects, Ding Liu, Xida Chen, Yee-Hong Yang
- Human Body Shape Estimation Using a Multi-Resolution Manifold Forest, Frank Perbet, Sam Johnson, Minh-Tri Pham, Björn Stenger
- Quality Dynamic Human Body Modeling Using a Single Low-cost Depth Camera, Qing Zhang, Bo Fu, Mao Ye, Ruigang Yang
- 8. Single-View 3D Scene Parsing by Attributed Grammar, Xiaobai Liu, Yibiao Zhao, Song-Chun Zhu
- Separation of Line Drawings Based on Split Faces for 3D Object Reconstruction, Changqing Zou, Heng Yang, Jianzhuang Liu
- 10. When 3D Reconstruction Meets Ubiquitous RGB-D Images, Quanshi Zhang, Xuan Song, Xiaowei Shao, Huijing Zhao, Ryosuke Shibasaki
- 11. Stable Template-Based Isometric 3D Reconstruction in All Imaging Conditions by Linear Least-Squares, Ajad Chhatkuli, Daniel Pizarro, Adrien Bartoli
- 12. Discrete-Continuous Depth Estimation from a Single Image, *Miaomiao Liu, Mathieu Salzmann, Xuming He*
- 13. Leveraging Hierarchical Parametric Networks for Skeletal Joints Based Action Segmentation and Recognition, Di Wu, Ling Shao
- 14. Seeing What You're Told: Sentence-Guided Activity Recognition In Video, Narayanaswamy Siddharth, Andrei Barbu, Jeffrey Mark Siskind

- Action Localization with Tubelets from Motion, Mihir Jain, Jan van Gemert, Hervé Jégou, Patrick Bouthemy, Cees G.M. Snoek
- Actionness Ranking with Lattice Conditional Ordinal Random Fields, Wei Chen, Caiming Xiong, Ran Xu, Jason J. Corso
- Multiple Granularity Analysis for Fine-grained Action Detection, Bingbing Ni, Vignesh R. Paramathayalan, Pierre Moulin
- 18. Human Action Recognition Across Datasets by Foreground-weighted Histogram Decomposition, Waqas Sultani, Imran Saleemi
- 19. Range-Sample Depth Feature for Action Recognition, Cewu Lu, Jiaya Jia, Chi-Keung Tang
- 20. The Language of Actions: Recovering the Syntax and Semantics of Goal-Directed Human Activities, Hilde Kuehne, Ali Arslan, Thomas Serre
- 21. Complex Activity Recognition using Granger Constrained DBN (GCDBN) in Sports and Surveillance Video, Eran Swears, Anthony Hoogs, Qiang Ji, Kim Boyer
- 22. Incremental Activity Modeling and Recognition in Streaming Videos, Mahmudul Hasan, Amit K. Roy-Chowdhury
- 23. Super Normal Vector for Activity Recognition Using Depth Sequences, Xiaodong Yang, YingLi Tian
- 24. Discriminative Hierarchical Modeling of Spatio-Temporally Composable Human Activities, Ivan Lillo, Alvaro Soto, Juan Carlos Niebles
- 25. A Multigraph Representation for Improved Unsupervised/Semi-supervised Learning of Human Actions, Simon Jones, Ling Shao
- 26. StoryGraphs: Visualizing Character Interactions as a Timeline, Makarand Tapaswi, Martin Bäuml, Rainer Stiefelhagen
- 27. Learning Receptive Fields for Pooling from Tensors of Feature Response, Can Xu, Nuno Vasconcelos
- 28. Towards Unified Human Parsing and Pose Estimation, Jian Dong, Qiang Chen, Xiaohui Shen, Jianchao Yang, Shuicheng Yan
- 29. Ask the Image: Supervised Pooling to Preserve Feature Locality, Sean Ryan Fanello, Nicoletta Noceti, Carlo Ciliberto, Giorgio Metta, Francesca Odone

### **Program**

- 30. Similarity Comparisons for Interactive Fine-Grained Categorization, Catherine Wah, Grant Van Horn, Steve Branson, Subhransu Maji, Pietro Perona, Serge Belongie
- 31. Continuous Manifold Based Adaptation for Evolving Visual Domains, Judy Hoffman, Trevor Darrell, Kate Saenko
- 32. Talking Heads: Detecting Humans and Recognizing Their Interactions, *Minh Hoai, Andrew Zisserman*
- 33. Salient Region Detection via High-Dimensional Color Transform, Jiwhan Kim, Dongyoon Han, Yu-Wing Tai, Junmo Kim
- 34. The Role of Context for Object Detection and Semantic Segmentation in the Wild, Roozbeh Mottaghi, Xianjie Chen, Xiaobai Liu, Nam-Gyu Cho, Seong-Whan Lee, Sanja Fidler, Raquel Urtasun, Alan Yuille
- 35. Switchable Deep Network for Pedestrian Detection, *Ping Luo, Yonglong Tian, Xiaogang Wang, Xiaoou Tang*
- 36. Compact Representation for Image Classification: To Choose or to Compress?, Yu Zhang, Jianxin Wu, Jianfei Cai
- 37. Capturing Long-tail Distributions of Object Subcategories, Xianqxin Zhu, Dragomir Anguelov, Deva Ramanan
- 38. Accurate Object Detection with Joint Classification-Regression Random Forests, Samuel Schulter, Christian Leistner, Paul Wohlhart, Peter M. Roth, Horst Bischof
- 39. Additive Quantization for Extreme Vector Compression, Artem Babenko, Victor Lempitsky
- 40. Product Sparse Coding, Tiezheng Ge, Kaiming He, Jian Sun
- 41. Informed Haar-like Features Improve Pedestrian
  Detection, Shanshan Zhang, Christian Bauckhage, Armin B.
  Cremers
- 42. Image Reconstruction from Bag-of-Visual-Words, *Hiroharu Kato, Tatsuya Harada*
- 43. Beta Process Multiple Kernel Learning, *Bingbing Ni, Teng Li, Pierre Moulin*
- 44. Random Laplace Feature Maps for Semigroup Kernels on Histograms, Jiyan Yang, Vikas Sindhwani, Quanfu Fan, Haim Avron, Michael W. Mahoney
- 45. Hash-SVM: Scalable Kernel Machines for Large-Scale Visual Classification, *Yadong Mu, Gang Hua, Wei Fan, Shih-Fu Chang*
- 46. Transitive Distance Clustering with K-Means Duality, Zhiding Yu, Chunjing Xu, Deyu Meng, Zhuo Hui, Fanyi Xiao, Wenbo Liu, Jianzhuana Liu

- 47. Simultaneous Twin Kernel Learning using Polynomial Transformations for Structured Prediction, Chetan Tonde, Ahmed Elgammal
- 48. Bregman Divergences for Infinite Dimensional Covariance Matrices, Mehrtash Harandi, Mathieu Salzmann, Fatih Porikli
- 49. Optimizing Average Precision using Weakly Supervised Data, Aseem Behl, C. V. Jawahar, M. Pawan Kumar
- 50. Subspace Clustering for Sequential Data, Stephen Tierney, Junbin Gao, Yi Guo
- 51. Predicting Multiple Attributes via Relative Multi-task Learning, *Lin Chen, Qiang Zhang, Baoxin Li*
- 52. Learning Inhomogeneous FRAME Models for Object Patterns, Jianwen Xie, Wenze Hu, Song-Chun Zhu, Ying Nian Wu
- 53. Empirical Minimum Bayes Risk Prediction: How to Extract an Extra Few % Performance from Vision Models with Just Three More Parameters, Vittal Premachandran, Daniel Tarlow, Dhruv Batra
- 54. Fantope Regularization in Metric Learning, Marc T. Law, Nicolas Thome, Matthieu Cord
- 55. Kernel-PCA Analysis of Surface Normals for Shape-from-Shading, *Patrick Snape*, *Stefanos Zafeiriou*
- 56. Merging SVMs with Linear Discriminant Analysis: A Combined Model, Symeon Nikitidis, Stefanos Zafeiriou, Maja Pantic
- 57. Stable Learning in Coding Space for Multi-Class Decoding and Its Extension for Multi-Class Hypothesis Transfer Learning, Bang Zhang, Yi Wang, Yang Wang, Fang Chen
- 58. Finding the Subspace Mean or Median to Fit Your Need, Tim Marrinan, J. Ross Beveridge, Bruce Draper, Michael Kirby, Chris Peterson

### **1830–2030 Reception** (Battelle Grand)

### Wednesday, June 25 (Morning)

### **Program**

## Wednesday, June 25

**0700–1700 Registration** (Exhibit Hall C Lobby)

0730-0830 Breakfast (Exhibit Hall C)

**0830-1200** AM Video Spotlights (C213-215)

## 0830-1000 Oral 2A: Motion & Tracking (Battelle Grand South)

Poster IDs for this session: O-2A-# where # is the paper #.

Chairs: Simon Lucey (CSIRO)
Ming-Hsuan Yang (UC Merced)

Format (13 min. for presentation + 2 min. for questions)

- Adaptive Color Attributes for Real-Time Visual Tracking, Martin Danelljan, Fahad Shahbaz Khan, Michael Felsberg, Joost van de Weijer
- Local Layering for Joint Motion Estimation and Occlusion Detection, Deging Sun, Ce Liu, Hanspeter Pfister
- 3. Realtime and Robust Hand Tracking from Depth, Chen Qian, Xiao Sun, Yichen Wei, Xiaoou Tang, Jian Sun
- Multi-Output Learning for Camera Relocalization, Abner Guzman-Rivera, Pushmeet Kohli, Ben Glocker, Jamie Shotton, Toby Sharp, Andrew Fitzgibbon, Shahram Izadi
- MAP Visibility Estimation for Large-Scale Dynamic 3D Reconstruction, Hanbyul Joo, Hyun Soo Park, Yaser Sheikh
- Multi-Object Tracking via Constrained Sequential Labeling, Sheng Chen, Alan Fern, Sinisa Todorovic

## **0830–1000 Oral 2B: Discrete Optimization** (Battelle Grand North)

Poster IDs for this session: O-2B-# where # is the paper #.

Chairs: Olga Veksler (Univ. of Western Ontario) Hiroshi Ishikawa (Waseda Univ.)

Format (13 min. for presentation + 2 min. for questions)

- A Primal-Dual Algorithm for Higher-Order Multilabel Markov Random Fields, Alexander Fix, Chen Wang, Ramin Zabih
- Energy Based Multi-model Fitting & Matching for 3D Reconstruction, Hossam Isack, Yuri Boykov
- Submodularization for Binary Pairwise Energies, Lena Gorelick, Yuri Boykov, Olga Veksler, Ismail Ben Ayed, Andrew Delong
- Maximum Persistency in Energy Minimization, Alexander Shekhovtsov
- Partial Optimality by Pruning for MAP-inference with General Graphical Models, Paul Swoboda, Bogdan Savchynskyy, Jörq H. Kappes, Christoph Schnörr
- 6. Scene Labeling Using Beam Search Under Mutex Constraints, *Anirban Roy*, *Sinisa Todorovic*

### 1000-1030 Break (Grand Ballroom Prefunction)

### 1000-1200 Exhibits (Grand Ballrooms 1-3)

Same as Tuesday morning Exhibits (see pg. 15)

### 1000-1200 Demos (C110-115)

- Automatic Façade Parsing and LOD3 Model Generation from 3D Point Clouds, William Nguatem, Martin Drauschke, Helmut Mayer (Bundeswehr Univ. Munich)
- Story-based Video Retrieval in TV series using Plot Synopses, Makarand Tapaswi, Martin Bäuml, Rainer Stiefelhagen (Karlsruhe Inst. of Technology)
- Authentication Using Sketches with Biometeric Information, Benjamin S. Riggan, Wesley E. Snyder, Cliff Wang (NC State Univ., US Army Research Office)
- Tracking Multiple Interacting Targets in a Camera Network, Shu Zhang, Amit K. Roy-Chowdhury (UC Riverside)

### 1000-1200 Poster 2A: Motion & Tracking, Optimization, Statistical Methods & Learning, Stereo & SFM (Grand Ballrooms 1-3)

### Poster IDs for this session: P-2A-# where # is the paper #.

- Persistent Tracking for Wide Area Aerial Surveillance, Jan Prokaj, Gérard Medioni
- Multi-Cue Visual Tracking Using Robust Feature-Level Fusion Based on Joint Sparse Representation, Xiangyuan Lan, Andy J. Ma, Pong C. Yuen
- 3. Multi-Forest Tracker: A Chameleon in Tracking, *David Joseph Tan, Slobodan Ilic*
- Rigid Motion Segmentation using Randomized Voting, Heechul Jung, Jeongwoo Ju, Junmo Kim
- Robust Online Multi-Object Tracking based on Tracklet Confidence and Online Discriminative Appearance Learning, Seung-Hwan Bae, Kuk-Jin Yoon
- Pyramid-based Visual Tracking Using Sparsity Represented Mean Transform, Zhe Zhang, Kin Hong Wong
- Tracklet Association with Online Target-Specific Metric Learning, Bing Wang, Gang Wang, Kap Luk Chan, Li Wang
- 8. An Online Learned Elementary Grouping Model for Multitarget Tracking, *Xiaojing Chen, Zhen Qin, Le An, Bir Bhanu*
- Diversity-Enhanced Condensation Algorithm and Its Application for Robust and Accurate Endoscope Three-Dimensional Motion Tracking, Xiongbiao Luo, Ying Wan, Xiangjian He, Jie Yang, Kensaku Mori
- 10. Partial Occlusion Handling for Visual Tracking via Robust Part Matching, Tianzhu Zhang, Kui Jia, Changsheng Xu, Yi Ma, Narendra Ahuja
- 11. Speeding Up Tracking by Ignoring Features, Lu Zhang, Hamdi Dibeklioğlu, Laurens van der Maaten
- 12. Subspace Tracking under Dynamic Dimensionality for Online Background Subtraction, Matthew Berger, Lee M. Seversky
- 13. Multiple Target Tracking Based on Undirected Hierarchical Relation Hypergraph, Longyin Wen, Wenbo Li, Junjie Yan, Zhen Lei, Dong Yi, Stan Z. Li
- 14. Bi-label Propagation for Generic Multiple Object Tracking, Wenhan Luo, Tae-Kyun Kim, Björn Stenger, Xiaowei Zhao, Roberto Cipolla

- 15. A Probabilistic Framework for Multitarget Tracking with Mutual Occlusions, Menglong Yang, Yiguang Liu, Longyin Wen, Zhisheng You, Stan Z. Li
- Occlusion Geodesics for Online Multi-Object Tracking, Horst Possegger, Thomas Mauthner, Peter M. Roth, Horst Bischof
- 17. Efficient Nonlinear Markov Models for Human Motion, Andreas M. Lehrmann, Peter V. Gehler, Sebastian Nowozin
- 18. A Compositional Model for Low-Dimensional Image Set Representation, Hossein Mobahi, Ce Liu, William T. Freeman
- 19. A Principled Approach for Coarse-to-Fine MAP Inference, Christopher Zach
- 20. Fast Approximate Inference in Higher Order MRF-MAP Labeling Problems, Chetan Arora, Subhashis Banerjee, Prem Kalra, S.N. Maheshwari
- 21. Scanline Sampler without Detailed Balance: An Efficient MCMC for MRF Optimization, Wonsik Kim, Kyoung Mu Lee
- 22. Higher-Order Clique Reduction Without Auxiliary Variables, *Hiroshi Ishikawa*
- 23. Topic Modeling of Multimodal Data: An Autoregressive Approach, Yin Zheng, Yu-Jin Zhang, Hugo Larochelle
- 24. Model Transport: Towards Scalable Transfer Learning on Manifolds, Oren Freifeld, Søren Hauberg, Michael J. Black
- 25. Learning Fine-grained Image Similarity with Deep Ranking, Jiang Wang, Yang Song, Thomas Leung, Chuck Rosenberg, Jingbin Wang, James Philbin, Bo Chen, Ying Wu
- 26. Attributed Graph Mining and Matching: An Attempt to Define and Extract Soft Attributed Patterns, Quanshi Zhang, Xuan Song, Xiaowei Shao, Huijing Zhao, Ryosuke Shibasaki
- 27. Deep Fisher Kernels End to End Learning of the Fisher Kernel GMM Parameters, Vladyslav Sydorov, Mayu Sakurada, Christoph H. Lampert
- 28. Transfer Joint Matching for Unsupervised Domain Adaptation, Mingsheng Long, Jianmin Wang, Guiguang Ding, Jiaguang Sun, Philip S. Yu
- 29. Recognizing RGB Images by Learning from RGB-D Data, Lin Chen, Wen Li, Dong Xu
- 30. Instance-weighted Transfer Learning of Active Appearance Models, Daniel Haase, Erik Rodner, Joachim Denzler

- 31. Scalable Multitask Representation Learning for Scene Classification, Maksim Lapin, Bernt Schiele, Matthias Hein
- 32. Learning to Learn, from Transfer Learning to Domain Adaptation: A Unifying Perspective, Novi Patricia, Barbara Caputo
- 33. Constructing Robust Affinity Graphs for Spectral Clustering, Xiatian Zhu, Chen Change Loy, Shaogang Gong
- 34. A Fast and Robust Algorithm to Count Topologically Persistent Holes in Noisy Clouds, *Vitaliy Kurlin*
- 35. Co-localization in Real-World Images, Kevin Tang, Armand Joulin, Li-Jia Li, Li Fei-Fei
- 36. Spectral Clustering with Jensen-type Kernels and their Multi-point Extensions, Debarghya Ghoshdastidar, Ambedkar Dukkipati, Ajay P. Adsul, Aparna S. Vijayan
- 37. Fast and Robust Archetypal Analysis for Representation Learning, Yuansi Chen, Julien Mairal, Zaid Harchaoui
- 38. Photometric Bundle Adjustment for Dense Multi-View 3D Modeling, Amaël Delaunoy, Marc Pollefeys
- 39. The Photometry of Intrinsic Images, Marc Serra, Olivier Penacchio, Robert Benavente, Maria Vanrell, Dimitris Samaras
- 40. High Resolution 3D Shape Texture from Multiple Videos, Vagia Tsiminaki, Jean-Sébastien Franco, Edmond Boyer
- 41. PatchMatch Based Joint View Selection and Depthmap Estimation, Enliang Zheng, Enrique Dunn, Vladimir Jojic, Jan-Michael Frahm
- 42. Light Field Stereo Matching Using Bilateral Statistics of Surface Cameras, Can Chen, Haiting Lin, Zhan Yu, Sing Bing Kang, Jingyi Yu
- 43. Recovering Surface Details under General Unknown Illumination Using Shading and Coarse Multi-view Stereo, Di Xu, Qi Duan, Jianming Zheng, Juyong Zhang, Jianfei Cai, Tat-Jen Cham
- 44. Probabilistic Labeling Cost for High-Accuracy Multi-View Reconstruction, *Ilya Kostrikov*, *Esther Horbert, Bastian Leibe*
- 45. Complex Non-Rigid Motion 3D Reconstruction by Union of Subspaces, Yingying Zhu, Dong Huang, Fernando De La Torre, Simon Lucey
- 46. A Procrustean Markov Process for Non-Rigid Structure Recovery, *Minsik Lee, Chong-Ho Choi, Songhwai Oh*

- 47. Good Vibrations: A Modal Analysis Approach for Sequential Non-Rigid Structure from Motion, *Antonio* Agudo, Lourdes Agapito, Begoña Calvo, Jose M. M. Montiel
- 48. Robust Scale Estimation in Real-Time Monocular SFM for Autonomous Driving, Shiyu Song, Manmohan Chandraker
- 49. On the Quotient Representation for the Essential Manifold, *Roberto Tron, Kostas Daniilidis*
- 50. Efficient High-Resolution Stereo Matching using Local Plane Sweeps, Sudipta N. Sinha, Daniel Scharstein, Richard Szeliski
- 51. Cross-Scale Cost Aggregation for Stereo Matching, Kang Zhang, Yuqiang Fang, Dongbo Min, Lifeng Sun, Shiqiang Yang, Shuicheng Yan, Qi Tian
- 52. Asymmetrical Gauss Mixture Models for Point Sets Matching, *Wenbing Tao, Kun Sun*
- 53. Fast and Reliable Two-View Translation Estimation, Johan Fredriksson, Olof Enqvist, Fredrik Kahl
- 54. Graph Cut based Continuous Stereo Matching using Locally Shared Labels, Tatsunori Taniai, Yasuyuki Matsushita, Takeshi Naemura
- 55. Learning to Detect Ground Control Points for Improving the Accuracy of Stereo Matching, Aristotle Spyropoulos, Nikos Komodakis, Philippos Mordohai

### 1200-1330 Lunch (Exhibit Hall C)

### Wednesday, June 25 (Afternoon)

### **Program**

### 1330-1830 PM Video Spotlights (C213-215)

## 1330-1500 Special 2: PAMI/IJCV Special Journal Session (Battelle Grand South)

Chairs: Ramin Zabih (Cornell Univ.)
Ronen Basri (Weizmann Inst. of Science)

Format (13 min. for presentation + 2 min. for questions)

- Make<sub>3</sub>D: Learning <sub>3</sub>D Scene Structure from a Single Still Image, Ashutosh Saxena, Min Sun, Andrew Y. Na
- 2. Product Quantization for Nearest Neighbor Search , Hervé Jégou, Matthijs Douze, Cordelia Schmid
- The PASCAL Visual Object Classes (VOC) Challenge, Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John Winn, Andrew Zisserman
- 4. Convex and Semi-Nonnegative Matrix Factorizations , Chris HQ Ding, Tao Li, Michael I. Jordan
- Robust Face Recognition via Sparse Representation, John Wright, Allen Y. Yang, Arvind Ganesh, S. Shankar Sastry, Yi Ma
- Deep Learning with Hierarchical Convolutional Factor Analysis, Bo Chen, Gungor Polatkan, Guillermo Sapiro, David Blei, David Dunson, Lawrence Carin

### 1330-1500 Oral 2D: Attribute-Based Recognition & Human Pose Estimation (Battelle Grand North)

Poster IDs for this session: O-2D-# where # is the paper #.

Chairs: Christoph Lampert (IST Austria) Yi Li (NICTA)

Format (13 min. for presentation + 2 min. for questions)

- 1. Decorrelating Semantic Visual Attributes by Resisting the Urge to Share, *Dinesh Jayaraman*, *Fei Sha*, *Kristen Grauman*
- PANDA: Pose Aligned Networks for Deep Attribute Modeling, Ning Zhang, Manohar Paluri, Marc'Aurelio Ranzato, Trevor Darrell, Lubomir Bourdev
- Learning Scalable Discriminative Dictionary with Sample Relatedness, Jiashi Feng, Stefanie Jegelka, Shuicheng Yan, Trevor Darrell
- 4. DeepPose: Human Pose Estimation via Deep Neural Networks, *Alexander Toshev, Christian Szegedy*
- Iterated Second-Order Label Sensitive Pooling for 3D Human Pose Estimation, Catalin Ionescu, Joao Carreira, Cristian Sminchisescu

6. 3D Pictorial Structures for Multiple Human Pose Estimation, Vasileios Belagiannis, Sikandar Amin, Mykhaylo Andriluka, Bernt Schiele, Nassir Navab, Slobodan Ilic

### 1500-1530 Break (Battelle Grand Prefunction)

### 1530-1630 Oral 2E: Face & Gesture

(Battelle Grand South)

Poster IDs for this session: O-2E-# where # is the paper #.

Chairs: Fernando De la Torre (CMU) Rama Chellappa (Univ. of Maryland)

Format (13 min. for presentation + 2 min. for questions)

- Learning Euclidean-to-Riemannian Metric for Point-to-Set Classification, Zhiwu Huang, Ruiping Wang, Shiguang Shan, Xilin Chen
- Face Alignment at 3000 FPS via Regressing Local Binary Features, Shaoqing Ren, Xudong Cao, Yichen Wei, Jian Sun
- A Compact and Discriminative Face Track Descriptor, Omkar M. Parkhi, Karen Simonyan, Andrea Vedaldi, Andrew Zisserman
- DeepFace: Closing the Gap to Human-Level Performance in Face Verification, Yaniv Taigman, Ming Yang, Marc'Aurelio Ranzato, Lior Wolf

## **1530–1630 Oral 2F: Convolutional Neural Networks** (Battelle Grand North)

Poster IDs for this session: O-2F-# where # is the paper #.

Chairs: Zhuowen Tu (UC San Diego) Fatih Porikli (MERL)

Format (13 min. for presentation + 2 min. for questions)

- Filter Forests for Learning Data-Dependent Convolutional Kernels, Sean Ryan Fanello, Cem Keskin, Pushmeet Kohli, Shahram Izadi, Jamie Shotton, Antonio Criminisi, Ugo Pattacini, Tim Paek
- Learning and Transferring Mid-Level Image Representations using Convolutional Neural Networks, Maxime Oquab, Leon Bottou, Ivan Laptev, Josef Sivic
- Large-scale Video Classification with Convolutional Neural Networks, Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, Li Fei-Fei

## Wednesday, June 25 (Afternoon)

### **Program**

 Convolutional Neural Networks for No-Reference Image Quality Assessment, Le Kang, Peng Ye, Yi Li, David Doermann

### 1630-1830 Exhibits (Grand Ballrooms 1-3)

• Same as Tuesday morning Exhibits (see pg. 15)

### 1630-1830 Demos (C110-115)

- Real-time Face Detection and Recognition on Google Glass, Shue-Ching, Bappaditya Mandal, Vijay Chandrasekhar, Cheston Tan, Liyuan Li, Joo Hwee Lim (Inst. for Infocomm Research)
- Rapid and Accurate Avatar Capture using a Single Mounted Kinect, Jongmoo Choi, Gérard Medioni (USC)
- Photo Recall: Using the Internet to Label Your Photos, Neeraj Kumar, Steven Seitz (Univ. of Washington)
- The Chameleon Tracker in 3D, David Joseph Tan, Nassir Navab, Slobodan Ilic (Technical Univ. of Munich)

### 1630–1830 Poster 2B: Face & Gesture, Recognition (Grand Ballrooms 1-3)

### Poster IDs for this session: P-2B-# where # is the paper #.

- Nonparametric Context Modeling of Local Appearance for Pose- and Expression-Robust Facial Landmark Localization, Brandon M. Smith, Jonathan Brandt, Zhe Lin, Li Zhang
- Learning Expressionlets on Spatio-Temporal Manifold for Dynamic Facial Expression Recognition, Mengyi Liu, Shiguang Shan, Ruiping Wang, Xilin Chen
- Who Do I Look Like? Determining Parent-Offspring Resemblance via Gated Autoencoders, Afshin Dehghan, Enrique G. Ortiz, Ruben Villegas, Mubarak Shah
- Unified Face Analysis by Iterative Multi-Output Random Forests, Xiaowei Zhao, Tae-Kyun Kim, Wenhan Luo
- Geometric Generative Gaze Estimation (G<sup>3</sup>E) for Remote RGB-D Cameras, Kenneth Alberto Funes Mora, Jean-Marc Odobez
- 6. A Hierarchical Probabilistic Model for Facial Feature Detection, *Yue Wu, Ziheng Wang, Qiang Ji*
- RAPS: Robust and Efficient Automatic Construction of Person-Specific Deformable Models, Christos Sagonas, Yannis Panagakis, Stefanos Zafeiriou, Maja Pantic

- 8. Non-Parametric Bayesian Constrained Local Models, Pedro Martins, Rui Caseiro, Jorge Batista
- Facial Expression Recognition via a Boosted Deep Belief Network, Ping Liu, Shizhong Han, Zibo Meng, Yan Tong
- 10. Automatic Construction of Deformable Models In-The-Wild, Epameinondas Antonakos, Stefanos Zafeiriou
- Learning-by-Synthesis for Appearance-based 3D Gaze Estimation, Yusuke Sugano, Yasuyuki Matsushita, Yoichi Sato
- 12. Towards Multi-view and Partially-Occluded Face Alignment, Junliang Xing, Zhiheng Niu, Junshi Huang, Weiming Hu, Shuicheng Yan
- 13. Head Pose Estimation Based on Multivariate Label Distribution, Xin Geng, Yu Xia
- 14. Efficient Boosted Exemplar-based Face Detection, Haoxiang Li, Zhe Lin, Jonathan Brandt, Xiaohui Shen, Gang Hua
- Gauss-Newton Deformable Part Models for Face Alignment in-the-Wild, Georgios Tzimiropoulos, Maja Pantic
- 16. Incremental Face Alignment in the Wild, Akshay Asthana, Stefanos Zafeiriou, Shiyang Cheng, Maja Pantic
- 17. One Millisecond Face Alignment with an Ensemble of Regression Trees, Vahid Kazemi, Josephine Sullivan
- 18. Discriminative Deep Metric Learning for Face Verification in the Wild, *Junlin Hu*, *Jiwen Lu*, *Yap-Peng Tan*
- 19. Stacked Progressive Auto-Encoders (SPAE) for Face Recognition Across Poses, Meina Kan, Shiguang Shan, Hong Chang, Xilin Chen
- 20. Deep Learning Face Representation from Predicting 10,000 Classes, *Yi Sun, Xiaogang Wang, Xiaoou Tang*
- 21. Occlusion Coherence: Localizing Occluded Faces with a Hierarchical Deformable Part Model, Golnaz Ghiasi, Charless C. Fowlkes
- 22. 3D-aided Face Recognition Robust to Expression and Pose Variations, Baptiste Chu, Sami Romdhani, Liming Chen
- 23. Learning Non-Linear Reconstruction Models for Image Set Classification, Munawar Hayat, Mohammed Bennamoun, Senjian An
- 24. Gesture Recognition Portfolios for Personalization, Angela Yao, Luc Van Gool, Pushmeet Kohli

### **Program**

- 25. Sign Spotting using Hierarchical Sequential Patterns with Temporal Intervals, Eng-Jon Ong, Oscar Koller, Nicolas Puqeault, Richard Bowden
- 26. Automatic Feature Learning for Robust Shadow Detection, Salman Hameed Khan, Mohammed Bennamoun, Ferdous Sohel, Roberto Togneri
- 27. Packing and Padding: Coupled Multi-index for Accurate Image Retrieval, Liang Zheng, Shengjin Wang, Ziqiong Liu, Oi Tian
- 28. Adaptive Object Retrieval with Kernel Reconstructive Hashing, *Haichuan Yang, Xiao Bai, Jun Zhou, Peng Ren, Zhihong Zhang, Jian Cheng*
- 29. Bayes Merging of Multiple Vocabularies for Scalable Image Retrieval, Liang Zheng, Shengjin Wang, Wengang Zhou, Qi Tian
- 30. Fast Supervised Hashing with Decision Trees for High-Dimensional Data, Guosheng Lin, Chunhua Shen, Qinfeng Shi, Anton van den Hengel, David Suter
- 31. Detect What You Can: Detecting and Representing Objects using Holistic Models and Body Parts, Xianjie Chen, Roozbeh Mottaghi, Xiaobai Liu, Sanja Fidler, Raquel Urtasun, Alan Yuille
- 32. Associative Embeddings for Large-scale Knowledge Transfer with Self-assessment, Alexander Vezhnevets, Vittorio Ferrari
- 33. Detecting Objects using Deformation Dictionaries, Bharath Hariharan, C. Lawrence Zitnick, Piotr Dollár
- 34. Persistence-based Structural Recognition, Chunyuan Li, Maks Ovsjanikov, Frederic Chazal
- 35. Inferring Unseen Views of People, Chao-Yeh Chen, Kristen Grauman
- 36. Birdsnap: Large-scale Fine-grained Visual Categorization of Birds, Thomas Berg, Jiongxin Liu, Seung Woo Lee, Michelle L. Alexander, David W. Jacobs, Peter N. Belhumeur
- 37. Predicting Object Dynamics in Scenes, *David F. Fouhey, C. Lawrence Zitnick*
- 38. Enriching Visual Knowledge Bases via Object Discovery and Segmentation, Xinlei Chen, Abhinav Shrivastava, Abhinav Gupta
- 39. Seeing the Arrow of Time, Lyndsey C. Pickup, Zheng Pan, Donglai Wei, YiChang Shih, Changshui Zhang, Andrew Zisserman, Bernhard Schölkopf, William T. Freeman

- 40. Hierarchical Feature Hashing for Fast Dimensionality Reduction, *Bin Zhao, Eric P. Xing*
- 41. Modeling Image Patches with a Generic Dictionary of Mini-Epitomes, George Papandreou, Liang-Chieh Chen, Alan L. Yuille
- 42. Simplex-Based 3D Spatio-Temporal Feature Description for Action Recognition, Hao Zhang, Wenjun Zhou, Christopher Reardon, Lynne E. Parker
- 43. In Search of Inliers: 3D Correspondence by Local and Global Voting, Anders Glent Buch, Yang Yang, Norbert Krüger, Henrik Gordon Petersen
- 44. Collective Matrix Factorization Hashing for Multimodal Data, *Guiguang Ding, Yuchen Guo, Jile Zhou*
- 45. Finding Matches in a Haystack: A Max-Pooling Strategy for Graph Matching in the Presence of Outliers, Minsu Cho, Jian Sun, Olivier Duchenne, Jean Ponce
- 46. Locality in Generic Instance Search from One Example, Ran Tao, Efstratios Gavves, Cees G.M. Snoek, Arnold W.M. Smeulders
- 47. Congruency-Based Reranking, Itai Ben-Shalom, Noga Levy, Lior Wolf, Nachum Dershowitz, Adiel Ben-Shalom, Roni Shweka, Yaacov Choueka, Tamir Hazan, Yaniv Bar
- 48. Asymmetric Sparse Kernel Approximations for Large-scale Visual Search, *Damek Davis, Jonathan Balzer, Stefano* Soatto
- 49. Locally Linear Hashing for Extracting Non-Linear Manifolds, Go Irie, Zhenguo Li, Xiao-Ming Wu, Shih-Fu Chang
- 50. Active Frame, Location, and Detector Selection for Automated and Manual Video Annotation, Vasiliy Karasev, Avinash Ravichandran, Stefano Soatto
- 51. Distance Encoded Product Quantization, Jae-Pil Heo, Zhe Lin, Sung-Eui Yoon
- 52. Collaborative Hashing, Xianglong Liu, Junfeng He, Cheng Deng, Bo Lang
- 53. Scalable Object Detection using Deep Neural Networks, Dumitru Erhan, Christian Szegedy, Alexander Toshev, Dragomir Anguelov

### 1830-2030 PAMI TC Meeting

(Battelle Grand South)

## Thursday, June 26 (Morning)

### **Program**

## Thursday, June 26

**0700–1700 Registration** (Exhibit Hall C Lobby)

0730-0830 Breakfast (Exhibit Hall C)

**0830-1200** AM Video Spotlights (C213-215)

## 0830-1000 Oral 3A: Physics-Based Vision & Shape-from-X (Battelle Grand South)

Poster IDs for this session: O-3A-# where # is the paper #.

Chairs: Robert Pless (Washington University)
Yoichi Sato (Univ. of Tokyo)

Format (13 min. for presentation + 2 min. for questions)

- Multiview Shape and Reflectance from Natural Illumination, Geoffrey Oxholm, Ko Nishino
- Reflectance and Fluorescent Spectra Recovery based on Fluorescent Chromaticity Invariance under Varying Illumination, Ying Fu, Antony Lam, Yasuyuki Kobashi, Imari Sato, Takahiro Okabe, Yoichi Sato
- 3. What Camera Motion Reveals About Shape With Unknown BRDF, Manmohan Chandraker
- Photometric Stereo using Constrained Bivariate Regression for General Isotropic Surfaces, Satoshi Ikehata, Kiyoharu Aizawa
- 5. Robust Separation of Reflection from Multiple Images, Xiaojie Guo, Xiaochun Cao, Yi Ma
- Surface-from-Gradients: An Approach Based on Discrete Geometry Processing, Wuyuan Xie, Yunbo Zhang, Charlie C. L. Wang, Ronald C.-K. Chung

## 0830-1000 Oral 3B: Video: Events, Activities & Surveillance (Battelle Grand North)

Poster IDs for this session: O-3B-# where # is the paper #.

Chairs: Tal Hassner (Open Univ. of Israel)
Rahul Sukthankar (Google Research)

Format (13 min. for presentation + 2 min. for questions)

- Socially-aware Large-scale Crowd Forecasting, Alexandre Alahi, Vignesh Ramanathan, Li Fei-Fei
- L<sub>o</sub> Regularized Stationary Time Estimation for Crowd Group Analysis, Shuai Yi, Xiaogang Wang, Cewu Lu, Jiaya Jia
- 3. Scene-Independent Group Profiling in Crowd, *Jing Shao*, Chen Change Loy, Xiaogang Wang
- Temporal Sequence Modeling for Video Event Detection, Yu Cheng, Quanfu Fan, Sharath Pankanti, Alok Choudhary
- Recognition of Complex Events: Exploiting Temporal Dynamics between Underlying Concepts, Subhabrata Bhattacharya, Mahdi M. Kalayeh, Rahul Sukthankar, Mubarak Shah
- Video Event Detection by Inferring Temporal Instance Labels, Kuan-Ting Lai, Felix X. Yu, Ming-Syan Chen, Shih-Fu Chang

1000-1030 Break (Grand Ballroom Prefunction)

### 1000-1200 Exhibits (Grand Ballrooms 1-3)

Same as Tuesday morning Exhibits (see pg. 15)

### 1000-1200 Demos (C110-115)

- Facial Analysis for BMI Estimation using a Mobile Device, Yu Zhu, Lingyun Wen, Guodong Guo (West Virginia Univ.)
- Real-Time Video Magnification, Neal Wadhwa, Michael Rubinstein, Frédo Durand, William T. Freeman (MIT CSAIL, Microsoft Research)
- Real Time Facial Expression Recognition on Android, Ankit Sharma, Oliver Nina, Lucas Pasqualin (Univ. of Central Florida)
- Ultra-Fast Attribute-Based Transfer Learning Using Images on the Internet, Daiki Kimura, Osamu Hasegawa, (Tokyo Inst. of Technology)

### 1000-1200 Poster 3A: Physics-Based Vision, Recognition, Video: Events, Activities & Surveillance (Grand Ballrooms 1-3)

### Poster IDs for this session: P-3A-# where # is the paper #.

- Backscatter Compensated Photometric Stereo with 3 Sources, Chourmouzios Tsiotsios, Maria E. Angelopoulou, Tae-Kyun Kim, Andrew J. Davison
- Calibrating a Non-isotropic Near Point Light Source using a Plane, Jaesik Park, Sudipta N. Sinha, Yasuyuki Matsushita, Yu-Wing Tai, In So Kweon
- A New Perspective on Material Classification and Ink Identification, Rakesh Shiradkar, Li Shen, George Landon, Sim Heng Ong, Ping Tan
- High Quality Photometric Reconstruction using a Depth Camera, Sk. Mohammadul Haque, Avishek Chatterjee, Venu Madhav Govindu
- 5. Robust Surface Reconstruction via Triple Sparsity, Hicham Badri, Hussein Yahia, Driss Aboutajdine
- Scattering Parameters and Surface Normals from Homogeneous Translucent Materials using Photometric Stereo, Bo Dong, Kathleen D. Moore, Weiyi Zhang, Pieter Peers
- Better Shading for Better Shape Recovery, Moumen T. El-Melegy, Aly S. Abdelrahim, Aly A. Faraq
- Stable and Informative Spectral Signatures for Graph Matching, Nan Hu, Raif M. Rustamov, Leonidas Guibas
- Deformable Object Matching via Deformation
   Decomposition based 2D Label MRF, Kangwei Liu, Junge Zhang, Kaiqi Huang, Tieniu Tan
- Locally Optimized Product Quantization for Approximate Nearest Neighbor Search, Yannis Kalantidis, Yannis Avrithis
- 11. Multi-source Deep Learning for Human Pose Estimation,
  Wanli Ouyang, Xiao Chu, Xiaogang Wang
- 12. Posebits for Monocular Human Pose Estimation, Gerard Pons-Moll, David J. Fleet, Bodo Rosenhahn
- Real-time Simultaneous Pose and Shape Estimation for Articulated Objects Using a Single Depth Camera, Mao Ye, Ruigang Yang

- 14. Mixing Body-Part Sequences for Human Pose Estimation, Anoop Cherian, Julien Mairal, Karteek Alahari, Cordelia Schmid
- Robust Estimation of 3D Human Poses from a Single Image, Chunyu Wang, Yizhou Wang, Zhouchen Lin, Alan L. Yuille, Wen Gao
- 16. Fisher and VLAD with FLAIR, Koen E. A. van de Sande, Cees G. M. Snoek, Arnold W. M. Smeulders
- 17. Immediate, Scalable Object Category Detection, Yusuf Aytar, Andrew Zisserman
- 18. Word Channel Based Multiscale Pedestrian Detection Without Image Resizing and Using Only One Classifier, Arthur Daniel Costea, Sergiu Nedevschi
- 19. Parsing Occluded People, Golnaz Ghiasi, Yi Yang, Deva Ramanan, Charless C. Fowlkes
- 20. Multi-fold MIL Training for Weakly Supervised Object Localization, Ramazan Gokberk Cinbis, Jakob Verbeek, Cordelia Schmid
- Generating Object Segmentation Proposals using Global and Local Search, Pekka Rantalankila, Juho Kannala, Esa Rahtu
- 22. A Novel Chamfer Template Matching Method Using Variational Mean Field, Duc Thanh Nguyen
- 23. Confidence-Rated Multiple Instance Boosting for Object Detection, Karim Ali, Kate Saenko
- 24. COSTA: Co-Occurrence Statistics for Zero-Shot Classification, Thomas Mensink, Efstratios Gavves, Cees G.M. Snoek
- 25. Analysis by Synthesis: 3D Object Recognition by Object Reconstruction, *Mohsen Hejrati, Deva Ramanan*
- 26. Submodular Object Recognition, Fan Zhu, Zhuolin Jiang, Ling Shao
- 27. Multimodal Learning in Loosely-organized Web Images, Kun Duan, David J. Crandall, Dhruv Batra
- 28. Generalized Max Pooling, Naila Murray, Florent Perronnin
- 29. Domain Adaptation on the Statistical Manifold, Mahsa Baktashmotlagh, Mehrtash T. Harandi, Brian C. Lovell, Mathieu Salzmann
- 30. Nonparametric Part Transfer for Fine-grained Recognition, Christoph Göring, Erik Rodner, Alexander Freytag, Joachim Denzler

- 31. The Fastest Deformable Part Model for Object Detection,
  Junjie Yan, Zhen Lei, Longyin Wen, Stan Z. Li
- 32. Unsupervised Learning of Dictionaries of Hierarchical Compositional Models, Jifeng Dai, Yi Hong, Wenze Hu, Song-Chun Zhu, Ying Nian Wu
- 33. Quasi Real-Time Summarization for Consumer Videos, Bin Zhao, Eric P. Xing
- 34. Gait Recognition under Speed Transition, Al Mansur, Yasushi Makihara, Rasyid Aqmar, Yasushi Yagi
- 35. Video Classification using Semantic Concept Cooccurrences, Shayan Modiri Assari, Amir Roshan Zamir, Mubarak Shah
- 36. Temporal Segmentation of Egocentric Videos, Yair Poleg, Chetan Arora, Shmuel Peleg
- 37. Efficient Action Localization with Approximately Normalized Fisher Vectors, Dan Oneata, Jakob Verbeek, Cordelia Schmid
- 38. Unsupervised Trajectory Modelling using Temporal Information via Minimal Paths, Brais Cancela, Alberto Iglesias, Marcos Ortega, Manuel G. Penedo
- 39. A Hierarchical Context Model for Event Recognition in Surveillance Video, *Xiaoyang Wang, Qiang Ji*
- 40. DISCOVER: Discovering Important Segments for Classification of Video Events and Recounting, Chen Sun, Ram Nevatia
- 41. Towards Good Practices for Action Video Encoding, *Jianxin Wu, Yu Zhang, Weiyao Lin*
- 42. Improving Semantic Concept Detection through the Dictionary of Visually-distinct Elements, Afshin Dehghan, Haroon Idrees, Mubarak Shah
- 43. Efficient Feature Extraction, Encoding and Classification for Action Recognition, Vadim Kantorov, Ivan Laptev
- 44.3D Pose from Motion for Cross-view Action Recognition via Non-linear Circulant Temporal Encoding, Ankur Gupta, Julieta Martinez, James J. Little, Robert J. Woodham
- 45. Human Action Recognition Based on Context-Dependent Graph Kernels, Baoxin Wu, Chunfeng Yuan, Weiming Hu
- 46. Depth and Skeleton Associated Action Recognition without Online Accessible RGB-D Cameras, Yen-Yu Lin, Ju-Hsuan Hua, Nick C. Tang, Min-Hung Chen, Hong-Yuan Mark Liao

- 47. DL-SFA: Deeply-Learned Slow Feature Analysis for Action Recognition, *Lin Sun, Kui Jia, Tsung-Han Chan, Yuqiang* Fanq, Ganq Wanq, Shuichenq Yan
- 48. A Cause and Effect Analysis of Motion Trajectories for Modeling Actions, Sanath Narayan, Kalpathi R. Ramakrishnan
- 49. From Stochastic Grammar to Bayes Network: Probabilistic Parsing of Complex Activity, Nam N. Vo, Aaron F. Bobick
- 50. Cross-view Action Modeling, Learning and Recognition, Jiang Wang, Xiaohan Nie, Yin Xia, Ying Wu, Song-Chun Zhu
- 51. Visual Semantic Search: Retrieving Videos via Complex Textual Queries, Dahua Lin, Sanja Fidler, Chen Kong, Raquel Urtasun
- 52. Zero-shot Event Detection using Multi-modal Fusion of Weakly Supervised Concepts, Shuang Wu, Sravanthi Bondugula, Florian Luisier, Xiaodan Zhuang, Pradeep Natarajan
- 53. Dual Linear Regression Based Classification for Face Cluster Recognition, Liang Chen
- 54. Bags of Spacetime Energies for Dynamic Scene Recognition, Christoph Feichtenhofer, Axel Pinz, Richard P. Wildes
- 55. Feature-Independent Action Spotting Without Human Localization, Segmentation or Frame-wise Tracking, Chuan Sun, Marshall Tappen, Hassan Foroosh

### **1200–1330 Lunch** (Exhibit Hall C)

### Thursday, June 26 (Afternoon)

### **Program**

### 1330-1830 PM Video Spotlights (C213-215)

### 1330-1500 Oral 3C: Medical & Biological Image Analysis (Battelle Grand South)

Poster IDs for this session: O-3C-# where # is the paper #.

Chairs: Petia Radeva (Univ. of Barcelona)
Ioannis Kakadiaris (Univ. of Houston)

Format (13 min. for presentation + 2 min. for questions)

- Multiscale Centerline Detection by Learning a Scale-Space Distance Transform, Amos Sironi, Vincent Lepetit, Pascal Fua
- Multivariate General Linear Models (MGLM) on Riemannian Manifolds with Applications to Statistical Analysis of Diffusion Weighted Images, Hyunwoo J. Kim, Nagesh Adluru, Maxwell D. Collins, Moo K. Chung, Barbara B. Bendlin, Sterling C. Johnson, Richard J. Davidson, Vikas Singh
- Preconditioning for Accelerated Iteratively Reweighted Least Squares in Structured Sparsity Reconstruction, Chen Chen, Junzhou Huang, Lei He, Hongsheng Li
- Joint Coupled-Feature Representation and Coupled Boosting for AD Diagnosis, Yinghuan Shi, Heung-Il Suk, Yang Gao, Dinggang Shen
- Deformable Registration of Feature-Endowed Point Sets Based on Tensor Fields, Demian Wassermann, James Ross, George Washko, William M. Wells III, Raul San Jose-Estepar
- Tracking Indistinguishable Translucent Objects over Time using Weakly Supervised Structured Learning, Luca Fiaschi, Ferran Diego, Konstantin Gregor, Martin Schiegg, Ullrich Koethe, Marta Zlatic, Fred A. Hamprecht

### 1330–1500 Oral 3D: Low-Level Vision & Image Processing (Battelle Grand North)

Poster IDs for this session: O-3D-# where # is the paper #.

Chairs: Michael Brown (National Univ. of Singapore)
Stefan Roth (TU Darmstadt)

Format (13 min. for presentation + 2 min. for questions)

- Scale-space Processing Using Polynomial
   Representations, Gou Koutaki, Keiichi Uchimura
- 2. Single Image Layer Separation using Relative Smoothness, Yu Li, Michael S. Brown

- Image Fusion with Local Spectral Consistency and Dynamic Gradient Sparsity, Chen Chen, Yeqing Li, Wei Liu, Junzhou Huanq
- 4. Segmentation-Free Dynamic Scene Deblurring, *Tae Hyun Kim, Kyoung Mu Lee*
- Shrinkage Fields for Effective Image Restoration, Uwe Schmidt, Stefan Roth
- Camouflaging an Object from Many Viewpoints, Andrew Owens, Connelly Barnes, Alex Flint, Hanumant Singh, William Freeman

### 1500-1530 Break (Battelle Grand Prefunction)

## **1530–1630 Special 3: Plenary Session** (Battelle Grand)

Chairs: Rene Vidal (Johns Hopkins Univ.)
Ronen Basri (Weizmann Inst. of Science)

 Plenary Talk: Are Deep Networks a Solution to the Curse of Dimensionality?, Stéphane Mallat (École Normale Supérieure)

**Abstract:** Learning gave a considerable and surprising boost to computer vision, and deep neural networks appear to be the new winners of the fierce race on classification errors. Algorithm refinements are now going well beyond our understanding of the problem, and seem to make irrelevant any study of computer vision models. Yet, learning from high-dimensional data such as images, suffers from a curse of dimensionality which predicts a combinatorial explosion. Why are these neural architectures avoiding this curse? Is this rooted in properties of images and visual tasks? Can these properties be related to high-dimensional problems in other fields? We shall explore the mathematical roots of these questions, and tell a story where invariants, contractions, sparsity, dimension reduction and multiscale analysis play important roles. Images and examples will give a colorful background to the talk.

### **1630–1830 Exhibits** (Grand Ballrooms 1-3)

• Same as Tuesday morning Exhibits (see pg. 15)

### **1630–1830 Demos** (C110-115)

 PanOptus: Automatic Video Editing for iPhone, Google Glass and Surveillance Camera, Bin Zhao, Bin Shu, Eric Xing (CMU)

- Surface-from-Gradient (SfG) by Discrete Geometry Processing (DGP), Wuyuan Xie, Yunbo Zhang, Harlie C.L. Wang, Ronald C.-K. Chung (Chinese Univ. of Hong Kong)
- Complex Activity Detection and Functional Scene Understanding in Video, Eran Swears, Anthony Hoogs, Sangmin Oh, Matt Leotta (Kitware Inc.)
- CloudCV: Large Scale Distributed Computer Vision as a Cloud Service, Harsh Agrawal, Neelima Chavali, Clint Solomon Mathialagan, Abdullah Alfadda, Prakriti Banik, Dhruv Batra (Virginia Tech)

### 1630–1830 Poster 3B: Biologically Inspired Vision, Low-Level Vision, Medical & Biological Image Analysis, Segmentation (Grand Ballrooms 1-3)

#### Poster IDs for this session: P-3B-# where # is the paper #.

- Learning Optimal Seeds for Diffusion-based Salient Object Detection, Song Lu, Vijay Mahadevan, Nuno Vasconcelos
- Large-Scale Optimization of Hierarchical Features for Saliency Prediction in Natural Images, Eleonora Vig, Michael Dorr, David Cox
- Saliency Detection on Light Field, Nianyi Li, Jinwei Ye, Yu Ji, Haibin Ling, Jingyi Yu
- Saliency Optimization from Robust Background Detection, Wangjiang Zhu, Shuang Liang, Yichen Wei, Jian Sun
- A Reverse Hierarchy Model for Predicting Eye Fixations, Tianlin Shi, Ming Liang, Xiaolin Hu
- 100+ Times Faster Weighted Median Filter (WMF), Qi Zhang, Li Xu, Jiaya Jia
- Edge-aware Gradient Domain Optimization Framework for Image Filtering by Local Propagation, Miao Hua, Xiaohui Bie, Minying Zhang, Wencheng Wang
- 8. Super-Resolving Noisy Images, Abhishek Singh, Fatih Porikli, Narendra Ahuja
- Sparse Dictionary Learning for Edit Propagation of High-Resolution Images, Xiaowu Chen, Dongqing Zou, Jianwei Li, Xiaochun Cao, Qinping Zhao, Hao Zhang
- Weighted Nuclear Norm Minimization with Application to Image Denoising, Shuhang Gu, Lei Zhang, Wangmeng Zuo, Xiangchu Feng
- 11. Using Projection Kurtosis Concentration Of Natural Images For Blind Noise Covariance Matrix Estimation, Xing Zhang, Siwei Lyu

- 12. Blind Image Quality Assessment using Semi-supervised Rectifier Networks, Huixuan Tang, Neel Joshi, Ashish Kapoor
- 13. Separable Kernel for Image Deblurring, Lu Fang, Haifeng Liu, Feng Wu, Xiaoyan Sun, Hougiang Li
- 14. Joint Depth Estimation and Camera Shake Removal from Single Blurry Image, *Zhe Hu, Li Xu, Ming-Hsuan Yang*
- 15. Deblurring Text Images via L<sub>o</sub>-Regularized Intensity and Gradient Prior, Jinshan Pan, Zhe Hu, Zhixun Su, Ming-Hsuan Yang
- 16. Total Variation Blind Deconvolution: The Devil is in the Details, Daniele Perrone, Paolo Favaro
- 17. Single Image Super-resolution using Deformable Patches, Yu Zhu, Yanning Zhang, Alan L. Yuille
- 18. Multi-Shot Imaging: Joint Alignment, Deblurring and Resolution-Enhancement, *Haichao Zhang, Lawrence Carin*
- CID: Combined Image Denoising in Spatial and Frequency Domains Using Web Images, Huanjing Yue, Xiaoyan Sun, Jingyu Yang, Feng Wu
- Multipoint Filtering with Local Polynomial Approximation and Range Guidance, Xiao Tan, Changming Sun, Tuan D. Pham
- 21. Decomposable Nonlocal Tensor Dictionary Learning for Multispectral Image Denoising, Yi Peng, Deyu Meng, Zongben Xu, Chenqiang Gao, Yi Yang, Biao Zhang
- 22. Robust 3D Features for Matching between Distorted Range Scans Captured by Moving Systems, Xiangqi Huang, Bo Zheng, Takeshi Masuda, Katsushi Ikeuchi
- 23. Discriminative Blur Detection Features, Jianping Shi, Li Xu, Jiaya Jia
- 24. Detection, Rectification and Segmentation of Coplanar Repeated Patterns, James Pritts, Ondřej Chum, Jiří Matas
- 25. Mirror Symmetry Histograms for Capturing Geometric Properties in Images, Marcelo Cicconet, Davi Geiger, Kristin C Gunsalus, Michael Werman
- 26. A Learning-to-Rank Approach for Image Color Enhancement, Jianzhou Yan, Stephen Lin, Sing Bing Kang, Xiaoou Tang
- 27. Investigating Haze-relevant Features in A Learning Framework for Image Dehazing, *Ketan Tang, Jianchao Yang, Jue Wang*

- 28. Quality Assessment for Comparing Image Enhancement Algorithms, *Zhengying Chen, Tingting Jiang, Yonghong Tian*
- 29. Shadow Removal from Single RGB-D Images, Yao Xiao, Efstratios Tsougenis, Chi-Keung Tang
- 30. Manifold Based Dynamic Texture Synthesis from Extremely Few Samples, Hongteng Xu, Hongyuan Zha, Mark A. Davenport
- 31. The Synthesizability of Texture Examples, *Dengxin Dai, Hayko Riemenschneider, Luc Van Gool*
- 32. Reconstructing Evolving Tree Structures in Time Lapse Sequences, Przemysław Głowacki, Miguel Amável Pinheiro, Engin Türetken, Raphael Sznitman, Daniel Lebrecht, Jan Kybic, Anthony Holtmaat, Pascal Fua
- 33. Total-Variation Minimization on Unstructured Volumetric Mesh: Biophysical Applications on Reconstruction of 3D Ischemic Myocardium, Jingjia Xu, Azar Rahimi Dehaghani, Fei Gao, Linwei Wang
- 34. Tracking on the Product Manifold of Shape and Orientation for Tractography from Diffusion MRI, Yuanxiang Wang, Hesamoddin Salehian, Guang Cheng, Baba C. Vemuri
- 35. Curvilinear Structure Tracking by Low Rank Tensor Approximation with Model Propagation, Erkang Cheng, Yu Pang, Ying Zhu, Jingyi Yu, Haibin Ling
- 36. Patch-based Evaluation of Image Segmentation, *Christian Ledig, Wenzhe Shi, Wenjia Bai, Daniel Rueckert*
- 37. Evaluation of Scan-Line Optimization for 3D Medical Image Registration, *Simon Hermann*
- 38. Classification of Histology Sections via Multispectral Convolutional Sparse Coding, Yin Zhou, Hang Chang, Kenneth Barner, Paul Spellman, Bahram Parvin
- 39. Matrix-Similarity Based Loss Function and Feature Selection for Alzheimer's Disease Diagnosis, Xiaofeng Zhu, Heung-Il Suk, Dinggang Shen
- 40. Discriminative Sparse Inverse Covariance Matrix:
  Application in Brain Functional Network Classification,
  Luping Zhou, Lei Wang, Philip Ogunbona
- 41. A Bayesian Framework For the Local Configuration of Retinal Junctions, Touseef Ahmad Qureshi, Andrew Hunter, Bashir Al-Diri

- 42. Learning-Based Atlas Selection for Multiple-Atlas Segmentation, *Gerard Sanroma*, *Guorong Wu*, *Yaozong Gao*, *Dinggang Shen*
- 43. Fully Automated Non-rigid Segmentation with Distance Regularized Level Set Evolution Initialized and Constrained by Deep-structured Inference, Tuan Anh Ngo, Gustavo Carneiro
- 44. FAST LABEL: Easy and Efficient Solution of Joint Multi-Label and Estimation Problems, *Ganesh Sundaramoorthi, Byung-Woo Hong*
- 45. Learning to Group Objects, *Victoria Yanulevskaya, Jasper Uijlings, Nicu Sebe*
- 46. Unsupervised Multi-Class Joint Image Segmentation, Fan Wang, Qixing Huang, Maks Ovsjanikov, Leonidas J. Guibas
- 47. Semantic Object Selection, *Ejaz Ahmed, Scott Cohen, Brian Price*
- 48. Discrete-Continuous Gradient Orientation Estimation for Faster Image Segmentation, *Michael Donoser, Dieter Schmalstieg*
- Object-based Multiple Foreground Video Cosegmentation, Huazhu Fu, Dong Xu, Bao Zhang, Stephen Lin
- Lin 50. Parsing World's Skylines using Shape-Constrained MRFs, Rashmi Tonge, Subhransu Maji, C. V. Jawahar
- Clothing Co-Parsing by Joint Image Segmentation and Labeling, Wei Yang, Ping Luo, Liang Lin
- 52. Tell Me What You See and I will Show You Where It Is, Jia Xu, Alexander G. Schwing, Raquel Urtasun
- 53. Beat the MTurkers: Automatic Image Labeling from Weak 3D Supervision, Liang-Chieh Chen, Sanja Fidler, Alan L. Yuille, Raquel Urtasun
- 54. Efficient Structured Parsing of Façades Using Dynamic Programming, Andrea Cohen, Alexander G. Schwing, Marc Pollefeys
- 55. Dense Semantic Image Segmentation with Objects and Attributes, Shuai Zheng, Ming-Ming Cheng, Jonathan Warrell, Paul Sturgess, Vibhav Vineet, Carsten Rother, Philip H. S. Torr

### 1830-2030 Reception (Battelle Grand)

## Friday, June 27

0700-1700 Registration (Exhibit Hall C Lobby)

0730-0830 Breakfast (Exhibit Hall C)

**0830-1200** AM Video Spotlights (C213-215)

# 0830-1000 Oral 4A: Computational Photography: Sensing and Display (Battelle Grand South)

Poster IDs for this session: O-4A-# where # is the paper #.

Chairs: Hui Ji (National Univ. of Singapore) Kyros Kutulakos (Univ. of Toronto)

Format (13 min. for presentation + 2 min. for questions)

- Diffuse Mirrors: 3D Reconstruction from Diffuse Indirect Illumination Using Inexpensive Time-of-Flight Sensors, Felix Heide, Lei Xiao, Wolfgang Heidrich, Matthias B. Hullin
- Fourier Analysis on Transient Imaging with a Multifrequency Time-of-Flight Camera, Jingyu Lin, Yebin Liu, Matthias B. Hullin, Qionghai Dai
- Transparent Object Reconstruction via Coded Transport of Intensity, Chenguang Ma, Xing Lin, Jinli Suo, Qionghai Dai, Gordon Wetzstein
- 3D Shape and Indirect Appearance by Structured Light Transport, Matthew O'Toole, John Mather, Kiriakos N. Kutulakos
- Shape-Preserving Half-Projective Warps for Image Stitching, Che-Han Chang, Yoichi Sato, Yung-Yu Chuang
- 6. Parallax-tolerant Image Stitching, Fan Zhang, Feng Liu

### 0830-1000 Oral 4B: Recognition: Detection, Categorization, Classification (Battelle Grand North)

Poster IDs for this session: O-4B-# where # is the paper #.

Chairs: Haibin Ling (Temple Univ.)
Jiri Matas (Czech Technical Univ.)

Format (13 min. for presentation + 2 min. for questions)

- Learning Everything about Anything: Webly-Supervised Visual Concept Learning, Santosh K. Divvala, Ali Farhadi, Carlos Guestrin
- Dirichlet-based Histogram Feature Transform for Image Classification, Takumi Kobayashi
- BING: Binarized Normed Gradients for Objectness
   Estimation at 300fps, Ming-Ming Cheng, Ziming Zhang,
   Wen-Yan Lin, Philip Torr
- Context Driven Scene Parsing with Attention to Rare Classes, Jimei Yang, Brian Price, Scott Cohen, Ming-Hsuan Yang
- 5. Patch to the Future: Unsupervised Visual Prediction, *Jacob Walker, Abhinav Gupta, Martial Hebert*
- Triangulation Embedding and Democratic Aggregation for Image Search, Hervé Jégou, Andrew Zisserman

1000-1030 Break (Grand Ballroom Prefunction)

### 1000-1200 Exhibits (Grand Ballrooms 1-3)

• Same as Tuesday morning Exhibits (see pg. 15)

### 1000-1200 Demos (C110-115)

- Mobile Vision for Community Care: A Framework for People with Dementia, Moi Hoon Yap, Choon-Ching Ng, Bill Cassidy, Gemma Stringer (Manchester Metropolitan Univ.)
- Large-Scale, Real Time Visual-Inertial Navigation and Mapping, Konstantine Tsotsos, Stephen Phillips, Stefano Soatto (UCLA)

## 1000-1200 Poster 4A: Computational Photography, Motion & Tracking, Recognition (Grand Ballrooms 1-3)

#### Poster IDs for this session: P-4A-# where # is the paper #.

- Low-Cost Compressive Sensing for Color Video and Depth, Xin Yuan, Patrick Llull, Xuejun Liao, Jianbo Yang, David J. Brady, Guillermo Sapiro, Lawrence Carin
- 2. Aliasing Detection and Reduction in Plenoptic Imaging, Zhaolin Xiao, Qing Wang, Guoqing Zhou, Jingyi Yu
- 3. Illumination-Aware Age Progression, Ira Kemelmacher-Shlizerman, Supasorn Suwajanakorn, Steven M. Seitz
- Color Transfer Using Probabilistic Moving Least Squares, Youngbae Hwang, Joon-Young Lee, In So Kweon, Seon Joo Kim
- Image Pre-compensation: Balancing Contrast and Ringing, Yu Ji, Jinwei Ye, Sing Bing Kang, Jingyi Yu
- 6. Time-Mapping Using Space-Time Saliency, Feng Zhou, Sing Bing Kang, Michael F. Cohen
- Gyro-Based Multi-Image Deconvolution for Removing Handshake Blur, Sung Hee Park, Marc Levoy
- Similarity-Aware Patchwork Assembly for Depth Image Super-Resolution, Jing Li, Zhichao Lu, Gang Zeng, Rui Gan, Hongbin Zha
- 9. Deblurring Low-light Images with Light Streaks, *Zhe Hu,* Sunghyun Cho, Jue Wang, Ming-Hsuan Yang
- 10. Depth Enhancement via Low-rank Matrix Completion, Si Lu, Xiaofeng Ren, Feng Liu
- 11. Raw-to-Raw: Mapping between Image Sensor Color Responses, Rang Nguyen, Dilip K. Prasad, Michael S. Brown
- 12. DAISY Filter Flow: A Generalized Discrete Approach to Dense Correspondences, Hongsheng Yang, Wen-Yan Lin, Jiangbo Lu
- 13. Robust 3D Tracking with Descriptor Fields, Alberto Crivellaro, Vincent Lepetit
- 14. Evolutionary Quasi-random Search for Hand Articulations Tracking, Iason Oikonomidis, Manolis I.A. Lourakis, Antonis A. Argyros
- 15. Scalable 3D Tracking of Multiple Interacting Objects, Nikolaos Kyriazis, Antonis Argyros
- 16. Bayesian Active Appearance Models, Joan Alabort-i-Medina, Stefanos Zafeiriou

- 17. Human Shape and Pose Tracking Using Keyframes, Chun-Hao Huang, Edmond Boyer, Nassir Navab, Slobodan Ilic
- 18. Better Feature Tracking Through Subspace Constraints, Bryan Poling, Gilad Lerman, Arthur Szlam
- 19. Online Object Tracking, Learning and Parsing with And-Or Graphs, Yang Lu, Tianfu Wu, Song Chun Zhu
- 20. Region-based Particle Filter for Video Object Segmentation, *David Varas, Ferran Marques*
- 21. Visual Tracking via Probability Continuous Outlier Model, Dong Wang, Huchuan Lu
- 22. Visual Tracking Using Pertinent Patch Selection and Masking, Dae-Youn Lee, Jae-Young Sim, Chang-Su Kim
- 23. Interval Tracker: Tracking by Interval Analysis, *Junseok Kwon, Kyoung Mu Lee*
- 24. Unifying Spatial and Attribute Selection for Distracter-Resilient Tracking, Nan Jiang, Ying Wu
- 25. Pedestrian Detection in Low-resolution Imagery by Learning Multi-scale Intrinsic Motion Structures (MIMS), Jiejie Zhu, Omar Javed, Jingen Liu, Qian Yu, Hui Cheng, Harpreet Sawhney
- 26. Multi-target Tracking with Motion Context in Tensor Power Iteration, Xinchu Shi, Haibin Ling, Weiming Hu, Chunfeng Yuan, Junliang Xing
- 27. SphereFlow: 6 DoF Scene Flow from RGB-D Pairs, Michael Hornáček, Andrew Fitzgibbon, Carsten Rother
- 28. Fast Edge-Preserving PatchMatch for Large Displacement Optical Flow, *Linchao Bao, Qingxiong Yang, Hailin Jin*
- 29. Learning an Image-based Motion Context for Multiple People Tracking, Laura Leal-Taixé, Michele Fenzi, Alina Kuznetsova, Bodo Rosenhahn, Silvio Savarese
- 3o. Semi-Supervised Coupled Dictionary Learning for Person Re-identification, Xiao Liu, Mingli Song, Dacheng Tao, Xingchen Zhou, Chun Chen, Jiajun Bu
- 31. What are You Talking About? Text-to-Image Coreference, Chen Kong, Dahua Lin, Mohit Bansal, Raquel Urtasun, Sanja Fidler
- 32. Predicting Failures of Vision Systems, Peng Zhang, Jiuling Wang, Ali Farhadi, Martial Hebert, Devi Parikh
- 33. Three Guidelines of Online Learning for Large-Scale Visual Recognition, Yoshitaka Ushiku, Masatoshi Hidaka, Tatsuya Harada

- 34. Using k-Poselets for Detecting People and Localizing Their Keypoints, Georgia Gkioxari, Bharath Hariharan, Ross Girshick, Jitendra Malik
- 35. Randomized Max-Margin Compositions for Visual Recognition, Angela Eigenstetter, Masato Takami, Björn Ommer
- 36. Large-Scale Visual Font Recognition, Guang Chen, Jianchao Yang, Hailin Jin, Jonathan Brandt, Eli Shechtman, Aseem Agarwala, Tony X. Han
- 37. Describing Textures in the Wild, Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, Sammy Mohamed, Andrea Vedaldi
- 38. Relative Parts: Distinctive Parts for Learning Relative Attributes, Ramachandruni N. Sandeep, Yashaswi Verma, C. V. Jawahar
- 39. Understanding Objects in Detail with Fine-Grained Attributes, Andrea Vedaldi, Siddharth Mahendran, Stavros Tsogkas, Subhransu Maji, Ross Girshick, Juho Kannala, Esa Rahtu, Iasonas Kokkinos, Matthew B. Blaschko, David Weiss, Ben Taskar, Karen Simonyan, Naomi Saphra, Sammy Mohamed
- 40. Predicting User Annoyance Using Visual Attributes, Gordon Christie, Amar Parkash, Ujwal Krothapalli, Devi Parikh
- 41. Linear Ranking Analysis, Weihong Deng, Jiani Hu, Jun Guo
- 42. Transformation Pursuit for Image Classification, Mattis Paulin, Jérôme Revaud, Zaid Harchaoui, Florent Perronnin, Cordelia Schmid
- 43. Incremental Learning of NCM Forests for Large-Scale Image Classification, Marko Ristin, Matthieu Guillaumin, Juergen Gall, Luc Van Gool
- 44. Object Classification with Adaptable Regions, Hakan Bilen, Marco Pedersoli, Vinay P. Namboodiri, Tinne Tuytelaars, Luc Van Gool
- 45. Discriminative Ferns Ensemble for Hand Pose Recognition, Eyal Krupka, Alon Vinnikov, Ben Klein, Aharon Bar Hillel, Daniel Freedman, Simon Stachniak
- 46. Are Cars Just 3D Boxes? Jointly Estimating the 3D Shape of Multiple Objects, Muhammad Zeeshan Zia, Michael Stark, Konrad Schindler
- 47. 2D Human Pose Estimation: New Benchmark and State of the Art Analysis, Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, Bernt Schiele

- 48. Using a Deformation Field Model for Localizing Faces and Facial Points under Weak Supervision, Marco Pedersoli, Tinne Tuytelaars, Luc Van Gool
- 49. Active Annotation Translation, Steve Branson, Kristján Eldjárn Hjörleifsson, Pietro Perona
- 50. Looking Beyond the Visible Scene, Aditya Khosla, Byoungkwon An An, Joseph J. Lim, Antonio Torralba
- 51. Two-Class Weather Classification, Cewu Lu, Di Lin, Jiaya Jia, Chi-Keung Tang
- 52. Learning Important Spatial Pooling Regions for Scene Classification, Di Lin, Cewu Lu, Renjie Liao, Jiaya Jia
- 53. Orientational Pyramid Matching for Recognizing Indoor Scenes, Lingxi Xie, Jingdong Wang, Baining Guo, Bo Zhang, Oi Tian
- 54. Multilabel Ranking with Inconsistent Rankers, Xin Geng, Longrun Luo
- 55. Scene Parsing with Object Instances and Occlusion Ordering, Joseph Tighe, Marc Niethammer, Svetlana Lazebnik

## 1200-1330 Lunch (Exhibit Hall C)

## 1330-1830 PM Video Spotlights (C213-215)

# 1330-1500 Oral 4C: 3D Geometry & Shape (Battelle Grand South)

#### Poster IDs for this session: O-4C-# where # is the paper #.

Chairs: Lourdes Agapito (Univ. College London)
Paolo Favaro (Univ. of Bern)

Format (13 min. for presentation + 2 min. for questions)

- A Riemannian Framework for Matching Point Clouds Represented by the Schrödinger Distance Transform, Yan Deng, Anand Rangarajan, Stephan Eisenschenk, Baba C. Vemuri
- Seeing 3D Chairs: Exemplar Part-based 2D-3D Alignment using a Large Dataset of CAD Models, Mathieu Aubry, Daniel Maturana, Alexei A. Efros, Bryan C. Russell, Josef Sivic
- A Mixture of Manhattan Frames: Beyond the Manhattan World, Julian Straub, Guy Rosman, Oren Freifeld, John J. Leonard, John W. Fisher III
- 4. Local Regularity-driven City-scale Facade Detection from Aerial Images, *Jingchen Liu, Yanxi Liu*
- Latent Regression Forest: Structured Estimation of 3D Articulated Hand Posture, Danhang Tang, Hyung Jin Chang, Alykhan Tejani, Tae-Kyun Kim
- 6. FAUST: Dataset and Evaluation for 3D Mesh Registration, Federica Bogo, Javier Romero, Matthew Loper, Michael J. Black

## 1330-1500 Oral 4D: Statistical Methods and Learning II (Battelle Grand North)

Poster IDs for this session: O-4D-# where # is the paper #.

Chairs: Jason Corso (SUNY Buffalo) Zhouchen Lin (Peking Univ.)

Format (13 min. for presentation + 2 min. for questions)

- Optimizing Over Radial Kernels on Compact Manifolds, Sadeep Jayasumana, Richard Hartley, Mathieu Salzmann, Hongdong Li, Mehrtash Harandi
- 2. Grassmann Averages for Scalable Robust PCA, Søren Hauberg, Aasa Feragen, Michael J. Black
- 3. Robust Subspace Segmentation with Block-diagonal Prior, Jiashi Feng, Zhouchen Lin, Huan Xu, Shuicheng Yan
- 4. Unsupervised One-Class Learning for Automatic Outlier Removal, Wei Liu, Gang Hua, John R. Smith
- Smooth Representation Clustering, Han Hu, Zhouchen Lin, Jianjiang Feng, Jie Zhou
- Novel Methods for Multilinear Data Completion and Denoising Based on Tensor-SVD, Zemin Zhang, Gregory Ely, Shuchin Aeron, Ning Hao, Misha Kilmer

## 1500-1530 Break (Battelle Grand Prefunction)

# **1530–1630 Oral 4E: Optimization Methods** (Battelle Grand South)

Poster IDs for this session: O-4E-# where # is the paper #.

Chairs: Allen Yang (UC Berkeley) Yi Ma (ShanghaiTech Univ. & UIUC)

Format (13 min. for presentation + 2 min. for questions)

- 1. Second-Order Shape Optimization for Geometric Inverse Problems in Vision, *Jonathan Balzer, Stefano Soatto*
- \$\ell\_0\$ Norm Based Dictionary Learning by Proximal Methods with Global Convergence, Chenglong Bao, Hui Ji, Yuhui Quan, Zuowei Shen
- Adaptive Partial Differential Equation Learning for Visual Saliency Detection, Risheng Liu, Junjie Cao, Zhouchen Lin, Shiguang Shan
- Robust Orthonormal Subspace Learning: Efficient Recovery of Corrupted Low-rank Matrices, Xianbiao Shu, Fatih Porikli, Narendra Ahuja

## Friday, June 27 (Afternoon)

## **Program**

## 1530-1630 Oral 4F: View Synthesis & Other Applications (Battelle Grand North)

### Poster IDs for this session: O-4F-# where # is the paper #.

Chairs: Kostas Daniilidis (Univ. of Pennsylvania) Gérard Medioni (USC)

#### Format (13 min. for presentation + 2 min. for questions)

- Reconstructing Storyline Graphs for Image Recommendation from Web Community Photos, Gunhee Kim, Eric P. Xing
- Active Flattening of Curved Document Images via Two Structured Beams, Gaofeng Meng, Ying Wang, Shenquan Qu, Shiming Xiang, Chunhong Pan
- Image-based Synthesis and Re-Synthesis of Viewpoints Guided by 3D Models, Konstantinos Rematas, Tobias Ritschel, Mario Fritz, Tinne Tuytelaars
- 4. Bayesian View Synthesis and Image-Based Rendering Principles, Sergi Pujades, Frédéric Devernay, Bastian Goldluecke

## 1630-1830 Exhibits (Grand Ballrooms 1-3)

• Same as Tuesday morning Exhibits (see pg. 15)

## 1630–1830 Poster 4B: 3D Vision, Document Analysis, Optimization Methods, Shape, Vision for Graphics, Web & Vision Systems

(Grand Ballrooms 1-3)

## Poster IDs for this session: P-4B-# where # is the paper #.

- 1. Fast MRF Optimization with Application to Depth Reconstruction, *Qifeng Chen, Vladlen Koltun*
- Exploiting Shading Cues in Kinect IR Images for Geometry Refinement, Gyeongmin Choe, Jaesik Park, Yu-Wing Tai, In So Kweon
- Fast Rotation Search with Stereographic Projections for 3D Registration, Álvaro Parra Bustos, Tat-Jun Chin, David Suter
- Local Readjustment for High-Resolution 3D
   Reconstruction, Siyu Zhu, Tian Fang, Jianxiong Xiao, Long
   Quan

- 5. Turning Mobile Phones into 3D Scanners, *Kalin Kolev*, *Petri Tanskanen*, *Pablo Speciale*, *Marc Pollefeys*
- T-Linkage: A Continuous Relaxation of J-Linkage for Multi-Model Fitting, Luca Magri, Andrea Fusiello
- Motion-Depth: RGB-D Depth Map Enhancement with Motion and Depth in Complement, Tak-Wai Hui, King Ngi Ngan
- Generalized Pupil-Centric Imaging and Analytical Calibration for a Non-frontal Camera, Avinash Kumar, Narendra Ahuja
- Geometric Urban Geo-Localization, Mayank Bansal, Kostas Daniilidis
- 10. 3D Reconstruction from Accidental Motion, Fisher Yu, David Gallup
- 11. Real-time Model-based Articulated Object Pose Detection and Tracking with Variable Rigidity Constraints, Karl Pauwels, Leonardo Rubio, Eduardo Ros
- Occluding Contours for Multi-View Stereo, Qi Shan, Brian Curless, Yasutaka Furukawa, Carlos Hernandez, Steven M. Seitz
- 13. Aerial Reconstructions via Probabilistic Data Fusion, Randi Cabezas, Oren Freifeld, Guy Rosman, John W. Fisher III
- 14. 3D Modeling from Wide Baseline Range Scans using Contour Coherence, Ruizhe Wang, Jongmoo Choi, Gérard Medioni
- 15. Ground Plane Estimation using a Hidden Markov Model, Ralf Dragon, Luc Van Gool
- 16. Orientation Robust Text Line Detection in Natural Images, Le Kang, Yi Li, David, Doermann
- Strokelets: A Learned Multi-Scale Representation for Scene Text Recognition, Cong Yao, Xiang Bai, Baoguang Shi, Wenyu Liu
- 18. Region-based Discriminative Feature Pooling for Scene Text Recognition, Chen-Yu Lee, Anurag Bhardwaj, Wei Di, Vignesh Jagadeesh, Robinson Piramuthu
- 19. Fast and Exact: ADMM-Based Discriminative Shape Segmentation with Loopy Part Models, Haithem Boussaid, Iasonas Kokkinos
- 20. Pseudoconvex Proximal Splitting for L<sub>∞</sub> Problems in Multiview Geometry, Anders Eriksson, Mats Isaksson

- 21. A Convex Relaxation of the Ambrosio—Tortorelli Elliptic Functionals for the Mumford-Shah Functional, Youngwook Kee, Junmo Kim
- 22. Multi Label Generic Cuts: Optimal Inference in Multi Label Multi Clique MRF-MAP Problems, Chetan Arora, S.N. Maheshwari
- 23. Sequential Convex Relaxation for Mutual Information-Based Unsupervised Figure-Ground Segmentation, Youngwook Kee, Mohamed Souiai, Daniel Cremers, Junmo Kim
- 24. Decorrelated Vectorial Total Variation, Shunsuke Ono, Isao Yamada
- 25. Efficient Squared Curvature, Claudia Nieuwenhuis, Eno Toeppe, Lena Gorelick, Olga Veksler, Yuri Boykov
- 26. Multi-feature Spectral Clustering with Minimax Optimization, Hongxing Wang, Chaoqun Weng, Junsong Yuan
- 27. Quality-based Multimodal Classification using Tree-Structured Sparsity, Soheil Bahrampour, Asok Ray, Nasser M. Nasrabadi, Kenneth W. Jenkins
- 28. Newton Greedy Pursuit: A Quadratic Approximation Method for Sparsity-Constrained Optimization, Xiao-Tong Yuan, Qingshan Liu
- 29. Generalized Nonconvex Nonsmooth Low-Rank Minimization, Canyi Lu, Jinhui Tang, Shuicheng Yan, Zhouchen Lin
- 30. Latent Dictionary Learning for Sparse Representation based Classification, Meng Yang, Dengxin Dai, Lilin Shen, Luc Van Gool
- Is Rotation a Nuisance in Shape Recognition?, Qiuhong Ke, Yi Li
- 32. Dual-Space Decomposition of 2D Complex Shapes, Guilin Liu, Zhonghua Xi, Jyh-Ming Lien
- 33. Noising versus Smoothing for Vertex Identification in Unknown Shapes, Konstantinos A. Raftopoulos, Marin Ferecatu
- 34. Surface Registration by Optimization in Constrained Diffeomorphism Space, Wei Zeng, Lok Ming Lui, Xianfeng Gu
- 35. Dense Non-Rigid Shape Correspondence using Random Forests, Emanuele Rodolà, Samuel Rota Bulò, Thomas Windheuser, Matthias Vestner, Daniel Cremers

- 36. Covariance Descriptors for 3D Shape Matching and Retrieval, Hedi Tabia, Hamid Laga, David Picard, Philippe-Henri Gosselin
- 37. Symmetry-Aware Nonrigid Matching of Incomplete 3D Surfaces, Yusuke Yoshiyasu, Eiichi Yoshida, Kazuhito Yokoi, Ryusuke Sagawa
- 38. An Automated Estimator of Image Visual Realism Based on Human Cognition, Shaojing Fan, Tian-Tsong Ng, Jonathan S. Herberg, Bryan L. Koenig, Cheston Y.-C. Tan, Rangding Wang
- 39. SteadyFlow: Spatially Smooth Optical Flow for Video Stabilization, Shuaicheng Liu, Lu Yuan, Ping Tan, Jian Sun
- 40. Automatic Face Reenactment, Pablo Garrido, Levi Valgaerts, Ole Rehmsen, Thorsten Thormählen, Patrick Pérez, Christian Theobalt
- 41. Joint Summarization of Large-scale Collections of Web Images and Videos for Storyline Reconstruction, *Gunhee Kim, Leonid Sigal, Eric P. Xing*
- 42. Semi-supervised Relational Topic Model for Weakly Annotated Image Recognition in Social Media, *Zhenxing* Niu, Gang Hua, Xinbo Gao, Qi Tian
- 43. Beyond Human Opinion Scores: Blind Image Quality Assessment based on Synthetic Scores, Peng Ye, Jayant Kumar, David Doermann
- 44. Active Sampling for Subjective Image Quality Assessment, Peng Ye, David Doermann
- 45. A Study on Cross-Population Age Estimation, *Guodong Guo, Chao Zhang*
- 46. Remote Heart Rate Measurement From Face Videos Under Realistic Situations, Xiaobai Li, Jie Chen, Guoying Zhao, Matti Pietikäinen
- 47. 6 Seconds of Sound and Vision: Creativity in Micro-Videos, Miriam Redi, Neil O'Hare, Rossano Schifanella, Michele Trevisiol, Alejandro Jaimes
- 48. GPS-Tag Refinement using Random Walks with an Adaptive Damping Factor, Amir Roshan Zamir, Shervin Ardeshir, Mubarak Shah

## Saturday, June 28

**0700–1700 Registration** (Exhibit Hall C Lobby)

0730-0830 Breakfast (Exhibit Hall C)

## Web-scale Vision and Social Media

Organizers: Lamberto Ballan

Alex Berg Marco Bertini Thomas Mensink Rahul Sukthankar

**Location:** C111-112 **Schedule:** Full Day

0900 Welcome

0910 Invited Talk: Marc'Aurelio Ranzato (Facebook)

og5o Photo Recall: Using the Internet to Label Your Photos, Neeraj Kumar, Steve Seitz

## 1015 Morning Break

1040 Streetscore - Predicting the Perceived Safety of One Million Streetscapes, Nikhil Naik, Jade Philipoom, Ramesh Raskar, Cesar Hidalgo

1100 A Stream Algebra for Computer Vision Pipelines, Mohamed Helala, Ken Pu, Faisal Qureshi

1120 Invited Talk: Media, Community, and the Social
Photograph, David Ayman Shamma (Yahoo Research)

#### 1200 Lunch Break (Exhibit Hall C)

1330 Invited Talk: Julian McAuley (Stanford Univ.)

1410 What is usual in unusual videos? Trajectory snippet histograms for discovering unusualness, Ahmet Iscen, Anil Armagan, Pinar Duygulu

1430 Clustering Social Event Images using Kernel Canonical Correlation Analysis, *Unaiza Ahsan* 

1450 Panel Discussion

## Computational Models for Social Interactions & Behavior: Scientific Grounding, Sensing & Applications

Organizers: Ajay Divakaran

Maneesh Singh Mohamed Amer Saad Khan Behjat Siddiquie

**Location:** C115 **Schedule:** Full Day

0910 Welcome

#### S1: Invited Talks (0915-1025)

0915 Overview of Affective Computing, Rosalind Picard (MIT)

og45 Building Blocks of Social Interaction, *Brian Lande* (SCPD)

1015 Morning Break

#### S2: Invited Talks (1040-1200)

1040 Neuroscience of Social Interactions, William Casebeer (DARPA)

1120 Neuroscientific aspects of human perception, *Ido*Davidesco (Princeton Univ.)

1045 Invited Talk: IARPA Program, Mark Burge

1200 Lunch Break (Exhibit Hall C)

## S3: Invited Talks (1300-1445)

1300 A movement based perspective on Social Interactions, *Elizabeth Torres (Rutgers Univ.)* 

1340 Multimodal Analysis of Human Behavior, *Louis-Philippe Morency (USC)* 

#### 1445 Afternoon Break

## S4: Afternoon Session (1500-1725)

1500 **Invited Talk:** Machine Learning based Temporal Models, *Graham Taylor (Univ. of Guelph)* 

1540 Oral Presentations

1625 Panel Discussion

1725 Closing Remarks

## **Long-term Detection and Tracking**

**Organizers:** Octavia Camps

Rita Cucchiara Alberto Del Bimbo

Jiri Matas

Federico Pernici Stan Sclaroff

**Location:** C210-212 **Schedule:** Full day

0905 Opening

og10 Invited Talk: Cristoph Lampert (IST Austria)
og45 Invited Talk: Peter Meier (CTO Metaio)

## 1015 Morning Break

### 1045 Invited Abstracts

- Persistent People Tracking and Face Capture Over a Wide Area, Gérard Medioni, Yinghao Cai
- Tracklet Association in Detect-then-track Paradigm for Long-term Multi-Person Tracking (Extended Abstract), Bing Wang, Gang Wang, Kap Luk Chan, Li Wang
- On Fast Trackers that are Robust to Partial Occlusions, Lu Zhang, Hamdi Dibeklioglu, Laurens van der Maaten

### 1130 Dataset Papers

- The Matrioska Tracking Algorithm on LTDT2014
   Dataset, Mario Edoardo Maresca, Alfredo Petrosino
- On-line Video Motion Estimation by Invariant Receptive Inputs, Marco Gori, Marco Lippi, Marco Maggini, Stefano Melacci

1200 Lunch Break (Exhibit Hall C)

1330 Invited Talk: Arnold Smuelders (Univ. of Amsterdam)

1400 Panel discussion

1445 Wrap-up

## **Embedded Vision**

Organizers: Goksel Dedeoglu

Fridtjof Stein Stefano Mattoccia Jagadeesh Sankaran

**Location:** C220-222 **Schedule:** Full day

o845 Opening Remarks: Goksel Dedeoglu and Fridtjof Stein, General Chairs

ogoo **Keynote**: Project Tango: Giving Mobile Devices a Human-Scale Understanding of Space and Motion, *Johnny Lee (Google)* 

1000 Brain-inspired Classroom Occupancy Monitoring on a Low-Power Mobile Platform, Francesco Conti, Antonio Pullini, Luca Benini

1020 Fast LBP Face Detection on Low-Power SIMD Architectures, Olexa Bilaniuk, Ehsan Fazl-Ersi, Robert Laganiere, Christina Xu, Daniel Laroche, Craig Moulder

1040 Invited Talk: Deep Learning Architectures, Eugenio Culurciello (Purdue Univ.)

1140 Lightning Talks: 2-minute-per-poster/demo briefs

1200 Lunch Break (Exhibit Hall C)

## 1230 Posters and Demos

#### **Posters**

- A High-Performance Hardware Architecture for a Frameless Stereo Vision Algorithm Implemented on a FPGA Platform, Florian Eibensteiner, Jürgen Kogler, Josef Scharinger
- Towards Autonomous Navigation of Miniature UAV, Roland Brockers, Martin Hummenberger, Stephan Weiss, Larry Matthies
- A Train Station Surveillance System: Challenges and Solutions, Burak Ozer, Marilyn Wolf
- Addressing System-Level Optimization with OpenVX Graphs, Erik Rainey, Jesse Villarreal, Goksel Dedeoglu, Kari Pulli, Thierry Lepley, Frank Brill
- A Compute-Efficient Algorithm for Robust Eyebrow Detection, Supriya Sathyanarayana, Ravi Kumar

## Saturday, June 28

- Satzoda, Suchitra Sathyanarayana, Srikanthan Thambipillai
- An Embedded Solution to Visual Mapping for Consumer Drones, Guyue Zhou, Ang Liu, Kang Yang, Tao Wang, Zexiang Li
- A Surround View Camera Solution for Embedded Systems, Buyue Zhang, Vikram Appia, Ibrahim Pekkucuksen, Yucheng Liu, Aziz Umit Batur, Pavan Shastry, Stanley Liu, Shiju Sivasankaran, Kedar Chitnis
- FPGA-based Fast Response Image Analysis for Autonomous or Semi-Autonomous Indoor Flight, Robert Ladiq, Kazuhiro Shimonomura
- Exploiting Traffic Scene Disparity Statistics for Stereo Vision, Stefan K. Gehriq, Uwe Franke, Nicolai Schneider
- A 240 G-ops/s Mobile Coprocessor for Deep Neural Networks, Vinayak Gokhale, Jonghoon Jin, Aysegul Dundar, Berin Martini, Eugenio Culurciello

#### Demos

- Efficient Lane and Vehicle Detection with Integrated Synergies (ELVIS), Ravi Kumar Satzoda, Mohan M. Trivedi
- A 240 G-ops/s Mobile Coprocessor for Deep Neural Networks, Vinayak Gokhale, Jonghoon Jin, Aysegul Dundar, Berin Martini, Eugenio Culurciello

1400 **Keynote Talk:** EyeTap, *Steve Mann (Univ. of Toronto)* 1500 **Invited Talk:** OpenVX, *Victor Eruhimov (ItSeez)* 

## 1540 Afternoon Break

- 1600 Gesture Recognition in Ego-Centric Videos using Dense Trajectories and Hand Segmentation, Lorenzo Baraldi, Francesco Paci, Giuseppe Serra, Luca Benini, Rita Cucchiara
- 1620 Efficient Lane and Vehicle Detection with Integrated Synergies (ELVIS), Ravi Kumar Satzoda, Mohan M. Trivedi
- 1640 Invited Talk: Embedded Vision Challenges for Implementing Augmented Reality Applications, Peter Meier (MetalO)
- 1720 Closing Remarks

# Large Scale Visual Recognition and Retrieval (Big Vision)

Organizers: Jia Deng

Alex Berg Yuanqing Lin Jason Corso

**Location:** Grand Ballroom 1

**Schedule:** Full Day ogoo Opening Remark

ogo5 **Invited Talk:** Towards a Visual Memex, *Alexei Efros* (*Univ. of California, Berkeley*)

og4o Invited Talk: Video as Training Data for Object Class Detectors, Vittorio Ferrari (Univ. of Edinburgh)

#### 1015 Morning Break

1045 Invited Talk: Towards Large-Scale Semantic Representations, Trevor Darrell and Yangqing Jia (Univ. of California, Berkeley)

1120 Spotlights and Posters

1200 Lunch Break (Exhibit Hall C)

- 1415 Invited Talk: Learning from Descriptive Text, Tamara
  Berg (Univ. of North Carolina)
- 1450 Invited Talk: The Distributed Camera, Noah Snavely (Cornell Univ.)

## 1525 Afternoon Break

- 1555 Invited Talk: Large-Scale Image Understanding, *Drago*Anguelov (Google)
- 1630 Invited Talk: Toward a Universal Perception System, Yann LeCun (New York Univ. & Facebook)

## **Egocentric Vision**

Organizers: Kris Kitani

Yong Jae Lee Michael S. Ryoo Alireza Fathi

**Location:** C113-114 **Schedule:** Full Day

0930 **Keynote Talk:** *Takeo Kanade (CMU)* 0945 **Keynote Talk:** *Chieko Asakawa (IBM)* 

## Spotlights and Poster Session A (1030-1200)

- Action and Interaction Recognition in First-Person Videos, Sanath Narayan, Mohan S. Kankanhalli, Kalpathi R. Ramakrishnan
- Video-based Object Recognition Using Novel Set-of-Sets Representations, Yang Liu, Youngkyoon Jang, Woontack Woo, Tae-Kyun Kim
- Efficient Retrieval from Large-Scale Egocentric Visual Data Using a Sparse Graph Representation, Vijay Chandrasekhar, Wu Min, Xiao Li, Cheston Tan, Bappaditya Mandal, Liyuan Li, Joo Hwee Lim
- Understanding the Nature of First-Person Videos: Characterization and Classification using Low-Level Features, Cheston Tan, Hanlin Goh, Vijay Chandrasekhar, Liyuan Li, Joo-Hwee Lim
- This Hand Is My Hand: A Probabilistic Approach to Hand Disambiguation in Egocentric Video, Stefan Lee, Sven Bambach, David J. Crandall, John M. Franchak, Chen Yu
- An Attention-based Activity Recognition for Egocentric Video, Kenji Matsuo, Kentaro Yamada, Satoshi Ueno, Sei Naito
- Temporally-Dependent Dirichlet Process Mixtures for Egocentric Video Segmentation, Joseph W. Barker, James W. Davis
- Visual Navigation Aid for the Blind in Dynamic Environments, Tung-Sing Leung, Gérard Medioni
- Wisdom of the Crowd in Egocentric Video Curation, Yedid Hoshen, Gil Ben-Artzi, Shmuel Peleg

#### 1200 Lunch Break (Exhibit Hall C)

1315 Keynote Talk: A Measure and Theory of 3D Joint Attention from First Person Cameras, Yaser Shiekh (CMU) 1400 **Keynote Talk:** Who is "Ego" in Ego-centric Vision?, *Ben Kuipers (Univ. of Michigan)* 

#### Spotlights and Poster Session B (1445-1645)

- From Ego to Nos-vision: Detecting Social Relationships in First-Person Views, Stefano Alletto, Giuseppe Serra, Simone Calderara, Francesco Solera, Rita Cucchiara
- A Sequential Classifier for Hand Detection in the Framework of Egocentric Vision, Alejandro Betancourt, Miriam M. López, Carlo S. Regazzoni, Matthias Rauterberg
- Eye-Model-Based Gaze Estimation by RGB-D Camera, Li Jianfeng, Li Shiqang
- Experiments on an RGB-D Wearable Vision System for Egocentric Activity Recognition, Mohammad Moghimi, Pablo Azagra, Luis Montesano, Ana C. Murillo, Serge Belongie
- Parsing Videos of Actions with Segmental Grammars, Hamed Pirsiavash, Deva Ramanan
- Estimating Relative Social Status from Face-to-Face Interactions using First-person Vision, Mirai Higuchi, Kris M. Kitani, Yoichi Sato
- Wearable RGB-D Navigation System For The Blind, Young Hoon Lee, Gérard Medioni
- VideoSET: Video Summary Evaluation Toolkit, Serena Yeung, Alireza Fathi, Li Fei-Fei
- First-Person Activity Recognition from Animal Videos, Yumi Iwashita, Asamichi Takamine, Ryo Kurazume, Michael S. Rvoo
- 2D Hand Parsing for Egocentric Gesture Recognition, Akanksha Saran, Kris M. Kitani
- Object Recognition in Egocentric Videos with Saliencybased Non-uniform Sampling and Variable Resolution Space for Features Selection, Vincent Buso, Jenny Benois-Pineau, Iván González-Díaz
- Gaze Estimation using Fingertip Gaze Calibration, Takeshi Saitoh
- Indoor Trajectory Estimation from Wearable Camera for Activity Monitoring, Guillaume Bourmaud, Rémi Mégret, Audrey Giremus, Yannick Berthoumieu
- Summarization of Egocentric Moving Videos for Generating Walking Route Guidance, Masaya Okamoto, Yoshiyuki Kawano, Keiji Yanai

## Saturday, June 28

- 3D Hand Pose Detection in Egocentric RGB-D Images, Grégory Rogez, Maryam Khademi, James Steven Supančič III, J. M. M. Montiel, Deva Ramanan
- 3-D Gaze Scan Path by Inside-out Camera System, Hironobu Fujiyoshi, Makoto Kimura, Shoichi Shimizu, Yuji Yamauchi, Takayoshi Yamashita
- PlaceAvoider: Steering First-Person Cameras Away from Sensitive Spaces, Robert Templeman, Mohammed Korayem, David Crandall, Apu Kapadia

1645 Closing Remarks and Award Presentations

## **Computational Cameras & Displays**

Organizers: Ashok Veeraraghavan

Oliver Cossairt Kaushik Mitra

**Location:** Exhibit Hall C

Schedule: Half Day — Morning

o830 Welcome

- 0845 Invited Talk: Integrated Imaging: Creating Images from the Tight Integration of Algorithms, Computation, and Sensors, Charles A. Bouman (Purdue Univ.)
- og3o A Novel HDR Depth Camera for Real-time 3D 360° Panoramic Vision, Ahmed Nabil Belbachir, Stephan Schraml, Manfred Mayerhofer, Michael Hofstätter
- 0945 Separating Texture and Illumination for Single-Shot Structured Light Reconstruction, Minh Vo, Srinivasa G. Narasimhan, Yaser Sheikh

1000 Morning Break: Posters & Demos Session

- 1045 Invited Talk: Time of Flight Revolution, Mohit Gupta (Columbia Univ.)
- 1130 Light Field Scale-Depth Space Transform for Dense Depth Estimation, *Ivana Tošić, Kathrin Berkner*
- Projection Center Calibration for a Co-located Projector Camera System, *Toshiyuki Amano*
- 1200 Dictionary Learning based Color Demosaicing for Plenoptic Cameras, *Xiang Huang, Oliver Cossairt*
- 1215 Best Paper Award & Concluding Remarks

## Multi-Sensor Fusion for Outdoor Dynamic Scene Understanding

Organizers: Mubarak Shah

Wolfgang Förstner Alper Yilmaz Clément Mallet Michael Ying Yang Yury Vizilter

Location: C121-122

**Schedule:** Half Day — Morning

0830 Welcome

o84o Integrating LIDAR Range Scans and Photographs with Temporal Changes, *Brittany Morago*, *Giang Bui*, *Ye Duan* 

ogoo Guided Depth Upsampling via A Cosparse Analysis Model, Xiaojin Gong, Jianqiang Ren, Baisheng Lai, Chaohua Yan, Hui Qian

0920 **Keynote Talk:** Raquel Urtasun (Univ. of Toronto)

1015 Morning Break

1045 Poster Spotlights

1120 Poster Session

- Alignment of 3D Building Models with Satellite Images Using Extended Chamfer Matching, Xi Zhang, Gady Agam, Xin Chen
- Active Planning, Sensing and Recognition Using a Resource-Constrained Discriminant POMDP, Zhaowen Wang, Zhangyang Wang, Mark Moll, Po-Sen Huang, Devin Grady, Nasser Nasrabadi, Thomas Huang, Lydia Kavraki, Mark Hasegawa-Johnson
- Frame Rate Fusion and Upsampling of EO/LIDAR Data for Multiple Platforms, T. Nathan Mundhenk, Kyungnam Kim, Yuri Owechko
- Feature Regression for Multimodal Image Analysis, Michael Ying Yang, Xuanzi Yong, Bodo Rosenhahn
- 2D/3D Sensor Exploitation and Fusion for Enhanced Object Detection, Jiejun Xu, Kyungnam Kim, Zhiqi Zhang, Hai-wen Chen, Yuri Owechko

## **Change Detection**

Organizers: Pierre-Marc Jodoin

Janusz Konrad Prakash Ishwar Fatih Porikli

Location: C121-122

Schedule: Half Day — Afternoon

1300 Opening Remarks & Description of the Challenge

1325 A Fast Self-Tuning Background Subtraction Algorithm, Bin Wang, Piotr Dudek

1350 Spectral-360: A Physics-Based Technique for Change Detection, Mohamed Sedky, Mansour Moniri, Claude C. Chibelushi

#### 1415 Break

1430 Change Detection with Weightless Neural Networks, Massimo De Gregorio, Maurizio Giordano

1455 Flexible Background Subtraction With Self-Balanced Local Sensitivity, Pierre-Luc St-Charles, Guillaume-Alexandre Bilodeau, Robert Bergevin

## 1520 Afternoon Break

1550 Static and Moving Object Detection Using Flux Tensor with Split Gaussian Models, Rui Wang, Filiz Bunyak, Guna Seetharaman, Kannappan Palaniappan

1615 Conclusion & Future Work

# Deep Vision: Deep Learning in Computer Vision

Organizers: Jose M. Alvarez

Yann LeCun Fatih Porikli

Yi Li

Location: Exhibit Hall C

Schedule: Half Day — Afternoon

1300 Opening Remarks

1305 Invited Talk: Kay Yu (Baidu)

1400 Heterogeneous Multi-task Learning for Human Pose Estimation with Deep Convolutional Neural Network, Sijin LI, Zhi-Qiang Liu, Antoni B. Chan

1415 Generalized Autoencoder: A Neural Network Framework for Dimensionality Reduction, Wei Wang, Yan Huang, Yizhou Wang, Liang Wang

1430 Unrolling Loopy Top-down Semantic Feedback in Convolutional Deep Networks, Carlo Gatta, Adriana Romero, Joost van de Veijer

1445 CNN Features Off-the-Shelf: An Astounding Baseline for Recognition, Ali Sharif Razavian, Hossein Azizpour, Josephine Sullivan, Stefan Carlsson

1500 A Piggyback Representation for Action Recognition, Lior Wolf, Yair Hanani, Tal Hassner

1515 Poster Session (with Afternoon Break)

1555 Awards

1600 Invited Talk: Pierre Sermanet (New York Univ.)

1650 Concluding Remarks

# Towards Solving Real-world Vision Problems with RGB-D Cameras

**Organizer:** Juergen Gall

Xiaofeng Ren Pushmeet Kohli

Time: 0830-1700 (Full Day)
Location: Grand Ballroom 3

**Description:** RGB-D depth cameras have the potential to become a key component for solving real-world problems. Their low cost and widespread availability have made them a commercial success and their popularity in the research community has dramatically increased. In the meanwhile, the next generation of RGB-D sensors have been developed that are better than current sensors in terms of depth quality, frame-rate, and sensor size. It is therefore a good moment to summarize what has been achieved so far and to discuss interesting possible direction for future work. The proposed short course intends to discuss the basics, underlying principles and cutting-edge results of a comprehensive list of topics in RGB-D perception.

## Visual SLAM

Organizer: Frank Dellaert

Michael Kaess

Time: 0830-1700 (Full Day)

**Location:** C213-215

**Description:** This tutorial addresses Visual SLAM, the problem of building a sparse or dense 3D model of the scene while traveling through it, and simultaneously recovering the trajectory of the platform/camera. Visual SLAM has received much attention in the computer vision community in the last few years, as more challenging data sets become available, and visual SLAM is starting to be implemented on mobile cameras and used in AR and other applications. We will provide an introduction to the core concepts underlying current sparse, dense and semantic visual SLAM systems.

## **Large-Scale Visual Recognition**

Organizer: Zaid Harchaoui

Hervé Jégou Florent Perronnin

Time: 0830-1700 (Full Day)
Location: Grand Ballroom 2

**Description:** This tutorial addresses the topic of Large-Scale Visual Recognition (LSVR), the problem of understanding visual content (e.g. photos or videos) on a large-scale. This topic has received much attention in the computer vision community in the last few years: as larger datasets have become available, handling millions of images and thousands of label classes has become the norm rather than the exception. Since LSVR is a broad topic, we will mainly focus on two tasks: image retrieval and image classification.

The goals of this tutorial are four-fold:

- Provide the audience with the key "tools" to process such large datasets.
- Review the main families of features for visual recognition, from bag-of-visual-words to deep convolutional features, including VLAD and Fisher Vectors
- Show the convergence between large-scale retrieval and large-scale classification, two problems which have been traditionally addressed separately.
- Show that LSVR does not necessarily require massive computational resources (although such resources can help, of course...)

## Describing Images in Natural Language

Organizer: Julia Hockenmaier

**Time:** 0830-1230 (Half Day — Morning)

Location: C216

**Description:** The ability to associate images with natural language sentences that describe what is depicted in them is a hallmark of image understanding, and a prerequisite for applications such as sentence-based image search. The purpose of this tutorial is to give researchers in computer vision an overview of the issues involved in automatic image description, and to introduce them to natural language processing tools and ideas they can use for this purpose.

# Learning and Inference in Discrete Graphical Models

Organizer: Nikos Komodakis

Nikos Paragios Dhruv Batra Stephen Gould

Time: 1300-1700 (Half Day — Afternoon)

Location: C216

**Description:** Several problems in computer vision can be formulated using the discrete graphical models framework. The two main issues faced by researchers in this case are: (i) Learning: How to estimate the parameters of the model?; and (ii) Inference: How to find the best assignment for the variables of the model? In this tutorial we will discuss these two issues, starting from the basics and building up to the state of the art.

## **Founding a Computer Vision Startup**

Organizer: Till Quack

Ian Erik Solem

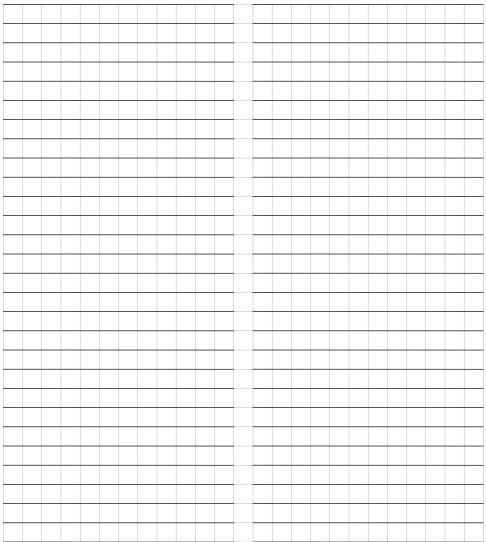
**Time:** 0830-1230 (Half Day — Morning)

**Location:** C123-125

Description: A few years ago Jan Erik and Till both founded startups in the computer vision field: Polar Rose and Kooaba, respectively. Both companies had their successful exits with a global player, and Jan Erik even went on to found his next company. Along the - sometimes bumpy - journey we learned a great deal of things which we would like to share with the computer vision community. We found that creating your own startup is a very fulfilling experience and an interesting career choice after a completed PhD. In addition, the time is right to create more startups in the vision field, where technology is now in a cycle where many novel (consumer) applications become feasible.

We gave a version of this course at CVPR 2010. Much has changed since, but the goal of this course remains to give a rather broad overview of practices and tools which turned out to be useful for us. Now refreshed and improved with additional learnings.

The course spans diverse topics as Funding, Hiring, Software Engineering, Business Models, Product Design. We present material for each topic, together with pointers and links to relevant resources. In this new edition of the tutorial we will also focus on conversations between Jan Erik and Till on the different topics.



# Silver Donors



# **SIEMENS**

# Honeywell





















# **Bronze Donors**





























# Platinum Donors























**Gold Donors** 

