



IEEE Computer Society  
Conference on

# Computer Vision and Pattern Recognition

Pocket Guide



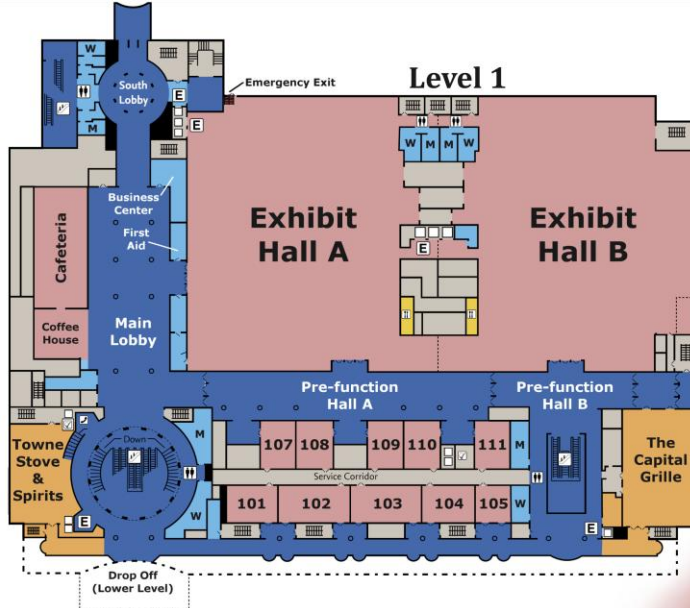
**CVPR**

June 7-12, 2015

**Boston, MA**



# Hynes Convention Center



**CVPR 2015 Sponsors:**





# Message from the General and Program Chairs

Welcome to Boston, Massachusetts and the 28<sup>th</sup> IEEE Conference on Computer Vision and Pattern Recognition (CVPR). In addition to the main three-day program of oral and poster presentations, plenary talks, demos, exhibitions, and social functions, CVPR 2015 has a number of co-located events, including 31 workshops and 18 tutorials. The conference includes 602 published papers (28.4% of valid submissions). 71 of these were allocated oral presentations (3.3% of valid submissions). All papers will include a poster presentation.

This year, we received a record 2123 valid submissions to the main conference, of which 1815 were fully reviewed (the others were either administratively rejected for technical or ethical reasons or withdrawn before review).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 66 researchers to act as Area Chairs (ACs). ACs were selected to provide a broad range of expertise, to balance junior and senior members, and to represent a variety of geographical locations.

We recruited a team of 1238 experienced reviewers from the broader computer vision and pattern recognition community. The original list of reviewers was augmented with reviewers recommended by the ACs to add expertise for papers where appropriate reviewers were not initially available.

The review process was similar to previous years, with a few salient differences. The Program Chairs made a special effort this year to assign papers on similar topics to the same set of ACs, so that related papers could be compared directly. To achieve this, papers were manually assigned to each AC, with the help of the Toronto Paper Matching System and while accounting for the author-provided suggestions. Each paper was reviewed by at least three reviewers and considered by at least two ACs before a decision was made. Borderline papers and candidate orals were discussed in groups of six non-conflicted ACs with common areas of expertise. Oral recommendations were made by panels of 12 or 18 ACs after extensive discussion.

The Program Chairs did not submit any papers to CVPR 2015, allowing them to work without any direct

conflicts throughout the review process. This year, General Chairs were allowed to submit papers, and therefore did not have any software access to the CMT system beyond that of an author. The General Chairs were consulted in a few special cases to evaluate double submissions and appeals, after conflicts of interest were taken into account. Additionally, 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 602 accepted papers will be available approximately two weeks before the conference on the Computer Vision Foundation web site. In addition, the proceedings of CVPR 2015 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.

Following a successful practice at the British Machine Vision Conference, the organizers asked authors to provide a one-page extended abstract for each paper. These abstracts will be collated by poster session and provided on the main CVPR web site.

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 2015 a first-rate conference. Our thanks also go to the organizers of previous CVPRs for their helpful advice and support. A few people working behind the scenes deserve special thanks. Cheni Chadowitz and Joseph Lim provided critical technical assistance with CMT and managing the largest CVPR ever. It could not have been a success without them. Mike Jones, Nicole Finn, Terry Boulton, and Ginger Boulton all provided critical support to the general conference organization and financial well-being of CVPR. Gary Huang has provided high-quality and rapid turnaround of the CVPR web pages. Eric Mortensen and Sanja Fidler have again lent expertise developed over many years to shepherd the papers through the publi-



## Message from the General and Program Chairs

cation process. We are grateful to the sponsors as well, and we are happy to report that CVPR 2015 has seen another record-breaking year of 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.

*Kristen Grauman, Erik Learned-Miller,  
Antonio Torralba, Andrew Zisserman  
Program Co-Chairs*

*Horst Bischof, David Forsyth,  
Cordelia Schmid, Stan Sclaroff  
General Co-Chairs*

## CVPR 2015 Organizing Committee

**General Chairs:** Horst Bischof  
David Forsyth  
Cordelia Schmid  
Stan Sclaroff

**Program Chairs:** Kristen Grauman  
Erik Learned-Miller  
Antonio Torralba  
Andrew Zisserman

**Workshops Chairs:** Alex Berg  
Tamara Berg

**Tutorials Chairs:** Deva Ramanan  
Octavia Camps

**Finance Chairs:** Terry Boulton  
Greg Mori

**Publications Chairs:** Eric Mortensen  
Sanja Fidler

**Demos/Exhibitions Chairs:** Kate Saenko  
Bruce A. Maxwell

**Doctoral Consortium Chairs:** Adriana Kovashka  
Christoph Lampert

**Website Chair:** Gary B. Huang

**Local Arrangements Chair:** Mike Jones

**Corporate Relations Chair:** Fatih Porikli

**Publicity Chair:** Kristin Dana

**Student Activities Chairs:** Devi Parikh  
James Hays

**Video Proceedings Chair:** Nathan Jacobs

**Logistics Advisor:** Ginger Boulton



## CVPR 2015 Area Chairs

Pablo Arbelaez  
Shai Avidan  
Serge Belongie  
Terry Boulton  
Yuri Boykov  
Shih-Fu Chang  
Piotr Dollar  
Fredo Durand  
Alyosha Efros  
Pedro Felzenszwalb  
Rob Fergus  
Rogerio Feris  
Vitorio Ferrari  
Abhinav Gupta  
Fred Hamprecht  
James Hays  
Martial Hebert

Derek Hoiem  
Gang Hua  
Michal Irani  
David Jacobs  
Herve Jegou  
Neel Joshi  
Frederic Jurie  
Iasonas Kokkinos  
Vladimir Kolmogorov  
M. Pawan Kumar  
Ivan Laptev  
Svetlana Lazebnik  
Kyoung Mu Lee  
Bastian Leibe  
Victor Lempitsky  
Vincent Lepetit  
Hongdong Li

Simon Lucey  
Jiri Matas  
Yasuyuki Matsushita  
Krystian Mikolajczyk  
Srinivasa Narasimhan  
Aude Oliva  
Devi Parikh  
Florent Perronnin  
Robert Pless  
Jean Ponce  
Deva Ramanan  
Marc Aurelio Ranzato  
James Rehg  
Xiaofeng Ren  
Arun Ross  
Ruslan Salakhutdinov  
Silvio Savarese

Eli Shechtman  
Yaser Sheikh  
Cristian Sminchisescu  
Noah Snavely  
Erik Sudderth  
Jian Sun  
Lorenzo Torresani  
Raquel Urtasun  
Luc Van Gool  
Andrea Vedaldi  
Olga Veksler  
Jakob Verbeek  
Lihi Zelnik-Manor  
S. Kevin Zhou  
Todd Zickler  
Larry Zitnick

## CVPR 2015 Outstanding Reviewers

We are pleased to recognize the following researchers as "Outstanding Reviewers for CVPR 2015". These reviewers were selected from more than 1200 reviewers for their hard work in providing detailed reviews

for the papers assigned to them. These reviewers were identified by one or more of the CVPR Area Chairs, who found their reviews of high quality.

Tal Arbel  
Dan Banica  
Adrien Bartoli  
Dhruv Batra  
Ismail Ben Ayed  
Rodrigo Benenson  
Ohad Ben-Shahar  
Joan Bruna  
Ayan Chakrabarti  
David Crandall  
Francisco Estrada  
Bin Fan  
Sanja Fidler  
Alexander Fix  
Boris Flach  
Francois Fleuret  
Wolfgang Forstner  
Friedrich Fraundorfer  
Adrien Gaidon

Ross Girshick  
Albert Gordo  
Bharath Hariharan  
Tal Hassner  
Soren Hauberg  
Kaiming He  
Adrian Hilton  
Jeremy Jancsary  
C. V. Jawahar  
Joerg Kappes  
Andrei Karpathy  
Piotr Koniusz  
Adriana Kovashka  
Jean-Francois Lalonde  
Yong Jae Lee  
Stephen Lin  
Michael Maire  
Subhansu Maji  
Tim Marks

Stefan Mathe  
Richard Newcombe  
Takayuki Okatani  
Chris Pal  
Vishal Patel  
Marco Pedersoli  
Leonid Pishchulin  
Jerry Prince  
Ian Reid  
Hayko Riemenschneider  
Emanuele Rodola  
Jose Rodriguez-Serrano  
Marcus Rohrbach  
Stefan Roth  
Olga Russakovsky  
Torsten Sattler  
Bogdan Savchynskyy  
Walter Scheirer  
Falk Schubert

Samuel Schulter  
Pradeep Sen  
Lavanya Sharan  
Alexander Shekhovtsov  
Abhinav Shrivastava  
Cees Snoek  
Anuj Srivastava  
Michael Stark  
Joost van de Weijer  
Jan van Gemert  
Rene Vidal  
Carl Vondrick  
Lei Wang  
David Wipf  
Paul Wohlhart  
Dijia Wu  
Stefanos Zafeiriou



## Sunday, June 7

---

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

**0730–0830 Breakfast** (Exhibit Hall B)

**1200–1400 Lunch** (Exhibit Hall B)

### Search and Planning for Inference and Learning in Computer Vision

**Organizer:** Iasonas Kokkinos  
Sinisa Todorovic  
Tianfu Wu

**Time:** 0830-1800 (Full Day)

**Location:** Room 102

**Description:** Visual understanding requires expressive representations, such as graphical models and structured predictors to account for context, compositionality, and the spatio-temporal structure of objects, activities, and events. However, brute-force inference, or learning, with such representations is intractable and requires organizing computation in a more efficient manner.

This long-standing problem of computer vision underlies an interrelated set of traditional methodologies such as top-down / bottom-up message passing, branch-and-bound, A\* search, sequential and cascade-based inference and learning, (inverse) reinforcement learning, scheduling of visual attention, visual routines and Gestalt-based grouping. In the past few years, new approaches have resurged as search- and planning-based formulations for solving complex, high-order structured prediction problems, including Generalized A\*, SEARN, HC-Search, and Monte Carlo Tree Search and have aroused increased interest in AI and machine learning.

SPIL aims at (i) reviewing the aforementioned search-, sequential-, cascade-, and message-passing-based approaches to inference and learning (ii) presenting the latest advances in

search- and planning-based formulations of complex structured prediction problems (iii) establishing connections between the recent advances in machine learning and artificial intelligence with early and existing work in computer vision (iv) providing extensive demonstrations of the algorithms presented with applications in object recognition and activity recognition, drawing from recent research activities of the organizers and other researchers.

### Energy Minimization and Discrete Optimization

**Organizer:** Stefanie Jegelka  
Lubor Ladicky  
Srikumar Ramlingam  
Christopher Russell

**Time:** 0830-1800 (Full Day)

**Location:** Room 103

**Description:** The connection between MRF problems in computer vision community and general discrete optimization concepts such as submodularity allowed us to study the properties of optimality of the existing algorithms and also enabled us to develop more principled discrete optimization algorithms. As more and more problems in vision community are formulated as energy minimization problems, there is also another side to the whole story. The problem formulation using maximization of submodular functions allows us to develop approximation algorithms for many important computer vision tasks with provable bounds. Further understanding of matroids and submodularity provide us a general class of maximization algorithms with a rich set of constraints that can be denoted by matroids. Numerous computer vision and machine learning problems are inherently discrete. Often, the resulting discrete optimization problems are computationally extremely challenging. While convexity is an important property when solving continuous optimization problems, submodularity, also viewed as a discrete analog of convexity, is



closely tied to tractability in the discrete world: its structure is key to solving many discrete optimization problems. Even more, the characterizing property of submodular functions, diminishing marginal returns, appears naturally in various settings and is a rich abstraction for a myriad of problems. The aim of the tutorial will be to give a basic introduction, to give a structured overview of recent results in this active research area (from within and beyond the computer vision and machine learning community), to provide intuition for which types of problems may profit from those results, and to offer pointers to material for further study.

## Open Source Structure-from-Motion

**Organizer:** Matt Leotta  
Sameer Agarwal  
Frank Dellaert  
Pierre Moulon  
Vincent Rabaud

**Time:** 0830-1800 (Full Day)

**Location:** Room 206

**Description:** This tutorial is a hands-on introduction to using open source software for solving structure-from-motion (SfM) estimation problems from imagery and video. The goal is to teach attendees how to combine existing open source libraries to build complete, customized SfM estimation pipelines. Topics include feature detection, feature matching/tracking, projective geometry, camera calibration, factor graphs, sparse bundle adjustment, etc. and how to combine all of these components for applications in processing video or unordered image collections. Example code for complete processing pipelines and sample data will be provided, and attendees will complete hands on exercises to modify, build, and run the code. The focus is on C++ libraries that are permissively licensed (BSD, Apache, MPL, etc.) to allow for use in commercial applications. Primary software packages covered are OpenCV, Ceres Solver, GTSAM, OpenMVG, and MAP-Tk.

## 3-D Vision with Robotics Applications in MATLAB

**Organizers:** Dima Lisin  
Witek Jachimczyk  
Zhen Wu  
Avi Nehemiah

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

**Location:** Room 101

**Description:** 3-D computer vision has rapidly evolved over the last decade and is now being widely used in fields such as autonomous robotics, automotive safety, gesture recognition and gaming. This half day tutorial will cover both theoretical and practical information on implementing 3D vision systems using passive sensors such as calibrated single and stereo cameras, and active RGB-D sensors such as the Microsoft Kinect. We will also show practical applications of 3D vision to robotics.

## Contactless 3D Fingerprint Identification

**Organizers:** Ajay Kumar

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

**Location:** Room 104

**Description:** Fingerprint identification is commonly used for the civilian and law-enforcement applications around the world. Traditional acquisition of fingerprint scans by pressing or rolling of finger against the hard surface like glass/polymer or paper often results in partial or degraded images due to skin deformation, smearing or due to sensor noise. As a result full potential from the fingerprints is not realized. Therefore, touchless 3D fingerprint systems have emerged to provide ideal solutions to above intrinsic problems. Such 3D approaches can also provide more accurate personal identification as rich information is available from 3D fingerprint images. Emerging solutions for the contactless 3D



fingerprint acquisition are largely based on shape from silhouette, structured lighting or photometric stereo based imaging. However, widely accepted standards or the representation of 3D fingerprint features is yet to emerge. The minutiae features are widely considered to be most reliable and widely employed by law enforcement experts and commercial 2D fingerprint systems available today. Therefore 3D fingerprint templates using minutiae representation extended in 3D space is expected to have wide acceptance and interoperability with the conventional 2D fingerprint templates. Accurate recovery, representation, selection, registration and matching of 3D fingerprints is still an open research problem with significant potential impacting our daily lives. This tutorial will provide algorithmic details and tutorial overview relating to contactless 3D fingerprint identification. Theoretical models to ascertain the uniqueness of 3D fingerprints can provide upper bound, on the expected improvement over matching accuracy using 2D fingerprints, and will also be covered in this tutorial.

## Computer Vision for Visual Effects

**Organizers:** Rich Radke

**Time:** 0830-1200 (Half Day — Morning)

**Location:** Room 105

**Description:** Modern blockbuster movies seamlessly introduce impossible characters and action into real-world settings using digital visual effects, which are largely made possible by research from computer vision. This tutorial will educate students, engineers, and researchers about the fundamental computer vision principles and state-of-the-art algorithms used to create cutting-edge visual effects for movies and television. The tutorial will begin with a general overview of computer vision algorithms used on a regular basis in Hollywood (such as blue-screen matting, structure from motion, optical flow, and feature tracking) as well as recent developments that form the basis for future effects (such as natural

image matting, multi-image compositing, view synthesis, and multiview stereo). We will also discuss the technologies behind motion capture and three-dimensional data acquisition. Next, several computer vision experts and visual effects artists will discuss some of their recent contributions to the VFX industry in depth. The tutorial will conclude with a panel discussion about the challenges of incorporating CVPR-level vision research into the real-world visual effects pipeline, which will include time for questions from the audience.

## Applied Deep Learning for Computer Vision with Torch

**Organizers:** Koray Kavukcuoglu  
Ronan Collobert  
Soumith Chintala

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

**Location:** Room 200

**Description:** This is a hands-on tutorial for Torch: a general purpose scientific computing platform that has great toolboxes for deep learning, computer vision and optimization, and fast multi-machine and multi-GPU backends. This tutorial is targeted at computer vision researchers who would like to use deep learning techniques in their research for classification, segmentation, transfer learning, and other common vision tasks. Torch is supported by Facebook, Google, Twitter and several smaller companies who actively open-source their code and packages.



## Group Behavior Analysis and Its Applications

**Organizers:** Hyun Soo Park  
Wongun Choi

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

**Location:** Room 201

**Description:** In social scenes, people form multiple groups and their behaviors are governed by mutual interactions. Their behaviors are often strongly correlated, e.g., they pay attention to the same point of interest, walk together in a crowd, or make reciprocal gestures during dyadic interactions. In this tutorial, we will review the existing group behavior analyses in computer vision. This will provide a complete overview of the area through three fundamentals: (1) computational representations of social signals in terms of gaze, poses, and gestures; (2) predictive models that encode the relationship between the social signals, and (3) various applications of such social scene understanding to vision, graphics, and robotics. The tutorial will also review social interactions datasets that allow us to empirically rediscover theories of social signals in psychology and sociology such as joint attention, *F*-formation, and proxemics.

**Invited Speaker:** Capturing Subtle Social Behaviors in the Panoptic Studio, *Yaser Sheikh (CMU)*

## OpenCV 3.0 Technical Tutorials: Beginner to Specialist

**Organizers:** Gary Bradski  
Vadim Pisarevsky  
Vincent Rabaud  
Grace Vesom

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

**Location:** Room 202

**Description:** We will overview the new OpenCV 3.0 Release. What's changed, what's improved, what's added. We will then step you through hands-on with a series of exercises from

simple "Hello world" getting started, to Android and iOS, calibration, turning pictures into calibration objects, stereo, line segment tracking, computational photography, Python, text detection, ARUCO marker reading, 3D object recognition, Random Forests and other advanced algorithms. We'll end with a "What's next" including connection to the main deep learning libraries, projection mapping, etc.

## ImageNet Large Scale Visual Recognition Challenge Tutorial

**Organizers:** Olga Russakovsky  
Jonathan Krause  
Karen Simonyan  
Yangqing Jia  
Jia Deng  
Alex Berg  
Fei-Fei Li

**Time:** 0830-1215 (Half Day — Morning)

**Location:** Room 203

**Description:** The ImageNet Large Scale Visual Recognition Challenge (ILSVRC) is a benchmark in object category classification and detection on hundreds of object categories and millions of images. In contrast to the annual ILSVRC workshop at ECCV/ICCV, the goal of this tutorial is to share some of the lessons learned by the organizers and participants in the past five years. The tutorial aims to (1) present the attendees with a historical background of the challenge, which would help them to place the latest research innovations in perspective, (2) familiarize the attendees with the process of large-scale crowdsourced data collection, enabling them to better understand the ImageNet dataset as well as collect their own large-scale datasets, and (3) provide the attendees with tools and techniques which would allow them to enter the next ImageNet challenge. The target audience is computer vision students and researchers (both from academia and industry) interested in better understanding the ImageNet challenge history, collecting their own large-scale datasets, or competing in ILSVRC2015.



## Fast Image Processing With Halide

**Organizer:** Jonathan Ragan-Kelley  
Andrew Adams  
Fredo Durand

**Time:** 1400-1800 (Half Day — Afternoon)

**Location:** Room 101

**Description:** Halide is a new programming language for image processing, applicable to many vision algorithms and applications, and now in widespread use at companies and research labs. This tutorial will introduce new programmers from the vision community to the core concepts in Halide, and how to use Halide to productively write high-performance image processing code. The tutorial will first introduce the major concepts in Halide, and then dive into them in more detail by example. We will bring users from setting up their first Halide environment, through writing and compiling basic programs, to implementing and optimizing a nontrivial image processing pipeline, all in a few hours and very few lines of code.

## OpenVX: A Framework for Accelerating Computer Vision

**Organizer:** Victor Erukhimov  
Kari Pulli  
Thierry Lepey

**Time:** 1400-1800 (Half Day — Afternoon)

**Location:** Room 104

**Description:** We will introduce the audience to OpenVX 1.0: a hardware abstraction layer for computer vision. OpenVX is an open standard released by the Khronos Group in Q4 2014. OpenVX enables performance and power-optimized computer vision functionality processing, especially important in embedded and real-time use cases. The first half of tutorial will cover the API, including the individual computer vision functions as well as a graph API that enables the OpenVX developer to efficiently run computer vision algorithms on

heterogeneous architectures. We will cover the relationship between OpenVX and OpenCV, as well as OpenCL. The second half of the tutorial will be a practice session, dedicated to solving a computer vision problem with OpenVX.

## Applied Math as Applied in Cinema

**Organizer:** Marcelo Bertalmio

**Time:** 1400-1800 (Half Day — Afternoon)

**Location:** Room 105

**Description:** This tutorial covers a wide range of topics showing how applied math techniques from image processing and computer vision have become ubiquitous in movie-making, from shooting to exhibition. It does not deal with visual effects or computer-generated images, but rather with all the ways in which applied math is used to enhance, restore, adapt or convert moving images, the purpose of these techniques being to make the images look as good as possible while exploiting all the capabilities of cameras, projectors and displays.

## DIY Deep Learning: A Hands-On Tutorial With Caffe

**Organizer:** Evan Shelhamer  
Jeff Donahue  
Yangqing Jia  
Jonathan Long  
Ross Girshick

**Time:** 1400-1800 (Half Day — Afternoon)

**Location:** Room 200

**Description:** This tutorial is designed to equip researchers and developers with the tools and know-how needed to incorporate deep learning into their work. Both the ideas and implementation of state-of-the-art deep learning models will be presented. While deep learning and deep features have



recently achieved strong results in many tasks, a common framework and shared models are needed to advance further research and applications and reduce the barrier to entry. To this end we present the Caffe framework that offers an open-source library, public reference models, and working examples for deep learning. Join our tour from the 1989 LeNet for digit recognition to today's top ILSVRC14 vision models and beyond to detection, vision + language, and segmentation models. Follow along with do-it-yourself code notebooks. While focusing on vision, general techniques are covered.

## Distance Metric Learning for Visual Recognition

**Organizer:** Jiwen Lu  
Ruiping Wang  
Wei-Shi Zheng  
Weihong Deng

**Time:** 1400-1800 (Half Day — Afternoon)

**Location:** Room 201

**Description:** Over the past decade, distance metric learning has been developed as one of the basic techniques for visual recognition and successfully applied to a wide range of vision tasks showing state-of-the-art performance. In this tutorial, we will overview the trend of distance metric learning techniques and discuss how they are employed to boost visual recognition performance. First, we briefly introduce the basic concept of distance metric learning, and show the key advantages and disadvantages of existing distance metric learning methods in different visual recognition tasks. Second, we introduce some of our newly proposed distance metric learning methods from three important visual recognition applications: face recognition, person re-identification & activity, and image set classification, respectively. Lastly, we will discuss some open problems in distance metric learning to show how to further develop more advanced metric learning algorithms for visual recognition in the future.

## 3D Indoor Scene Understanding

**Organizer:** Sanja Fidler  
Raquel Urtasun

**Time:** 1400-1800 (Half Day — Afternoon)

**Location:** Room 202

**Description:** This tutorial aims to provide the necessary background for understanding different indoor problems, their difficulties, the different sensors and data sources that one can use, how to exploit them, as well as how to formulate the problems such that efficient learning and inference is possible. In particular, we will review models for 3D object detection in both, mono-nocular and RGB-D imagery, semantic segmentation (class and instance level), room layout estimation, as well as depth and normal estimation from a single image. The tutorial will discuss different possible parameterizations, sensors, the use of generative vs discriminative models, learned vs hand-crafted features, use of extra-information (e.g., CAD models, furniture catalogues), possible inference techniques, and possible learning algorithms. This will allow the audience to get a bigger picture of what is happening in the field, and what is potential-ly missing.



## Sparse and Low-Rank Modeling for High-Dimensional Data Analysis

**Organizer:** Guillermo Sapiro  
Rene Vidal  
John Wright  
Ehsan Elhamifar

**Time:** 1400-1800 (Half Day — Afternoon)

**Location:** Room 203

**Description:** Description: The increasing amounts of high-dimensional data in computer vision and other science and engineering fields, requires robust tools and techniques for recovering the un-derlying low-dimensional structures in the data. While sub-space methods have been studied intensively (from PCA and beyond), new results have recently emerged from the areas of sparse coding, matrix factorization, matrix completion, sub-space clustering and manifold clustering. This tutorial will present the audience with a unifying perspective of the theo-retical foundations behind these methods together with its applications in computer vision including motion segmenta-tion, face recognition, registration, active learning, video summarization, and more.

### Syllabus:

- 1) Robust Principal Component Analysis and Learning Low-Rank Models
- 2) Subspace Clustering and Subset Selection via Sparse Representation.

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

**Organizer:** Torsten Sattler  
Akihiko Torii

**Time:** 1400-1800 (Half Day — Afternoon)

**Location:** Room 204

**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 query 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, enable place recognition systems to handle the repetitive structures prevalent in urban environments, learn distinct representation of different places, and synthesize images in order to overcome larger changes in viewpoint. 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 prioritized matching schemes that enable state-of-the-art localization systems to efficiently handle 3D models consisting of millions of 3D points. Next, we discuss methods for scalable localization based on advanced pose estimation techniques, compressed models, and place recognition systems. Finally, we address the issue of real-time localization on mobile devices.



## Monday, June 8

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

**0730–0830 Breakfast** (Exhibit Hall B)

**0830–0835 Welcome by the General Chairs**  
(Ballrooms A-B)

**0840–1010 Oral Mon-AM-1: CNN Architectures**  
(Ballrooms A-B)

**Papers in this session are also in Poster Session Mon-AM.**

**Chairs:** Rob Fergus (*New York Univ.*)  
Florent Perronnin (*Xerox Research*)

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

1. Hypercolumns for Object Segmentation and Fine-Grained Localization, *Bharath Hariharan, Pablo Arbeláez, Ross Girshick, Jitendra Malik*
2. Modeling Local and Global Deformations in Deep Learning: Epitomic Convolution, Multiple Instance Learning, and Sliding Window Detection, *George Papandreou, Iasonas Kokkinos, Pierre-André Savalle*
3. Improving Object Detection With Deep Convolutional Networks via Bayesian Optimization and Structured Prediction, *Yuting Zhang, Kihyuk Sohn, Ruben Villegas, Gang Pan, Honglak Lee*
4. Going Deeper With Convolutions, Christian Szegedy, *Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, Andrew Rabinovich*
5. Understanding Image Representations by Measuring Their Equivariance and Equivalence, *Karel Lenc, Andrea Vedaldi*
6. Deep Neural Networks Are Easily Fooled: High Confidence Predictions for Unrecognizable Images, *Anh Nguyen, Jason Yosinski, Jeff Clune*

**0840–1010 Oral Mon-AM-2: Depth/3D/Surfaces**  
(Rooms 302-306)

**Papers in this session are also in Poster Session Mon-AM.**

**Chairs:** Yaser Sheikh (*Carnegie Mellon Univ.*)  
TBA

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

1. DynamicFusion: Reconstruction and Tracking of Non-Rigid Scenes in Real-Time, *Richard A. Newcombe, Dieter Fox, Steven M. Seitz*
2. 3D Scanning Deformable Objects With a Single RGBD Sensor, *Mingsong Dou, Jonathan Taylor, Henry Fuchs, Andrew Fitzgibbon, Shahram Izadi*
3. An Efficient Volumetric Framework for Shape Tracking, *Benjamin Allain, Jean-Sébastien Franco, Edmond Boyer*
4. Part-Based Modelling of Compound Scenes From Images, *Anton van den Hengel, Chris Russell, Anthony Dick, John Bastian, Daniel Pooley, Lachlan Fleming, Lourdes Agapito*
5. SUN RGB-D: A RGB-D Scene Understanding Benchmark Suite, *Shuran Song, Samuel P. Lichtenberg, Jianxiang Xiao*
6. Small-Variance Nonparametric Clustering on the Hypersphere, *Julian Straub, Trevor Campbell, Jonathan P. How, John W. Fisher III*

**1010–1045 Break** (Exhibit Hall B)

**1010–1230 Exhibits** (Exhibit Hall A)

- Megvii
- Google
- Ag
- Intel
- Amazon
- Facebook
- Metiao
- Orbeus
- Minieye
- Itseez
- Curalate
- Wolfram
- Kitware
- URC Ventures
- Microsoft
- Cognex
- NVIDIA
- DAQRI
- Cogtu
- Avigilon
- MERL
- iRobot
- CogniVue
- Snapchat
- MathWorks
- Omron
- Disney Research
- Vision Systems Inc.



- 4D View Solutions
- LensBricks
- Lucidyne
- KAUST
- NREC
- Rethink Robotics
- 3dMD
- FacioMetrics
- Second Spectrum
- FastVis
- Morgan & Claypool
- CRC press
- Springer
- Elsevier
- Point Grey
- Body Labs
- Datatang Technology
- now publishers

## 1010-1230 Demos (Room 108)

- Energy Efficient Structured Light Imaging, *Supreeth Achar, Matthew O'Toole, Srinivasa G. Narasimhan, Kiriakos N. Kutlakos (CMU; Univ. of Toronto)*
- Privacy Preserving Vision Sensors, *Francesco Pittaluga, Sanjeev J. Koppal (Univ. of Florida)*
- RFLOW: User Interaction Beyond Walls, *Hisham Bedri, Yonkrist Gupta, Micha Feigin, Andrew Temme, Greg Charvat, Ramesh Raskar (MIT Media Lab; Michigan State University)*

## 1010-1230 Poster Mon-AM (Exhibit Hall A)

1. Going Deeper With Convolutions, *Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, Andrew Rabinovich*
2. Propagated Image Filtering, *Jen-Hao Rick Chang, Yu-Chiang Frank Wang*
3. Web Scale Photo Hash Clustering on A Single Machine, *Yunchao Gong, Marcin Pawlowski, Fei Yang, Louis Brandy, Lubomir Bourdev, Rob Fergus*
4. Expanding Object Detector's HORIZON: Incremental Learning Framework for Object Detection in Videos, *Alina Kuznetsova, Sung Ju Hwang, Bodo Rosenhahn, Leonid Sigal*
5. Supervised Discrete Hashing, *Fumin Shen, Chunhua Shen, Wei Liu, Heng Tao Shen*
6. What do 15,000 Object Categories Tell Us About Classifying and Localizing Actions?, *Mihir Jain, Jan C. van Gemert, Cees G. M. Snoek*

7. Landmarks-Based Kernelized Subspace Alignment for Unsupervised Domain Adaptation, *Rahaf Aljundi, Rémi Emonet, Damien Mueselet, Marc Sebban*
8. Blur Kernel Estimation Using Normalized Color-Line Prior, *Wei-Sheng Lai, Jian-Jiun Ding, Yen-Yu Lin, Yung-Yu Chuang*
9. A Light Transport Model for Mitigating Multipath Interference in Time-of-Flight Sensors, *Nikhil Naik, Achuta Kadambi, Christoph Rhemann, Shahram Izadi, Ramesh Raskar, Sing Bing Kang*
10. Traditional Saliency Reloaded: A Good Old Model in New Shape, *Simone Frintrap, Thomas Werner, Germán Martín García*
11. Automatic Construction Of Robust Spherical Harmonic Subspaces, *Patrick Snape, Yannis Panagakis, Stefanos Zafeiriou*
12. Leveraging Stereo Matching With Learning-Based Confidence Measures, *Min-Gyu Park, Kuk-Jin Yoon*
13. Saliency Detection via Cellular Automata, *Yao Qin, Huchuan Lu, Yiqun Xu, He Wang*
14. Efficient Sparse-to-Dense Optical Flow Estimation Using a Learned Basis and Layers, *Jonas Wulff, Michael J. Black*
15. Learning Multiple Visual Tasks While Discovering Their Structure, *Carlo Ciliberto, Lorenzo Rosasco, Silvia Villa*
16. Projection Metric Learning on Grassmann Manifold With Application to Video Based Face Recognition, *Zhiwu Huang, Ruiping Wang, Shiguang Shan, Xilin Chen*
17. Structural Sparse Tracking, *Tianzhu Zhang, Si Liu, Changsheng Xu, Shuicheng Yan, Bernard Ghanem, Narendra Ahuja, Ming-Hsuan Yang*
18. Data-Driven Depth Map Refinement via Multi-Scale Sparse Representation, *HyeokHyen Kwon, Yu-Wing Tai, Stephen Lin*
19. Uncalibrated Photometric Stereo Based on Elevation Angle Recovery From BRDF Symmetry of Isotropic Materials, *Feng Lu, Imari Sato, Yoichi Sato*
20. Attributes and Categories for Generic Instance Search From One Example, *Ran Tao, Arnold W.M. Smeulders, Shih-Fu Chang*
21. Heat Diffusion Over Weighted Manifolds: A New Descriptor for Textured 3D Non-Rigid Shapes, *Mostafa Abdelrahman, Aly Farag, David Swanson, Moumen T. El-Melegy*



22. A Dynamic Programming Approach for Fast and Robust Object Pose Recognition From Range Images, *Christopher Zach, Adrian Penate-Sanchez, Minh-Tri Pham*
23. Beyond Gaussian Pyramid: Multi-Skip Feature Stacking for Action Recognition, *Zhengzhong Lan, Ming Lin, Xuanchong Li, Alex G. Hauptmann, Bhiksha Raj*
24. A Geodesic-Preserving Method for Image Warping, *Dongping Li, Kaiming He, Jian Sun, Kun Zhou*
25. Shape Driven Kernel Adaptation in Convolutional Neural Network for Robust Facial Traits Recognition, *Shaoxin Li, Junliang Xing, Zhiheng Niu, Shiguang Shan, Shuicheng Yan*
26. From Categories to Subcategories: Large-Scale Image Classification With Partial Class Label Refinement, *Marko Ristin, Juergen Gall, Matthieu Guillaumin, Luc Van Gool*
27. Combination Features and Models for Human Detection, *Yunsheng Jiang, Jinwen Ma*
28. Improving Object Detection With Deep Convolutional Networks via Bayesian Optimization and Structured Prediction, *Yuting Zhang, Kihyuk Sohn, Ruben Villegas, Gang Pan, Honglak Lee*
29. A Metric Parametrization for Trifocal Tensors With Non-Colinear Pinholes, *Spyridon Leonardos, Roberto Tron, Kostas Daniilidis*
30. An Efficient Volumetric Framework for Shape Tracking, *Benjamin Allain, Jean-Sébastien Franco, Edmond Boyer*
31. Structured Sparse Subspace Clustering: A Unified Optimization Framework, *Chun-Guang Li, René Vidal*
32. Delving Into Egocentric Actions, *Yin Li, Zhefan Ye, James M. Rehg*
33. Latent Trees for Estimating Intensity of Facial Action Units, *Sebastian Kaltwang, Sinisa Todorovic, Maja Pantic*
34. Robust Regression on Image Manifolds for Ordered Label Denoising, *Hui Wu, Richard Souvenir*
35. Privacy Preserving Optics for Miniature Vision Sensors, *Francesco Pittaluga, Sanjeev J. Koppal*
36. Deep Transfer Metric Learning, *Junlin Hu, Jiwen Lu, Yap-Peng Tan*
37. Small-Variance Nonparametric Clustering on the Hypersphere, *Julian Straub, Trevor Campbell, Jonathan P. How, John W. Fisher III*
38. DynamicFusion: Reconstruction and Tracking of Non-Rigid Scenes in Real-Time, *Richard A. Newcombe, Dieter Fox, Steven M. Seitz*
39. Reliable Patch Trackers: Robust Visual Tracking by Exploiting Reliable Patches, *Yang Li, Jianke Zhu, Steven C.H. Hoi*
40. Predicting Eye Fixations Using Convolutional Neural Networks, *Nian Liu, Junwei Han, Dingwen Zhang, Shifeng Wen, Tianming Liu*
41. Kernel Fusion for Better Image Deblurring, *Long Mai, Feng Liu*
42. Direction Matters: Depth Estimation With a Surface Normal Classifier, *Christian Häne, Lubor Ladický, Marc Pollefeys*
43. Modeling Local and Global Deformations in Deep Learning: Epitomic Convolution, Multiple Instance Learning, and Sliding Window Detection, *George Papandreou, Iasonas Kokkinos, Pierre-André Savalle*
44. Grasp Type Revisited: A Modern Perspective on a Classical Feature for Vision, *Yezhou Yang, Cornelia Fermüller, Yi Li, Yiannis Aloimonos*
45. Learning Hypergraph-Regularized Attribute Predictors, *Sheng Huang, Mohamed Elhoseiny, Ahmed Elgammal, Dan Yang*
46. A Coarse-to-Fine Model for 3D Pose Estimation and Sub-Category Recognition, *Roosbeh Mottaghi, Yu Xiang, Silvio Savarese*
47. Deep Neural Networks Are Easily Fooled: High Confidence Predictions for Unrecognizable Images, *Anh Nguyen, Jason Yosinski, Jeff Clune*
48. Deformable Part Models are Convolutional Neural Networks, *Ross Girshick, Forrest Iandola, Trevor Darrell, Jitendra Malik*
49. Hypercolumns for Object Segmentation and Fine-Grained Localization, *Bharath Hariharan, Pablo Arbeláez, Ross Girshick, Jitendra Malik*
50. Mapping Visual Features to Semantic Profiles for Retrieval in Medical Imaging, *Johannes Hofmanninger, Georg Langs*
51. Event-Driven Stereo Matching for Real-Time 3D Panoramic Vision, *Stephan Schraml, Ahmed Nabil Belbachir, Horst Bischof*



52. Graph-Based Simplex Method for Pairwise Energy Minimization With Binary Variables, *Daniel Průša*
53. Image Denoising via Adaptive Soft-Thresholding Based on Non-Local Samples, *Hangfan Liu, Ruiqin Xiong, Jian Zhang, Wen Gao*
54. 3D Scanning Deformable Objects With a Single RGBD Sensor, *Mingsong Dou, Jonathan Taylor, Henry Fuchs, Andrew Fitzgibbon, Shahram Izadi*
55. Nested Motion Descriptors, *Jeffrey Byrne*
56. Efficient Minimal-Surface Regularization of Perspective Depth Maps in Variational Stereo, *Gottfried Graber, Jonathan Balzer, Stefano Soatto, Thomas Pock*
57. Maximum Persistency via Iterative Relaxed Inference With Graphical Models, *Alexander Shekhovtsov, Paul Swoboda, Bogdan Savchynskyy*
58. Deep Hierarchical Parsing for Semantic Segmentation, *Abhishek Sharma, Oncel Tuzel, David W. Jacobs*
59. Designing Deep Networks for Surface Normal Estimation, *Xiaolong Wang, David Fouhey, Abhinav Gupta*
60. Layered RGBD Scene Flow Estimation, *Deqing Sun, Erik B. Sudderth, Hanspeter Pfister*
61. Hashing With Binary Autoencoders, *Miguel Á. Carreira-Perpiñán, Ramin Raziperchikolaei*
62. SUN RGB-D: A RGB-D Scene Understanding Benchmark Suite, *Shuran Song, Samuel P. Lichtenberg, Jianxiang Xiao*
63. Collaborative Feature Learning From Social Media, *Chen Fang, Hailin Jin, Jianchao Yang, Zhe Lin*
64. Diversity-Induced Multi-View Subspace Clustering, *Xiaochun Cao, Changqing Zhang, Huazhu Fu, Si Liu, Hua Zhang*
65. Building a Bird Recognition App and Large Scale Dataset With Citizen Scientists: The Fine Print in Fine-Grained Dataset Collection, *Grant Van Horn, Steve Branson, Ryan Farrell, Scott Haber, Jessie Barry, Panos Ipeirotis, Pietro Perona, Serge Belongie*
66. Early Burst Detection for Memory-Efficient Image Retrieval, *Miaojing Shi, Yannis Avrithis, Hervé Jégou*
67. Indoor Scene Structure Analysis for Single Image Depth Estimation, *Wei Zhuo, Mathieu Salzmann, Xuming He, Miaoqiao Liu*
68. Light Field Layer Matting, *Juliet Fiss, Brian Curless, Rick Szeliski*
69. Depth Camera Tracking With Contour Cues, *Qian-Yi Zhou, Vladlen Koltun*
70. Radial Distortion Homography, *Zuzana Kukelova, Jan Heller, Martin Bujnak, Tomas Pajdla*
71. Efficient Object Localization Using Convolutional Networks, *Jonathan Tompson, Ross Goroshin, Arjun Jain, Yann LeCun, Christoph Bregler*
72. Just Noticeable Defocus Blur Detection and Estimation, *Jianping Shi, Li Xu, Jiaya Jia*
73. How Do We Use Our Hands? Discovering a Diverse Set of Common Grasps, *De-An Huang, Minghuang Ma, Wei-Chiu Ma, Kris M. Kitani*
74. Rotating Your Face Using Multi-Task Deep Neural Network, *Junho Yim, Heechul Jung, ByungIn Yoo, Changkyu Choi, Dusik Park, Junmo Kim*
75. Is Object Localization for Free? – Weakly-Supervised Learning With Convolutional Neural Networks, *Maxime Oquab, Léon Bottou, Ivan Laptev, Josef Sivic*
76. Super-Resolution Person Re-Identification With Semi-Coupled Low-Rank Discriminant Dictionary Learning, *Xiao-Yuan Jing, Xiaoke Zhu, Fei Wu, Xinge You, Qinglong Liu, Dong Yue, Ruimin Hu, Baowen Xu*
77. Dual Domain Filters Based Texture and Structure Preserving Image Non-Blind Deconvolution, *Hang Yang, Ming Zhu, Yan Niu, Yujing Guan, Zhongbo Zhang*
78. Region-Based Temporally Consistent Video Post-Processing, *Xuan Dong, Boyan Bonev, Yu Zhu, Alan L. Yuille*
79. Global Refinement of Random Forest, *Shaoqing Ren, Xudong Cao, Yichen Wei, Jian Sun*
80. Adaptive Region Pooling for Object Detection, *Yi-Hsuan Tsai, Onur C. Hamsici, Ming-Hsuan Yang*
81. Discriminative and Consistent Similarities in Instance-Level Multiple Instance Learning, *Mohammad Rastegari, Hannaneh Hajishirzi, Ali Farhadi*
82. Multi-Store Tracker (MUSTer): A Cognitive Psychology Inspired Approach to Object Tracking, *Zhibin Hong, Zhe Chen, Chaohui Wang, Xue Mei, Danil Prokhorov, Dacheng Tao*
83. Finding Action Tubes, *Georgia Gkioxari, Jitendra Malik*



84. Learning a Convolutional Neural Network for Non-Uniform Motion Blur Removal, *Jian Sun, Wenfei Cao, Zongben Xu, Jean Ponce*
85. Complexity-Adaptive Distance Metric for Object Proposals Generation, *Yao Xiao, Cewu Lu, Efstratios Tsougenis, Yongyi Lu, Chi-Keung Tang*
86. High-Fidelity Pose and Expression Normalization for Face Recognition in the Wild, *Xiangyu Zhu, Zhen Lei, Junjie Yan, Dong Yi, Stan Z. Li*
87. Transformation of Markov Random Fields for Marginal Distribution Estimation, *Masaki Saito, Takayuki Okatani*
88. Sparse Convolutional Neural Networks, *Baoyuan Liu, Min Wang, Hassan Foroosh, Marshall Tappen, Marianna Pensky*
89. FaceNet: A Unified Embedding for Face Recognition and Clustering, *Florian Schroff, Dmitry Kalenichenko, James Philbin*
90. Cascaded Hand Pose Regression, *Xiao Sun, Yichen Wei, Shuang Liang, Xiaoou Tang, Jian Sun*
91. Cross-Scene Crowd Counting via Deep Convolutional Neural Networks, *Cong Zhang, Hongsheng Li, Xiaogang Wang, Xiaokang Yang*
92. The Application of Two-Level Attention Models in Deep Convolutional Neural Network for Fine-Grained Image Classification, *Tianjun Xiao, Yichong Xu, Kuiyuan Yang, Jiaxing Zhang, Yuxin Peng, Zheng Zhang*
93. End-to-End Integration of a Convolution Network, Deformable Parts Model and Non-Maximum Suppression, *Li Wan, David Eigen, Rob Fergus*
94. A Mixed Bag of Emotions: Model, Predict, and Transfer Emotion Distributions, *Kuan-Chuan Peng, Tshuan Chen, Amir Sadvovnik, Andrew C. Gallagher*
95. Neuroaesthetics in Fashion: Modeling the Perception of Fashionability, *Edgar Simo-Serra, Sanja Fidler, Francesc Moreno-Noguer, Raquel Urtasun*
96. Part-Based Modelling of Compound Scenes From Images, *Anton van den Hengel, Chris Russell, Anthony Dick, John Bastian, Daniel Pooley, Lachlan Fleming, Lourdes Agapito*
97. Efficient Parallel Optimization for Potts Energy With Hierarchical Fusion, *Olga Veksler*
98. Pooled Motion Features for First-Person Videos, *Michael S. Ryoo, Brandon Rothrock, Larry Matthies*
99. Functional Correspondence by Matrix Completion, *Artiom Kovnatsky, Michael M. Bronstein, Xavier Bresson, Pierre Vandergheynst*
100. Elastic-Net Regularization of Singular Values for Robust Subspace Learning, *Eunwoo Kim, Minsik Lee, Songhwa Oh*
101. Hardware Compliant Approximate Image Codes, *Da Kuang, Alex Gittens, Raffay Hamid*
102. Photometric Refinement of Depth Maps for Multi-Albedo Objects, *Avishek Chatterjee, Venu Madhav Govindu*
103. Predicting the Future Behavior of a Time-Varying Probability Distribution, *Christoph H. Lampert*
104. Classifier Based Graph Construction for Video Segmentation, *Anna Khoreva, Fabio Galasso, Matthias Hein, Bernt Schiele*
105. ActivityNet: A Large-Scale Video Benchmark for Human Activity Understanding, *Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, Juan Carlos Nieves*
106. Mid-Level Deep Pattern Mining, *Yao Li, Lingqiao Liu, Chunhua Shen, Anton van den Hengel*
107. Prediction of Search Targets From Fixations in Open-World Settings, *Hosnieh Sattar, Sabine Müller, Mario Fritz, Andreas Bulling*
108. Understanding Image Representations by Measuring Their Equivariance and Equivalence, *Karel Lenc, Andrea Vedaldi*
109. Effective Learning-Based Illuminant Estimation Using Simple Features, *Dongliang Cheng, Brian Price, Scott Cohen, Michael S. Brown*
110. PAIGE: PAirwise Image Geometry Encoding for Improved Efficiency in Structure-From-Motion, *Johannes L. Schönberger, Alexander C. Berg, Jan-Michael Frahm*
111. Dense, Accurate Optical Flow Estimation With Piecewise Parametric Model, *Jiaolong Yang, Hongdong Li*
112. Single-Image Estimation of the Camera Response Function in Near-Lighting, *Pedro Rodrigues, João P. Barreto*
113. Multispectral Pedestrian Detection: Benchmark Dataset and Baseline, *Soonmin Hwang, Jaesik Park, Namil Kim, Yukyung Choi, In So Kweon*



- 114. A Low-Dimensional Step Pattern Analysis Algorithm With Application to Multimodal Retinal Image Registration, *Jimmy Addison Lee, Jun Cheng, Beng Hai Lee, Ee Ping Ong, Guozhen Xu, Damon Wing Kee Wong, Jiang Liu, Augustinus Laude, Tock Han Lim*
- 115. Bilinear Heterogeneous Information Machine for RGB-D Action Recognition, *Yu Kong, Yun Fu*
- 116. MRF Optimization by Graph Approximation, *Wonsik Kim, Kyoung Mu Lee*
- 117. SALICON: Saliency in Context, *Ming Jiang, Shengsheng Huang, Juanyong Duan, Qi Zhao*
- 118. Weakly Supervised Object Detection With Convex Clustering, *Hakan Bilen, Marco Pedersoli, Tinne Tuytelaars*
- 119. Interleaved Text/Image Deep Mining on a Very Large-Scale Radiology Database, *Hoo-Chang Shin, Le Lu, Lauren Kim, Ari Seff, Jianhua Yao, Ronald M. Summers*
- 120. Learning Semantic Relationships for Better Action Retrieval in Images, *Vignesh Ramanathan, Congcong Li, Jia Deng, Wei Han, Zhen Li, Kunlong Gu, Yang Song, Samy Bengio, Charles Rosenberg, Li Fei-Fei*
- 121. Hierarchical Recurrent Neural Network for Skeleton Based Action Recognition, *Yong Du, Wei Wang, Liang Wang*

## 1230–1400 Lunch (Exhibit Hall B)



## 1400–1530 Oral Mon-PM-1: Discovery & Dense Correspondences in Image Datasets & Videos (Ballrooms A-B)

Papers in this session are also in Poster Session Mon-PM.

**Chairs:** Devi Parikh (*Virginia Tech*)  
Jakob Verbeek (*INRIA*)

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

1. Discovering States and Transformations in Image Collections, *Phillip Isola, Joseph J. Lim, Edward H. Adelson*
2. Unsupervised Object Discovery and Localization in the Wild: Part-Based Matching With Bottom-Up Region Proposals, *Minsu Cho, Suha Kwak, Cordelia Schmid, Jean Ponce*
3. FlowWeb: Joint Image Set Alignment by Weaving Consistent, Pixel-Wise Correspondences, *Tinghui Zhou, Yong Jae Lee, Stella X. Yu, Alyosha A. Efros*
4. EpicFlow: Edge-Preserving Interpolation of Correspondences for Optical Flow, *Jerome Revaud, Philippe Weinzaepfel, Zaid Harchaoui, Cordelia Schmid*
5. Phase-Based Frame Interpolation for Video, *Simone Meyer, Oliver Wang, Henning Zimmer, Max Grosse, Alexander Sorkine-Hornung*
6. Towards Open World Recognition, *Abhijit Bendale, Terrance Boulton*

## 1400–1530 Oral Mon-PM-2: 3D Shape — Matching, Recognition, Reconstruction (Rooms 302-306)

Papers in this session are also in Poster Session Mon-PM.

**Chairs:** Derek Hoiem (*UIUC*)  
Raquel Urtasun (*Univ. of Toronto*)

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

1. Category-Specific Object Reconstruction From a Single Image, *Abhishek Kar, Shubham Tulsiani, João Carreira, Jitendra Malik*
2. Discriminative Shape From Shading in Uncalibrated Illumination, *Stephan R. Richter, Stefan Roth*
3. Learning to Generate Chairs With Convolutional Neural Networks, *Alexey Dosovitskiy, Jost Tobias Springenberg, Thomas Brox*

4. 3D ShapeNets: A Deep Representation for Volumetric Shapes, *Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, Jianxiong Xiao*
5. Sketch-Based 3D Shape Retrieval Using Convolutional Neural Networks, *Fang Wang, Le Kang, Yi Li*
6. Data-Driven 3D Voxel Patterns for Object Category Recognition, *Yu Xiang, Wongun Choi, Yuanqing Lin, Silvio Savarese*

## 1530–1615 Break (Exhibit Hall B)

## 1530–1800 Exhibits (Exhibit Hall A)

- Same as Monday morning Exhibits (see pg. 11)

## 1530–1800 Demos (Room 108)

- Same as Monday morning Demos (see pg. 12)

## 1530–1800 Poster Mon-PM (Exhibit Hall A)

1. Depth and Surface Normal Estimation From Monocular Images Using Regression on Deep Features and Hierarchical CRFs, *Bo Li, Chunhua Shen, Yuchao Dai, Anton van den Hengel, Mingyi He*
2. Discriminative Shape From Shading in Uncalibrated Illumination, *Stephan R. Richter, Stefan Roth*
3. Multi-Manifold Deep Metric Learning for Image Set Classification, *Jiwen Lu, Gang Wang, Weihong Deng, Pierre Moulin, Jie Zhou*
4. Target Identity-Aware Network Flow for Online Multiple Target Tracking, *Afshin Dehghan, Yicong Tian, Philip H. S. Torr, Mubarak Shah*
5. Adaptive As-Natural-As-Possible Image Stitching, *Chung-Ching Lin, Sharathchandra U. Pankanti, Karthikeyan Natesan Ramamurthy, Aleksandr Y. Aravkin*
6. EpicFlow: Edge-Preserving Interpolation of Correspondences for Optical Flow, *Jerome Revaud, Philippe Weinzaepfel, Zaid Harchaoui, Cordelia Schmid*
7. Learning Coarse-to-Fine Sparselets for Efficient Object Detection and Scene Classification, *Gong Cheng, Junwei Han, Lei Guo, Tianming Liu*
8. Continuous Visibility Feature, *Guilin Liu, Yotam Gingold, Jyh-Ming Lien*



9. FlowWeb: Joint Image Set Alignment by Weaving Consistent, Pixel-Wise Correspondences, *Tinghui Zhou, Yong Jae Lee, Stella X. Yu, Alyosha A. Efros*
10. Unsupervised Object Discovery and Localization in the Wild: Part-Based Matching With Bottom-Up Region Proposals, *Minsu Cho, Suha Kwak, Cordelia Schmid, Jean Ponce*
11. Supervised Descriptor Learning for Multi-Output Regression, *Xiantong Zhen, Zhijie Wang, Mengyang Yu, Shuo Li*
12. A Statistical Model of Riemannian Metric Variation for Deformable Shape Analysis, *Andrea Gasparetto, Andrea Torsello*
13. Temporally Coherent Interpretations for Long Videos Using Pattern Theory, *Fillipe Souza, Sudeep Sarkar, Anuj Srivastava, Jingyong Su*
14. Line-Sweep: Cross-Ratio For Wide-Baseline Matching and 3D Reconstruction, *Srikumar Ramalingam, Michel Antunes, Dan Snow, Gim Hee Lee, Sudeep Pillai*
15. Simplified Mirror-Based Camera Pose Computation via Rotation Averaging, *Gucan Long, Laurent Kneip, Xin Li, Xiaohu Zhang, Qifeng Yu*
16. On the Relationship Between Visual Attributes and Convolutional Networks, *Victor Escorcia, Juan Carlos Niebles, Bernard Ghanem*
17. Saliency Detection by Multi-Context Deep Learning, *Rui Zhao, Wanli Ouyang, Hongsheng Li, Xiaogang Wang*
18. DeepShape: Deep Learned Shape Descriptor for 3D Shape Matching and Retrieval, *Jin Xie, Yi Fang, Fan Zhu, Edward Wong*
19. Bayesian Adaptive Matrix Factorization With Automatic Model Selection, *Peixian Chen, Naiyan Wang, Nevin L. Zhang, Dit-Yan Yeung*
20. Joint Action Recognition and Pose Estimation From Video, *Bruce Xiaohan Nie, Caiming Xiong, Song-Chun Zhu*
21. Fast Action Proposals for Human Action Detection and Search, *Gang Yu, Junsong Yuan*
22. Joint Multi-Feature Spatial Context for Scene Recognition on the Semantic Manifold, *Xinhang Song, Shuqiang Jiang, Luis Herranz*
23. Large-Scale Damage Detection Using Satellite Imagery, *Lionel Gueguen, Raffay Hamid*
24. A Novel Locally Linear KNN Model for Visual Recognition, *Qingfeng Liu, Chengjun Liu*
25. Bilinear Random Projections for Locality-Sensitive Binary Codes, *Saeheon Kim, Seungjin Choi*
26. Combining Local Appearance and Holistic View: Dual-Source Deep Neural Networks for Human Pose Estimation, *Xiaochuan Fan, Kang Zheng, Yuewei Lin, Song Wang*
27. Superpixel Segmentation Using Linear Spectral Clustering, *Zhengqin Li, Jiansheng Chen*
28. Person Count Localization in Videos From Noisy Foreground and Detections, *Sheng Chen, Alan Fern, Sinisa Todorovic*
29. Good Features to Track for Visual SLAM, *Guangcong Zhang, Patricio A. Vela*
30. Discovering States and Transformations in Image Collections, *Phillip Isola, Joseph J. Lim, Edward H. Adelson*
31. Generalized Deformable Spatial Pyramid: Geometry-Preserving Dense Correspondence Estimation, *Junhua Hur, Hwasup Lim, Changsoo Park, Sang Chul Ahn*
32. Classifier Adaptation at Prediction Time, *Amélie Royer, Christoph H. Lampert*
33. Phase-Based Frame Interpolation for Video, *Simone Meyer, Oliver Wang, Henning Zimmer, Max Grosse, Alexander Sorkine-Hornung*
34. Matching-CNN Meets KNN: Quasi-Parametric Human Parsing, *Si Liu, Xiaodan Liang, Luoqi Liu, Xiaohui Shen, Jianchao Yang, Changsheng Xu, Liang Lin, Xiaochun Cao, Shuicheng Yan*
35. Absolute Pose for Cameras Under Flat Refractive Interfaces, *Sebastian Haner, Kalle Åström*
36. Protecting Against Screenshots: An Image Processing Approach, *Alex Yong-Sang Chia, Udana Bandara, Xiangyu Wang, Hiromi Hirano*
37. Pose-Conditioned Joint Angle Limits for 3D Human Pose Reconstruction, *Ijaz Akhter, Michael J. Black*
38. VisKE: Visual Knowledge Extraction and Question Answering by Visual Verification of Relation Phrases, *Fereshteh Sadeghi, Santosh K. Kumar Divvala, Ali Farhadi*
39. A Graphical Model Approach for Matching Partial Signatures, *Xianzhi Du, David Doermann, Wael Abd-Almageed*



40. From Captions to Visual Concepts and Back, *Hao Fang, Saurabh Gupta, Forrest Iandola, Rupesh K. Srivastava, Li Deng, Piotr Dollár, Jianfeng Gao, Xiaodong He, Margaret Mitchell, John C. Platt, C. Lawrence Zitnick, Geoffrey Zweig*
41. Semi-Supervised Low-Rank Mapping Learning for Multi-Label Classification, *Liping Jing, Liu Yang, Jian Yu, Michael K. Ng*
42. ConceptLearner: Discovering Visual Concepts From Weakly Labeled Image Collections, *Bolei Zhou, Vignesh Jagadeesh, Robinson Piramuthu*
43. Computationally Bounded Retrieval, *Mohammad Rastegari, Cem Keskin, Pushmeet Kohli, Shahram Izadi*
44. Viewpoints and Keypoints, *Shubham Tulsiani, Jitendra Malik*
45. Discrete Hyper-Graph Matching, *Junchi Yan, Chao Zhang, Hongyuan Zha, Wei Liu, Xiaokang Yang, Stephen M. Chu*
46. Rolling Shutter Motion Deblurring, *Shuochen Su, Wolfgang Heidrich*
47. Learning to Generate Chairs With Convolutional Neural Networks, *Alexey Dosovitskiy, Jost Tobias Springenberg, Thomas Brox*
48. Accurate Depth Map Estimation From a Lenslet Light Field Camera, *Hae-Gon Jeon, Jaesik Park, Gyeongmin Choe, Jinsun Park, Yunsu Bok, Yu-Wing Tai, In So Kweon*
49. Deep Semantic Ranking Based Hashing for Multi-Label Image Retrieval, *Fang Zhao, Yongzhen Huang, Liang Wang, Tieniu Tan*
50. Similarity Learning on an Explicit Polynomial Kernel Feature Map for Person Re-Identification, *Dapeng Chen, Zejian Yuan, Gang Hua, Nanning Zheng, Jingdong Wang*
51. Learning to Propose Objects, *Philipp Krähenbühl, Vladlen Koltun*
52. Basis Mapping Based Boosting for Object Detection, *Haoyu Ren, Ze-Nian Li*
53. Computing the Stereo Matching Cost With a Convolutional Neural Network, *Jure Žbontar, Yann LeCun*
54. Recognize Complex Events From Static Images by Fusing Deep Channels, *Yuanjun Xiong, Kai Zhu, Dahua Lin, Xiaoou Tang*
55. Multi-Feature Max-Margin Hierarchical Bayesian Model for Action Recognition, *Shuang Yang, Chunfeng Yuan, Baoxin Wu, Weiming Hu, Fangshi Wang*
56. Model Recommendation: Generating Object Detectors From Few Samples, *Yu-Xiong Wang, Martial Hebert*
57. A Linear Least-Squares Solution to Elastic Shape-From-Template, *Abed Malti, Adrien Bartoli, Richard Hartley*
58. Robust Large Scale Monocular Visual SLAM, *Guillaume Bourmaud, Rémi Mège*
59. Membership Representation for Detecting Block-Diagonal Structure in Low-Rank or Sparse Subspace Clustering, *Minsik Lee, Jieun Lee, Hyeogjin Lee, Naeun Kwak*
60. Bayesian Inference for Neighborhood Filters With Application in Denoising, *Chao-Tsung Huang*
61. Deep LAC: Deep Localization, Alignment and Classification for Fine-Grained Recognition, *Di Lin, Xiaoyong Shen, Cewu Lu, Jiaya Jia*
62. Unconstrained Realtime Facial Performance Capture, *Pei-Lun Hsieh, Chongyang Ma, Jihun Yu, Hao Li*
63. Blind Optical Aberration Correction by Exploring Geometric and Visual Priors, *Tao Yue, Jinli Suo, Jue Wang, Xun Cao, Qionghai Dai*
64. Ontological Supervision for Fine Grained Classification of Street View Storefronts, *Yair Movshovitz-Attias, Qian Yu, Martin C. Stumpe, Vinay Shet, Sacha Arnoud, Liron Yatziv*
65. Finding Distractors In Images, *Ohad Fried, Eli Shechtman, Dan B. Goldman, Adam Finkelstein*
66. From Image-Level to Pixel-Level Labeling With Convolutional Networks, *Pedro O. Pinheiro, Ronan Collobert*
67. Semantic Alignment of LiDAR Data at City Scale, *Fisher Yu, Jianxiong Xiao, Thomas Funkhouser*
68. Oriented Edge Forests for Boundary Detection, *Sam Hallman, Charles C. Fowlkes*
69. Query-Adaptive Late Fusion for Image Search and Person Re-Identification, *Liang Zheng, Shengjin Wang, Lu Tian, Fei He, Ziqiong Liu, Qi Tian*
70. Filtered Feature Channels for Pedestrian Detection, *Shanshan Zhang, Rodrigo Benenson, Bernt Schiele*
71. GRSA: Generalized Range Swap Algorithm for the Efficient Optimization of MRFs, *Kangwei Liu, Junge Zhang, Peipei Yang, Kaiqi Huang*
72. PatchCut: Data-Driven Object Segmentation via Local Shape Transfer, *Jimei Yang, Brian Price, Scott Cohen, Zhe Lin, Ming-Hsuan Yang*



73. Illumination and Reflectance Spectra Separation of a Hyperspectral Image Meets Low-Rank Matrix Factorization, *Yinqiang Zheng, Imari Sato, Yoichi Sato*
74. Semantic Part Segmentation Using Compositional Model Combining Shape and Appearance, *Jianyu Wang, Alan L. Yuille*
75. A Discriminative CNN Video Representation for Event Detection, *Zhongwen Xu, Yi Yang, Alex G. Hauptmann*
76. 24/7 Place Recognition by View Synthesis, *Akihiko Torii, Relja Arandjelović, Josef Sivic, Masatoshi Okutomi, Tomas Pajdla*
77. Understanding Image Virality, *Arturo Deza, Devi Parikh*
78. Book2Movie: Aligning Video Scenes With Book Chapters, *Makarand Tapaswi, Martin Bäuml, Rainer Stiefelhofen*
79. 3D Model-Based Continuous Emotion Recognition, *Hui Chen, Jiangdong Li, Fengjun Zhang, Yang Li, Hongan Wang*
80. Learning to Rank in Person Re-Identification With Metric Ensembles, *Sakrapree Paisitkriangkrai, Chunhua Shen, Anton van den Hengel*
81. Making Better Use of Edges via Perceptual Grouping, *Yonggang Qi, Yi-Zhe Song, Tao Xiang, Honggang Zhang, Timothy Hospedales, Yi Li, Jun Guo*
82. Real-Time Joint Estimation of Camera Orientation and Vanishing Points, *Jeong-Kyun Lee, Kuk-Jin Yoon*
83. Sketch-Based 3D Shape Retrieval Using Convolutional Neural Networks, *Fang Wang, Le Kang, Yi Li*
84. Salient Object Detection via Bootstrap Learning, *Na Tong, Huchuan Lu, Xiang Ruan, Ming-Hsuan Yang*
85. Towards Open World Recognition, *Abhijit Bendale, Terrance Boulton*
86. Data-Driven 3D Voxel Patterns for Object Category Recognition, *Yu Xiang, Wongun Choi, Yuanqing Lin, Silvio Savarese*
87. 3D ShapeNets: A Deep Representation for Volumetric Shapes, *Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, Jianxiong Xiao*
88. Robust Image Alignment With Multiple Feature Descriptors and Matching-Guided Neighborhoods, *Kuang-Jui Hsu, Yen-Yu Lin, Yung-Yu Chuang*
89. Pushing the Frontiers of Unconstrained Face Detection and Recognition: IARPA Janus Benchmark A, *Brendan F. Klare, Ben Klein, Emma Taborsky, Austin Blanton, Jordan Cheney, Kristen Allen, Patrick Grother, Alan Mah, Mark Burge, Anil K. Jain*
90. Depth From Shading, Defocus, and Correspondence Using Light-Field Angular Coherence, *Michael W. Tao, Pratul P. Srinivasan, Jitendra Malik, Szymon Rusinkiewicz, Ravi Ramamoorthi*
91. New Insights Into Laplacian Similarity Search, *Xiao-Ming Wu, Zhenguo Li, Shih-Fu Chang*
92. Feature-Independent Context Estimation for Automatic Image Annotation, *Amara Tariq, Hassan Foroosh*
93. Category-Specific Object Reconstruction From a Single Image, *Abhishek Kar, Shubham Tulsiani, João Carneira, Jitendra Malik*
94. Active Sample Selection and Correction Propagation on a Gradually-Augmented Graph, *Hang Su, Zhaozheng Yin, Takeo Kanade, Seungil Huh*
95. Efficient and Accurate Approximations of Nonlinear Convolutional Networks, *Xiangyu Zhang, Jianhua Zou, Xiang Ming, Kaifeng He, Jian Sun*
96. Ranking and Retrieval of Image Sequences From Multiple Paragraph Queries, *Gunhee Kim, Seungwhan Moon, Leonid Sigal*
97. Casual Stereoscopic Panorama Stitching, *Fan Zhang, Feng Liu*
98. Superpixel Meshes for Fast Edge-Preserving Surface Reconstruction, *András Bádis-Szomorú, Hayko Riemenschneider, Luc Van Gool*
99. Best-Buddies Similarity for Robust Template Matching, *Tali Dekel, Shaul Oron, Michael Rubinstein, Shai Avidan, William T. Freeman*
100. Superdifferential Cuts for Binary Energies, *Tatsunori Tanai, Yasuyuki Matsushita, Takeshi Naemura*
101. The S-HOCK Dataset: Analyzing Crowds at the Stadium, *Davide Conigliaro, Paolo Rota, Francesco Setti, Chiara Bassetti, Nicola Conci, Nicu Sebe, Marco Cristani*
102. Discriminant Analysis on Riemannian Manifold of Gaussian Distributions for Face Recognition With Image Sets, *Wen Wang, Ruiping Wang, Zhiwu Huang, Shiguang Shan, Xilin Chen*



103. Texture Representations for Image and Video Synthesis, *Georgios Georgiadis, Alessandro Chiuso, Stefano Soatto*
104. Shadow Optimization From Structured Deep Edge Detection, *Li Shen, Teck Wee Chua, Karianto Leman*
105. Total Variation Regularization of Shape Signals, *Maximilian Baust, Laurent Demaret, Martin Storath, Nassir Navab, Andreas Weinmann*
106. Learning Similarity Metrics for Dynamic Scene Segmentation, *Damien Teney, Matthew Brown, Dmitry Kit, Peter Hall*
107. Subspace Clustering by Mixture of Gaussian Regression, *Baohua Li, Ying Zhang, Zhouchen Lin, Huchuan Lu*
108. DASC: Dense Adaptive Self-Correlation Descriptor for Multi-Modal and Multi-Spectral Correspondence, *Seungryoung Kim, Dongbo Min, Bumsu Ham, Seungchul Ryu, Minh N. Do, Kwanghoon Sohn*
109. In Defense of Color-Based Model-Free Tracking, *Horst Possegger, Thomas Mauthner, Horst Bischof*
110. Best of Both Worlds: Human-Machine Collaboration for Object Annotation, *Olga Russakovsky, Li-Jia Li, Li Fei-Fei*
111. Robust Multiple Homography Estimation: An Ill-Solved Problem, *Zygmunt L. Szpak, Wojciech Chojnacki, Anton van den Hengel*
112. Semi-Supervised Domain Adaptation With Subspace Learning for Visual Recognition, *Ting Yao, Yingwei Pan, Chong-Wah Ngo, Houqiang Li, Tao Mei*
113. Articulated Motion Discovery Using Pairs of Trajectories, *Luca Del Pero, Susanna Ricco, Rahul Sukthankar, Vittorio Ferrari*
114. A Solution for Multi-Alignment by Transformation Synchronisation, *Florian Bernard, Johan Thunberg, Peter Gemmar, Frank Hertel, Andreas Husch, Jorge Goncalves*
115. A Convex Optimization Approach to Robust Fundamental Matrix Estimation, *Yongfang Cheng, Jose A. Lopez, Octavia Camps, Mario Sznajder*
116. Simultaneous Pose and Non-Rigid Shape With Particle Dynamics, *Antonio Agudo, Francesc Moreno-Noguer*
117. Semi-Supervised Learning With Explicit Relationship Regularization, *Kwang In Kim, James Tompkin, Hanspeter Pfister, Christian Theobalt*
118. Person Re-Identification by Local Maximal Occurrence Representation and Metric Learning, *Shengcai Liao, Yang Hu, Xiangyu Zhu, Stan Z. Li*
119. Joint Patch and Multi-Label Learning for Facial Action Unit Detection, *Kaili Zhao, Wen-Sheng Chu, Fernando De la Torre, Jeffrey F. Cohn, Honggang Zhang*
120. Real-Time Visual Analysis of Microvascular Blood Flow for Critical Care, *Chao Liu, Hernando Gomez, Srinivasa Narasimhan, Artur Dubrawski, Michael R. Pinsky, Brian Zuckerbraun*

## 1800–1930 Reception & Awards (Ballrooms A-B)

## 1930–2030 PAMI Technical Committee Meeting (Rooms 302-306)



# Tuesday, June 9

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

**0730–0830 Breakfast** (Exhibit Hall B)

**0830–1000 Oral Tue-AM-1: Images/Videos & Words/Language; Humans in Loop** (Ballrooms A-B)

Papers in this session are also in **Poster Session Tue-AM**.

**Chairs:** Abhinav Gupta (*Carnegie Mellon Univ.*)  
Larry Zitnick (*Microsoft Research*)

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

1. Show and Tell: A Neural Image Caption Generator, *Oriol Vinyals, Alexander Toshev, Samy Bengio, Dumitru Erhan*
2. Deep Visual-Semantic Alignments for Generating Image Descriptions, *Andrej Karpathy, Li Fei-Fei*
3. Long-Term Recurrent Convolutional Networks for Visual Recognition and Description, *Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, Trevor Darrell*
4. Image Specificity, *Mainak Jas, Devi Parikh*
5. Don't Just Listen, Use Your Imagination: Leveraging Visual Common Sense for Non-Visual Tasks, *Xiao Lin, Devi Parikh*
6. Becoming the Expert - Interactive Multi-Class Machine Teaching, *Edward Johns, Oisin Mac Aodha, Gabriel J. Brostow*

**0830–1000 Oral Tue-AM-2: Multiple View Geometry** (Rooms 302-306)

Papers in this session are also in **Poster Session Tue-AM**.

**Chairs:** Noah Snavely (*Cornell Univ.*)  
Jean Ponce (*Ecole Normale Supérieure*)

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

1. Reconstructing the World\* in Six Days\* (As Captured by the Yahoo 100 Million Image Dataset), *Jared Heintz, Johannes L. Schönberger, Enrique Dunn, Jan-Michael Frahm*

2. Joint Vanishing Point Extraction and Tracking, *Till Kroeger, Dengxin Dai, Luc Van Gool*
3. Robust Camera Location Estimation by Convex Programming, *Onur Özyesil, Amit Singer*
4. Efficient Globally Optimal Consensus Maximisation With Tree Search, *Tat-Jun Chin, Pulak Purkait, Anders Eriksson, David Suter*
5. R6P - Rolling Shutter Absolute Camera Pose, *cenek Albl, Zuzana Kukelova, Tomas Pajdla*
6. Building Proteins in a Day: Efficient 3D Molecular Reconstruction, *Marcus A. Brubaker, Ali Punjani, David J. Fleet*

**1000–1045 Break** (Exhibit Hall B)

**1000–1230 Exhibits** (Exhibit Hall A)

- Same as Monday morning Exhibits (see pg. 11)

**1000–1230 Demos** (Room 108)

- Realtime Monocular Gaze Correction, *Daniil Kononenko, Victor Lempitsky (Skolkovo Institute of Science and Technology)*
- Multi-language Image Captioning and Retrieval, *Junhua Mao, Alan Yuille, Wei Xu, Yi Yang, Jiang Wang, Zhiheng Huang (Univ. of California, Los Angeles; Baidu Research)*
- Translating Videos to Natural Language, *Subhashini Venugopalan, Raymond Mooney, Huijuan Xu, Kate Saenko, Marcus Rohrbach, Trevor Darrell (Univ. of Texas, Austin; Univ. of Massachusetts, Lowell; UC Berkeley)*

**1000–1230 Poster Tue-AM** (Exhibit Hall A)

1. JOTS: Joint Online Tracking and Segmentation, *Longyin Wen, Dawei Du, Zhen Lei, Stan Z. Li, Ming-Hsuan Yang*
2. Gaze-Enabled Egocentric Video Summarization via Constrained Submodular Maximization, *Jia Xu, Lopamudra Mukherjee, Yin Li, Jamieson Warner, James M. Rehg, Vikas Singh*
3. Sparse Depth Super Resolution, *Jiajun Lu, David Forsyth*
4. Efficient Illuminant Estimation for Color Constancy Using Grey Pixels, *Kai-Fu Yang, Shao-Bing Gao, Yong-Jie Li*



5. Can Humans Fly? Action Understanding With Multiple Classes of Actors, *Chenliang Xu, Shao-Hang Hsieh, Caiming Xiong, Jason J. Corso*
6. Reweighted Laplace Prior Based Hyperspectral Compressive Sensing for Unknown Sparsity, *Lei Zhang, Wei Wei, Yanning Zhang, Chunna Tian, Fei Li*
7. Class Consistent Multi-Modal Fusion With Binary Features, *Ashish Shrivastava, Mohammad Rastegari, Sumit Shekhar, Rama Chellappa, Larry S. Davis*
8. R6P - Rolling Shutter Absolute Camera Pose, *Cenek Albl, Zuzana Kukelova, Tomas Pajdla*
9. Embedded Phase Shifting: Robust Phase Shifting With Embedded Signals, *Daniel Moreno, Kilho Son, Gabriel Taubin*
10. Shape and Light Directions From Shading and Polarization, *Trung Ngo Thanh, Hajime Nagahara, Rintaro Taniguchi*
11. 3D Deep Shape Descriptor, *Yi Fang, Jin Xie, Guoxian Dai, Meng Wang, Fan Zhu, Tiantian Xu, Edward Wong*
12. Cross-Age Face Verification by Coordinating With Cross-Face Age Verification, *Liang Du, Haibin Ling*
13. Beyond Mahalanobis Metric: Cayley-Klein Metric Learning, *Yanhong Bi, Bin Fan, Fuchao Wu*
14. From Dictionary of Visual Words to Subspaces: Locality-Constrained Affine Subspace Coding, *Peihua Li, Xiaoxiao Lu, Qilong Wang*
15. FPA-CS: Focal Plane Array-Based Compressive Imaging in Short-Wave Infrared, *Huajin Chen, M. Salman Asif, Aswin C. Sankaranarayanan, Ashok Veeraraghavan*
16. BOLD - Binary Online Learned Descriptor For Efficient Image Matching, *Vassileios Balntas, Lilian Tang, Krystian Mikolajczyk*
17. Defocus Deblurring and Superresolution for Time-of-Flight Depth Cameras, *Lei Xiao, Felix Heide, Matthew O'Toole, Andreas Kolb, Matthias B. Hullin, Kyros Kutulakos, Wolfgang Heidrich*
18. Burst Deblurring: Removing Camera Shake Through Fourier Burst Accumulation, *Mauricio Delbracio, Guillermo Sapiro*
19. SOM: Semantic Obviousness Metric for Image Quality Assessment, *Peng Zhang, Wengang Zhou, Lei Wu, Houqiang Li*
20. DeepID-Net: Deformable Deep Convolutional Neural Networks for Object Detection, *Wanli Ouyang, Xiaogang Wang, Xingyu Zeng, Shi Qiu, Ping Luo, Yonglong Tian, Hongsheng Li, Shuo Yang, Zhe Wang, Chen-Change Loy, Xiaoou Tang*
21. Efficient Globally Optimal Consensus Maximisation With Tree Search, *Tat-Jun Chin, Pulak Purkait, Anders Eriksson, David Suter*
22. Mind's Eye: A Recurrent Visual Representation for Image Caption Generation, *Xinlei Chen, C. Lawrence Zitnick*
23. Hierarchical Sparse Coding With Geometric Prior For Visual Geo-Location, *Raghuraman Gopalan*
24. P3.5P: Pose Estimation With Unknown Focal Length, *Changchang Wu*
25. Joint Vanishing Point Extraction and Tracking, *Till Kroeger, Dengxin Dai, Luc Van Gool*
26. Learning a Non-Linear Knowledge Transfer Model for Cross-View Action Recognition, *Hossein Rahmani, Ajmal Mian*
27. Random Tree Walk Toward Instantaneous 3D Human Pose Estimation, *Ho Yub Jung, Soochahn Lee, Yong Seok Heo, Il Dong Yun*
28. Deep Hashing for Compact Binary Codes Learning, *Venice Erin Liong, Jiwen Lu, Gang Wang, Pierre Moulin, Jie Zhou*
29. Completing 3D Object Shape From One Depth Image, *Jason Rock, Tanmay Gupta, Justin Thorsen, JunYoung Gwak, Daeyun Shin, Derek Hoiem*
30. Encoding Based Saliency Detection for Videos and Images, *Thomas Mauthner, Horst Possegger, Georg Waltner, Horst Bischof*
31. Online Sketching Hashing, *Cong Leng, Jiaxiang Wu, Jian Cheng, Xiao Bai, Hanqing Lu*
32. Enriching Object Detection With 2D-3D Registration and Continuous Viewpoint Estimation, *Christopher Bongsoo Choy, Michael Stark, Sam Corbett-Davies, Silvio Savarese*
33. Representing 3D Texture on Mesh Manifolds for Retrieval and Recognition Applications, *Naoufel Werghi, Claudio Tortorici, Stefano Berretti, Alberto Del Bimbo*
34. Saliency Propagation From Simple to Difficult, *Chen Gong, Dacheng Tao, Wei Liu, Stephen J. Maybank, Meng Fang, Keren Fu, Jie Yang*



35. Learning an Efficient Model of Hand Shape Variation From Depth Images, *Sameh Khamis, Jonathan Taylor, Jamie Shotton, Cem Keskin, Shahram Izadi, Andrew Fitzgibbon*
36. On the Minimal Problems of Low-Rank Matrix Factorization, *Fangyuan Jiang, Magnus Oskarsson, Kalle Åström*
37. Symmetry-Based Text Line Detection in Natural Scenes, *Zheng Zhang, Wei Shen, Cong Yao, Xiang Bai*
38. DevNet: A Deep Event Network for Multimedia Event Detection and Evidence Recounting, *Chuang Gan, Naiyan Wang, Yi Yang, Dit-Yan Yeung, Alex G. Hauptmann*
39. Learning to Detect Motion Boundaries, *Philippe Weinzaepfel, Jerome Revaud, Zaid Harchaoui, Cordelia Schmid*
40. Improving Object Proposals With Multi-Thresholding Straddling Expansion, *Xiaozi Chen, Huimin Ma, Xiang Wang, Zhichen Zhao*
41. Visual Recognition by Counting Instances: A Multi-Instance Cardinality Potential Kernel, *Hossein Hajimirsadeghi, Wang Yan, Arash Vahdat, Greg Mori*
42. Unconstrained 3D Face Reconstruction, *Joseph Roth, Yiyang Tong, Xiaoming Liu*
43. Becoming the Expert - Interactive Multi-Class Machine Teaching, *Edward Johns, Oisín Mac Aodha, Gabriel J. Brostow*
44. Long-Term Recurrent Convolutional Networks for Visual Recognition and Description, *Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, Trevor Darrell*
45. Zero-Shot Object Recognition by Semantic Manifold Distance, *Zhenyong Fu, Tao Xiang, Elyor Kodirov, Shaogang Gong*
46. Hyper-Class Augmented and Regularized Deep Learning for Fine-Grained Image Classification, *Saining Xie, Tianbao Yang, Xiaoyu Wang, Yuanqing Lin*
47. Direct Structure Estimation for 3D Reconstruction, *Nianjuan Jiang, Daniel Lin, Minh N. Do, Jiangbo Lu*
48. Global Supervised Descent Method, *Xuehan Xiong, Fernando De la Torre*
49. Robust Camera Location Estimation by Convex Programming, *Onur Özyesil, Amit Singer*
50. Practical Robust Two-View Translation Estimation, *Johan Fredriksson, Viktor Larsson, Carl Olsson*
51. Learning From Massive Noisy Labeled Data for Image Classification, *Tong Xiao, Tian Xia, Yi Yang, Chang Huang, Xiaogang Wang*
52. KL Divergence Based Agglomerative Clustering for Automated Vitiligo Grading, *Mithun Das Gupta, Srinidhi Srinivasa, Madhukara J., Meryl Antony*
53. Robust Saliency Detection via Regularized Random Walks Ranking, *Changyang Li, Yuchen Yuan, Weidong, Cai, Yong Xia, David Dagan Feng*
54. Weakly Supervised Semantic Segmentation for Social Images, *Wei Zhang, Sheng Zeng, Dequan Wang, Xiangyang Xue*
55. Image Specificity, *Mainak Jas, Devi Parikh*
56. A Multi-Plane Block-Coordinate Frank-Wolfe Algorithm for Training Structural SVMs With a Costly Max-Oracle, *Neel Shah, Vladimir Kolmogorov, Christoph H. Lampert*
57. Web-Scale Training for Face Identification, *Yaniv Taigman, Ming Yang, Marc'Aurelio Ranzato, Lior Wolf*
58. Dynamically Encoded Actions Based on Spacetime Saliency, *Christoph Feichtenhofer, Axel Pinz, Richard P. Wildes*
59. Three Viewpoints Toward Exemplar SVM, *Takumi Kobayashi*
60. Visual Recognition by Learning From Web Data: A Weakly Supervised Domain Generalization Approach, *Li Niu, Wen Li, Dong Xu*
61. Clustering of Static-Adaptive Correspondences for Deformable Object Tracking, *Georg Nebelbay, Roman Pflugfelder*
62. Geo-Semantic Segmentation, *Shervin Ardeshtir, Kofi Malcolm Collins-Sibley, Mubarak Shah*
63. Towards Unified Depth and Semantic Prediction From a Single Image, *Peng Wang, Xiaohui Shen, Zhe Lin, Scott Cohen, Brian Price, Alan L. Yuille*
64. Towards Force Sensing From Vision: Observing Hand-Object Interactions to Infer Manipulation Forces, *Tu-Hoa Pham, Abderrahmane Kheddar, Ammar Qammar, Antonis A. Argyros*



65. A MRF Shape Prior for Facade Parsing With Occlusions, *Mateusz Kozirski, Raghudeep Gadde, Sergey Zagoruyko, Guillaume Obozinski, Renaud Marlet*
66. Probability Occupancy Maps for Occluded Depth Images, *Timur Bagautdinov, François Fleuret, Pascal Fua*
67. Segment Based 3D Object Shape Priors, *Rabeeh Karimi Mahabadi, Christian Häne, Marc Pollefeys*
68. Shape-From-Template in Flatland, *Mathias Gallardo, Daniel Pizarro, Adrien Bartoli, Toby Collins*
69. Understanding Tools: Task-Oriented Object Modeling, Learning and Recognition, *Yixin Zhu, Yibiao Zhao, Song Chun Zhu*
70. Deep Roto-Translation Scattering for Object Classification, *Edouard Oyallon, Stéphane Mallat*
71. Non-Rigid Registration of Images With Geometric and Photometric Deformation by Using Local Affine Fourier-Moment Matching, *Hong-Ren Su, Shang-Hong Lai*
72. Detector Discovery in the Wild: Joint Multiple Instance and Representation Learning, *Judy Hoffman, Deepak Pathak, Trevor Darrell, Kate Saenko*
73. Deeply Learned Face Representations Are Sparse, Selective, and Robust, *Yi Sun, Xiaogang Wang, Xiaoou Tang*
74. Unsupervised Visual Alignment With Similarity Graphs, *Fatemeh Shokrollahi Yancheshmeh, Ke Chen, Joni-Kristian Kämäräinen*
75. Video Anomaly Detection and Localization Using Hierarchical Feature Representation and Gaussian Process Regression, *Kai-Wen Cheng, Yie-Tamg Chen, Wen-Hsien Fang*
76. Inferring 3D Layout of Building Facades From a Single Image, *Jiyan Pan, Martial Hebert, Takeo Kanade*
77. Evaluation of Output Embeddings for Fine-Grained Image Classification, *Zeynep Akata, Scott Reed, Daniel Walter, Honglak Lee, Bernt Schiele*
78. Virtual View Networks for Object Reconstruction, *João Carreira, Abhishek Kar, Shubham Tulsiani, Jitendra Malik*
79. Real-Time Coarse-to-Fine Topologically Preserving Segmentation, *Jian Yao, Marko Boben, Sanja Fidler, Raquel Urtasun*
80. Supervised Mid-Level Features for Word Image Representation, *Albert Gordo*
81. Learning Lightness From Human Judgement on Relative Reflectance, *Takuya Narihira, Michael Maire, Stella X. Yu*
82. Scene Classification With Semantic Fisher Vectors, *Mandar Dixit, Si Chen, Dashan Gao, Nikhil Rasiwasia, Nuno Vasconcelos*
83. Don't Just Listen, Use Your Imagination: Leveraging Visual Common Sense for Non-Visual Tasks, *Xiao Lin, Devi Parikh*
84. Co-Saliency Detection via Looking Deep and Wide, *Dingwen Zhang, Junwei Han, Chao Li, Jingdong Wang*
85. Adopting an Unconstrained Ray Model in Light-Field Cameras for 3D Shape Reconstruction, *Filippo Bergamasco, Andrea Albarelli, Luca Cosmo, Andrea Torsello, Emanuele Rodolà, Daniel Cremers*
86. Towards 3D Object Detection With Bimodal Deep Boltzmann Machines Over RGBD Imagery, *Wei Liu, Rongrong Ji, Shaozi Li*
87. An Active Search Strategy for Efficient Object Class Detection, *Abel Gonzalez-Garcia, Alexander Vezhnevets, Vittorio Ferrari*
88. Geodesic Exponential Kernels: When Curvature and Linearity Conflict, *Aasa Feragen, François Lauze, Søren Hauberg*
89. Transformation-Invariant Convolutional Jungles, *Dmitry Laptev, Joachim M. Buhmann*
90. Exemplar SVMs as Visual Feature Encoders, *Joaquin Zepeda, Patrick Pérez*
91. Object Scene Flow for Autonomous Vehicles, *Moritz Menze, Andreas Geiger*
92. Reflectance Hashing for Material Recognition, *Hang Zhang, Kristin Dana, Ko Nishino*
93. Joint Photo Stream and Blog Post Summarization and Exploration, *Gunhee Kim, Seungwhan Moon, Leonid Sigal*
94. Video Summarization by Learning Submodular Mixtures of Objectives, *Michael Gygli, Helmut Grabner, Luc Van Gool*
95. Building Proteins in a Day: Efficient 3D Molecular Reconstruction, *Marcus A. Brubaker, Ali Punjani, David J. Fleet*
96. Learning Descriptors for Object Recognition and 3D Pose Estimation, *Paul Wohlhart, Vincent Lepetit*



97. Image Partitioning Into Convex Polygons, *Liuyun Duan, Florent Lafarge*
98. Deep Visual-Semantic Alignments for Generating Image Descriptions, *Andrej Karpathy, Li Fei-Fei*
99. Unsupervised Learning of Complex Articulated Kinematic Structures Combining Motion and Skeleton Information, *Hyung Jin Chang, Yiannis Demiris*
100. Elastic Functional Coding of Human Actions: From Vector-Fields to Latent Variables, *Rushil Anirudh, Pavan Turaga, Jingyong Su, Anuj Srivastava*
101. Show and Tell: A Neural Image Caption Generator, *Oriol Vinyals, Alexander Toshev, Samy Bengio, Dumitru Erhan*
102. Descriptor Free Visual Indoor Localization With Line Segments, *Branislav Micusik, Horst Wildenauer*
103. Fixation Bank: Learning to Reweight Fixation Candidates, *Jiaping Zhao, Christian Siagian, Laurent Itti*
104. Deep Networks for Saliency Detection via Local Estimation and Global Search, *Lijun Wang, Huchuan Lu, Xiang Ruan, Ming-Hsuan Yang*
105. Reflection Removal Using Ghosting Cues, *YiChang Shih, Dilip Krishnan, Frédo Durand, William T. Freeman*
106. A Dataset for Movie Description, *Anna Rohrbach, Marcus Rohrbach, Niket Tandon, Bernt Schiele*
107. Fast and Robust Hand Tracking Using Detection-Guided Optimization, *Srinath Sridhar, Franziska Mueller, Antti Oulasvirta, Christian Theobalt*
108. Efficient SDP Inference for Fully-Connected CRFs Based on Low-Rank Decomposition, *Peng Wang, Chunhua Shen, Anton van den Hengel*
109. Discriminative Learning of Iteration-Wise Priors for Blind Deconvolution, *Wangmeng Zuo, Dongwei Ren, Shuhang Gu, Liang Lin, Lei Zhang*
110. Eye Tracking Assisted Extraction of Attentionally Important Objects From Videos, *Karthikeyan Shanmuga Vadivel, Thuyen Ngo, Miguel Eckstein, B.S. Manjunath*
111. Multi-View Feature Engineering and Learning, *Jingming Dong, Nikolaos Karianakis, Damek Davis, Joshua Hernandez, Jonathan Balzer, Stefano Soatto*
112. Self Scaled Regularized Robust Regression, *Yin Wang, Caglayan Dicle, Mario Szafer, Octavia Camps*
113. Simultaneous Feature Learning and Hash Coding With Deep Neural Networks, *Hanjiang Lai, Yan Pan, Ye Liu, Shuicheng Yan*
114. MatchNet: Unifying Feature and Metric Learning for Patch-Based Matching, *Xufeng Han, Thomas Leung, Yangqing Jia, Rahul Sukthankar, Alexander C. Berg*
115. Reconstructing the World\* in Six Days\* (As Captured by the Yahoo 100 Million Image Dataset), *Jared Heinly, Johannes L. Schönberger, Enrique Dunn, Jan-Michael Frahm*
116. Exact Bias Correction and Covariance Estimation for Stereo Vision, *Charles Freundlich, Michael Zavlanos, Philippos Mordohai*
117. Computing Similarity Transformations From Only Image Correspondences, *Chris Sweeney, Laurent Kneip, Tobias Höllerer, Matthew Turk*
118. Image Segmentation in Twenty Questions, *Christian Rupprecht, Loic Peter, Nassir Navab*
119. Interaction Part Mining: A Mid-Level Approach for Fine-Grained Action Recognition, *Yang Zhou, Bingbing Ni, Richang Hong, Meng Wang, Qi Tian*
120. Sparse Projections for High-Dimensional Binary Codes, *Yan Xia, Kaiming He, Pushmeet Kohli, Jian Sun*

**1230-1400 Lunch (Exhibit Hall B)**



## 1230–1400 Doctoral Consortium (Room 103) (by invitation only)

Supported by:



- Abhijit Bendale (*Univ. of Colorado Colorado Springs*)
- Liang-Chieh Chen (*Univ. of California, Los Angeles*)
- Chao-Yeh Chen (*Univ. of Texas at Austin*)
- Xiaochuan Fan (*Univ. of South Carolina*)
- Ying Fu (*Univ. of Tokyo*)
- Georgios Georgiadis (*Univ. of California, Los Angeles*)
- Munawar Hayat (*Univ. of Western Australia*)
- Jared Heinly (*Univ. of North Carolina at Chapel Hill*)
- Joao F. Henriques (*Univ. of Coimbra*)
- Jian-Fang Hu (*Sun Yat-Sen Univ.*)
- Taehwan Kim (*Toyota Technological Institute at Chicago*)
- Zhuwen Li (*National Univ. of Singapore*)
- Yen-Liang Lin (*National Taiwan Univ.*)
- Eslam Mostafa (*Univ. of Louisville*)
- Jaesik Park (*KAIST*)
- Konstantinos Rematas (*KU Leuven*)
- Olga Russakovsky (*Stanford Univ.*)
- Zhiyuan Shi (*Queen Mary Univ. of London*)
- Limin Wang (*Chinese Univ. of Hong Kong*)
- Ziheng Wang (*Rensselaer Polytechnic Institute*)
- Yue Wu (*Rensselaer Polytechnic Institute*)
- Hui Wu (*Univ. of North Carolina at Charlotte*)
- Shijie Xiao (*Nanyang Technological Univ.*)
- Jia Xu (*Univ. of Wisconsin, Madison*)
- Junjie Yan (*Chinese Academy of Sciences*)

- Heng Yang (*Queen Mary Univ. of London*)
- Jian Yao (*Univ. of Toronto*)
- Zhengwu Zhang (*Florida State Univ.*)
- Rui Zhao (*Chinese Univ. of Hong Kong*)
- Liang Zheng (*Tsinghua Univ.*)
- Enliang Zheng (*Univ. of North Carolina at Chapel Hill*)



## 1400–1530 Oral Tue-PM-1: Image & Video Segmentation (Ballrooms A-B)

Papers in this session are also in Poster Session Tue-PM.

**Chairs:** Cristian Sminchisescu (*Lund Univ.*)  
Kyoung Mu Lee (*Seoul National Univ.*)

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

1. Causal Video Object Segmentation From Persistence of Occlusions, *Brian Taylor, Vasily Karasev, Stefano Soatto*
2. Semantic Object Segmentation via Detection in Weakly Labeled Video, *Yu Zhang, Xiaowu Chen, Jia Li, Chen Wang, Changun Xia*
3. Fully Convolutional Networks for Semantic Segmentation, *Jonathan Long, Evan Shelhamer, Trevor Darrell*
4. Shape-Tailored Local Descriptors and Their Application to Segmentation and Tracking, *Naeemullah Khan, Marei Algarni, Anthony Yezzi, Ganesh Sundaramoorthi*
5. Deep Filter Banks for Texture Recognition and Segmentation, *Mircea Cimpoi, Subhansu Maji, Andrea Vedaldi*
6. Active Learning for Structured Probabilistic Models With Histogram Approximation, *Qing Sun, Ankit Laddha, Dhruv Batra*

## 1400–1530 Oral Tue-PM-2: 3D Models & Images (Rooms 302-306)

Papers in this session are also in Poster Session Tue-PM.

**Chairs:** Luc VanGool (*ETH Zurich*)  
Silvio Savarese (*Stanford Univ.*)

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

1. *Picture*: A Probabilistic Programming Language for Scene Perception, *Tejas D. Kulkarni, Pushmeet Kohli, Joshua B. Tenenbaum, Vikash Mansinghka*
2. Rent3D: Floor-Plan Priors for Monocular Layout Estimation, *Chenxi Liu, Alexander G. Schwing, Kaustav Kundu, Raquel Urtasun, Sanja Fidler*
3. The Stitched Puppet: A Graphical Model of 3D Human Shape and Pose, *Silvia Zuffi, Michael J. Black*
4. 3D Shape Estimation From 2D Landmarks: A Convex Relaxation Approach, *XiaoWei Zhou, Spyridon Leonardos, Xiaoyan Hu, Kostas Daniilidis*

5. Holistic 3D Scene Understanding From a Single Geo-Tagged Image, *Shenlong Wang, Sanja Fidler, Raquel Urtasun*
6. Joint SFM and Detection Cues for Monocular 3D Localization in Road Scenes, *Shiyu Song, Manmohan Chandraker*

## 1530–1615 Break (Exhibit Hall B)

## 1530–1800 Exhibits (Exhibit Hall A)

- Same as Monday morning Exhibits (see pg. 11)

## 1530–1800 Demos (Room 108)

- Same as Tuesday morning Demos (see pg. 22)

## 1530–1800 Poster Tue-PM (Exhibit Hall A)

1. Hierarchically-Constrained Optical Flow, *Ryan Kennedy, Camillo J. Taylor*
2. The  $k$ -Support Norm and Convex Envelopes of Cardinality and Rank, *Anders Eriksson, Trung Thanh Pham, Tat-Jun Chin, Ian Reid*
3. Matching Bags of Regions in RGBD images, *Hao Jiang*
4. Recurrent Convolutional Neural Network for Object Recognition, *Ming Liang, Xiaolin Hu*
5. Feedforward Semantic Segmentation With Zoom-Out Features, *Mohammadreza Mostajabi, Payman Yadollahpour, Gregory Shakhnarovich*
6. The Aperture Problem for Refractive Motion, *Tianfan Xue, Hossein Mobahi, Frédo Durand, William T. Freeman*
7. Saliency-Aware Geodesic Video Object Segmentation, *Wenguan Wang, Jianbing Shen, Fatih Porikli*
8. DEEP-CARVING: Discovering Visual Attributes by Carving Deep Neural Nets, *Sukrit Shankar, Vikas K. Garg, Roberto Cipolla*
9. Rent3D: Floor-Plan Priors for Monocular Layout Estimation, *Chenxi Liu, Alexander G. Schwing, Kaustav Kundu, Raquel Urtasun, Sanja Fidler*
10. Learning a Sequential Search for Landmarks, *Saurabh Singh, Derek Hoiem, David Forsyth*



11. Fully Convolutional Networks for Semantic Segmentation, *Jonathan Long, Evan Shelhamer, Trevor Darrell*
12. Deep Correlation for Matching Images and Text, *Fei Yan, Krystian Mikolajczyk*
13. Multi-Objective Convolutional Learning for Face Labeling, *Sifei Liu, Jimei Yang, Chang Huang, Ming-Hsuan Yang*
14. Deep Multiple Instance Learning for Image Classification and Auto-Annotation, *Jiajun Wu, Yanan Yu, Chang Huang, Kai Yu*
15. Multi-Instance Object Segmentation With Occlusion Handling, *Yi-Ting Chen, Xiaokai Liu, Ming-Hsuan Yang*
16. Material Recognition in the Wild With the Materials in Context Database, *Sean Bell, Paul Upchurch, Noah Snavely, Kavita Bala*
17. Understanding Pedestrian Behaviors From Stationary Crowd Groups, *Shuai Yi, Hongsheng Li, Xiaogang Wang*
18. Depth From Focus With Your Mobile Phone, *Supasorn Suwajanakorn, Carlos Hernandez, Steven M. Seitz*
19. Fusion Moves for Correlation Clustering, *Thorsten Beier, Fred A. Hamprecht, Jörg H. Kappes*
20. Second-Order Constrained Parametric Proposals and Sequential Search-Based Structured Prediction for Semantic Segmentation in RGB-D Images, *Dan Banica, Cristian Sminchisescu*
21. Metric Imitation by Manifold Transfer for Efficient Vision Applications, *Dengxin Dai, Till Kroeger, Radu Timofte, Luc Van Gool*
22. The Stitched Puppet: A Graphical Model of 3D Human Shape and Pose, *Silvia Zuffi, Michael J. Black*
23. Scene Labeling With LSTM Recurrent Neural Networks, *Wonmin Byeon, Thomas M. Breuel, Federico Raue, Marcus Lwicky*
24. FAemb: A Function Approximation-Based Embedding Method for Image Retrieval, *Thanh-Toan Do, Quang D. Tran, Ngai-Man Cheung*
25. Automatically Discovering Local Visual Material Attributes, *Gabriel Schwartz, Ko Nishino*
26. Depth Image Enhancement Using Local Tangent Plane Approximations, *Kiyoshi Matsuo, Yoshimitsu Aoki*
27. Video Co-Summarization: Video Summarization by Visual Co-Occurrence, *Wen-Sheng Chu, Yale Song, Alejandro Jaimés*
28. Watch and Learn: Semi-Supervised Learning for Object Detectors From Video, *Ishan Misra, Abhinav Shrivastava, Martial Hebert*
29. Generalized Tensor Total Variation Minimization for Visual Data Recovery, *Xiaoje Guo, Yi Ma*
30. Active Learning for Structured Probabilistic Models With Histogram Approximation, *Qing Sun, Ankit Laddha, Dhruv Batra*
31. Image Parsing With a Wide Range of Classes and Scene-Level Context, *Marian George*
32. Bayesian Sparse Representation for Hyperspectral Image Super Resolution, *Naveed Akhtar, Faisal Shafait, Ajmal Mian*
33. Semantic Object Segmentation via Detection in Weakly Labeled Video, *Yu Zhang, Xiaowu Chen, Jia Li, Chen Wang, Changqun Xia*
34. Learning With Dataset Bias in Latent Subcategory Models, *Dimitris Stamos, Samuele Martelli, Moin Nabi, Andrew McDonald, Vittorio Murino, Massimiliano Pontil*
35. Project-Out Cascaded Regression With an Application to Face Alignment, *Georgios Tzimiropoulos*
36. Image Retrieval Using Scene Graphs, *Justin Johnson, Ranjay Krishna, Michael Stark, Li-Jia Li, David Shamma, Michael Bernstein, Li Fei-Fei*
37. Unifying Holistic and Parts-Based Deformable Model Fitting, *Joan Alabort-i-Medina, Stefanos Zafeiriou*
38. Small Instance Detection by Integer Programming on Object Density Maps, *Zheng Ma, Lei Yu, Antoni B. Chan*
39. Motion Part Regularization: Improving Action Recognition via Trajectory Selection, *Bingbing Ni, Pierre Moulin, Xiaokang Yang, Shuicheng Yan*
40. Multi-Task Deep Visual-Semantic Embedding for Video Thumbnail Selection, *Wu Liu, Tao Mei, Yongdong Zhang, Cherry Che, Jiebo Luo*
41. Fine-Grained Visual Categorization via Multi-Stage Metric Learning, *Qi Qian, Rong Jin, Shenghuo Zhu, Yuanqing Lin*
42. Saturation-Preserving Specular Reflection Separation, *Yuanliu Liu, Zejian Yuan, Nanning Zheng, Yang Wu*
43. Joint SFM and Detection Cues for Monocular 3D Localization in Road Scenes, *Shiyu Song, Manmohan Chandraker*



44. Fisher Vectors Meet Neural Networks: A Hybrid Classification Architecture, *Florent Perronnin, Diane Larlus*
45. UniHIST: A Unified Framework for Image Restoration With Marginal Histogram Constraints, *Xing Mei, Weiming Dong, Bao-Gang Hu, Siwei Lu*
46. Human Action Segmentation With Hierarchical Supervoxel Consistency, *Jiasen Lu, ran Xu, Jason J. Corso*
47. Robust Manhattan Frame Estimation From a Single RGB-D Image, *Bernard Ghanem, Ali Thabet, Juan Carlos Nibbles, Fabian Caba Heilbron*
48. Learning to Segment Under Various Forms of Weak Supervision, *Jia Xu, Alexander G. Schwing, Raquel Urtasun*
49. Fast and Accurate Image Upscaling With Super-Resolution Forests, *Samuel Schuler, Christian Leistner, Horst Bischof*
50. Light Field From Micro-Baseline Image Pair, *Zhoutong Zhang, Yebin Liu, Qionghai Dai*
51. Efficient ConvNet-Based Marker-Less Motion Capture in General Scenes With a Low Number of Cameras, *Ahmed Elhayek, Edison de Aguiar, Arjun Jain, Jonathan Tompson, Leonid Pishchulin, Micha Andriluka, Chris Bregler, Bernt Schiele, Christian Theobalt*
52. Learning Scene-Specific Pedestrian Detectors Without Real Data, *Hironori Hattori, Vishnu Naresh Boddeti, Kris M. Kitani, Takeo Kanade*
53. Deep Filter Banks for Texture Recognition and Segmentation, *Mircea Cimpoi, Subhransu Maji, Andrea Vedaldi*
54. Multiple Random Walkers and Their Application to Image Cosegmentation, *Chulwoo Lee, Won-Dong Jang, Jae-Young Sim, Chang-Su Kim*
55. Beyond the Shortest Path : Unsupervised Domain Adaptation by Sampling Subspaces Along the Spline Flow, *Rui Caseiro, João F. Henriques, Pedro Martins, Jorge Batista*
56. Spherical Embedding of Inlier Silhouette Dissimilarities, *Etai Littwin, Hadar Averbuch-Elor, Daniel Cohen-Or*
57. Semantics-Preserving Hashing for Cross-View Retrieval, *Zijia Lin, Guiguang Ding, Mingqing Hu, Jianmin Wang*
58. Object Proposal by Multi-Branch Hierarchical Segmentation, *Chaoyang Wang, Long Zhao, Shuang Liang, Liqing Zhang, Jinyuan Jia, Yichen Wei*
59. Ambient Occlusion via Compressive Visibility Estimation, *Wei Yang, Yu Ji, Haiting Lin, Yang Yang, Sing Bing Kang, Jingyi Yu*
60. Shape-Tailored Local Descriptors and Their Application to Segmentation and Tracking, *Naeemullah Khan, Marei Algarni, Anthony Yezzi, Ganesh Sundaramoorthi*
61. Scalable Object Detection by Filter Compression With Regularized Sparse Coding, *Ting-Hsuan Chao, Yen-Liang Lin, Yin-Hsi Kuo, Winston H. Hsu*
62. An Improved Deep Learning Architecture for Person Re-Identification, *Ejaz Ahmed, Michael Jones, Tim K. Marks*
63. Understanding Classifier Errors by Examining Influential Neighbors, *Mayank Kabra, Alice Robie, Kristin Branson*
64. Riemannian Coding and Dictionary Learning: Kernels to the Rescue, *Mehrtash Harandi, Mathieu Salzmann*
65. Scalable Structure From Motion for Densely Sampled Videos, *Benjamin Resch, Hendrik P. A. Lensch, Oliver Wang, Marc Pollefeys, Alexander Sorkine-Hornung*
66. Parsing Occluded People by Flexible Compositions, *Xianjie Chen, Alan L. Yuille*
67. Joint Calibration of Ensemble of Exemplar SVMs, *Davide Modolo, Alexander Vezhnevets, Olga Russakovsky, Vittorio Ferrari*
68. Holistic 3D Scene Understanding From a Single Geo-Tagged Image, *Shenlong Wang, Sanja Fidler, Raquel Urtasun*
69. A Large-Scale Car Dataset for Fine-Grained Categorization and Verification, *Linjie Yang, Ping Luo, Chen Change Loy, Xiaoou Tang*
70. DeepContour: A Deep Convolutional Feature Learned by Positive-Sharing Loss for Contour Detection, *Wei Shen, Xinggang Wang, Yan Wang, Xiang Bai, Zhijiang Zhang*
71. Convolutional Feature Masking for Joint Object and Stuff Segmentation, *Jifeng Dai, Kaiming He, Jian Sun*
72. A Fixed Viewpoint Approach for Dense Reconstruction of Transparent Objects, *Kai Han, Kwan-Yee K. Wong, Miaomiao Liu*
73. Low-Level Vision by Consensus in a Spatial Hierarchy of Regions, *Ayan Chakrabarti, Ying Xiong, Steven J. Gortler, Todd Zickler*
74. Line Drawing Interpretation in a Multi-View Context, *Jean-Dominique Favreau, Florent Lafarge, Adrien Bousseau*



75. Toward User-Specific Tracking by Detection of Human Shapes in Multi-Cameras, *Chun-Hao Huang, Edmond Boyer, Bibiana do Canto Angonese, Nassir Navab, Slobodan Ilic*
76. Intra-Frame Deblurring by Leveraging Inter-Frame Camera Motion, *Haichao Zhang, Jianchao Yang*
77. Salient Object Subitizing, *Jianming Zhang, Shugao Ma, Mehrnoosh Sameki, Stan Sclaroff, Margrit Betke, Zhe Lin, Xiaohui Shen, Brian Price, Radomir Měch*
78. Hierarchical-PEP Model for Real-World Face Recognition, *Haoxiang Li, Gang Hua*
79. The Common Self-Polar Triangle of Concentric Circles and Its Application to Camera Calibration, *Haifei Huang, Hui Zhang, Yiu-ming Cheung*
80. Taking a Deeper Look at Pedestrians, *Jan Hosang, Mohamed Omran, Rodrigo Benenson, Bernt Schiele*
81. Learning to Segment Moving Objects in Videos, *Katerina Fragkiadaki, Pablo Arbeláez, Panna Felsen, Jitendra Malik*
82. GMMCP Tracker: Globally Optimal Generalized Maximum Multi Clique Problem for Multiple Object Tracking, *Afshin Dehghan, Shayan Modiri Assari, Mubarak Shah*
83. Learning Graph Structure for Multi-Label Image Classification via Clique Generation, *Mingkui Tan, Qinfeng Shi, Anton van den Hengel, Chunhua Shen, Junbin Gao, Fuyuan Hu, Zhen Zhang*
84. Matrix Completion for Resolving Label Ambiguity, *Ching-Hui Chen, Vishal M. Patel, Rama Chellappa*
85. Video Magnification in Presence of Large Motions, *Mohamed Elgharib, Mohamed Hefeeda, Frédo Durand, William T. Freeman*
86. Flying Objects Detection From a Single Moving Camera, *Artem Rozantsev, Vincent Lepetit, Pascal Fua*
87. Line-Based Multi-Label Energy Optimization for Fisheye Image Rectification and Calibration, *Mi Zhang, Jian Yao, Menghan Xia, Kai Li, Yi Zhang, Yaping Liu*
88. Adaptive Eye-Camera Calibration for Head-Worn Devices, *David Perra, Rohit Kumar Gupta, Jan-Michael Frahm*
89. Modeling Object Appearance Using Context-Conditioned Component Analysis, *Daniyar Turmukhambetov, Neill D.F. Campbell, Simon J.D. Prince, Jan Kautz*
90. Displets: Resolving Stereo Ambiguities Using Object Knowledge, *Fatma Güney, Andreas Geiger*
91. Time-to-Contact From Image Intensity, *Yukitoshi Watanabe, Fumihiko Sakaue, Jun Sato*
92. Transferring a Semantic Representation for Person Re-Identification and Search, *Zhiyuan Shi, Timothy M. Hospedales, Tao Xiang*
93. Robust Video Segment Proposals With Painless Occlusion Handling, *Zhengyang Wu, Fuxin Li, Rahul Sukthankar, James M. Rehg*
94. Face Alignment Using Cascade Gaussian Process Regression Trees, *Donghoon Lee, Hyunsin Park, Chang D. Yoo*
95. Regularizing Max-Margin Exemplars by Reconstruction and Generative Models, *Jose C. Rubio, Björn Ommer*
96. A Fast Algorithm for Elastic Shape Distances Between Closed Planar Curves, *Günay Doğan, Javier Bernal, Charles R. Hagwood*
97. Reflection Removal for In-Vehicle Black Box Videos, *Christian Simon, In Kyu Park*
98. Tree Quantization for Large-Scale Similarity Search and Classification, *Artem Babenko, Victor Lempitsky*
99. Integrating Parametric and Non-Parametric Models For Scene Labeling, *Bing Shuai, Gang Wang, Zhen Zuo, Bing Wang, Lifan Zhao*
100. Mining Semantic Affordances of Visual Object Categories, *Yu-Wei Chao, Zhan Wang, Rada Mihalcea, Jia Deng*
101. Causal Video Object Segmentation From Persistence of Occlusions, *Brian Taylor, Vasilij Karasev, Stefano Soatto*
102. Multiple Instance Learning for Soft Bags via Top Instances, *Weixin Li, Nuno Vasconcelos*
103. Multiclass Semantic Video Segmentation With Object-Level Active Inference, *Buyu Liu, Xuming He*
104. Effective Face Frontalization in Unconstrained Images, *Tal Hassner, Shai Harel, Eran Paz, Roei Enbar*
105. Action Recognition With Trajectory-Pooled Deep-Convolutional Descriptors, *Limin Wang, Yu Qiao, Xiaoou Tang*
106. Weakly Supervised Localization of Novel Objects Using Appearance Transfer, *Mrigank Rochan, Yang Wang*
107. First-Person Pose Recognition Using Egocentric Workspaces, *Grégory Rogez, James S. Supančič III, Deva Ramanan*



108. Simultaneous Time-of-Flight Sensing and Photometric Stereo With a Single ToF Sensor, *Changpeng Ti, Ruigang Yang, James Davis, Zhigeng Pan*
109. Active Learning and Discovery of Object Categories in the Presence of Unnameable Instances, *Christoph Kuding, Alexander Freytag, Erik Rodner, Paul Bodesheim, Joachim Denzler*
110. Learning to Compare Image Patches via Convolutional Neural Networks, *Sergey Zagoruyko, Nikos Komodakis*
111. Watch-n-Patch: Unsupervised Understanding of Actions and Relations, *Chenxia Wu, Jiemi Zhang, Silvio Savarese, Ashutosh Saxena*
112. Optimal Graph Learning With Partial Tags and Multiple Features for Image and Video Annotation, *Lianli Gao, Jingkuan Song, Feiping Nie, Yan Yan, Nicu Sebe, Heng Tao Shen*
113. DeepEdge: A Multi-Scale Bifurcated Deep Network for Top-Down Contour Detection, *Gedas Bertasius, Jianbo Shi, Lorenzo Torresani*
114. Picture: A Probabilistic Programming Language for Scene Perception, *Tejas D. Kulkarni, Pushmeet Kohli, Joshua B. Tenenbaum, Vikash Mansinghka*
115. Exploiting Uncertainty in Regression Forests for Accurate Camera Relocalization, *Julien Valentin, Matthias Niesner, Jamie Shotton, Andrew Fitzgibbon, Shahram Izadi, Philip H. S. Torr*
116. Fusing Subcategory Probabilities for Texture Classification, *Yang Song, Weidong Cai, Qing Li, Fan Zhang, David Dagan Feng, Heng Huang*
117. Video Event Recognition With Deep Hierarchical Context Model, *Xiaoyang Wang, Qiang Ji*
118. Object-Based RGBD Image Co-Segmentation With Mutex Constraint, *Huazhu Fu, Dong Xu, Stephen Lin, Jiang Liu*
119. Associating Neural Word Embeddings With Deep Image Representations Using Fisher Vectors, *Benjamin Klein, Guy Lev, Gil Sadeh, Lior Wolf*
120. 3D Shape Estimation From 2D Landmarks: A Convex Relaxation Approach, *Xiaowei Zhou, Spyridon Leonardos, Xiaoyan Hu, Kostas Daniilidis*

**1800–2100 Banquet Dinner**  
(Sheraton Grand Ballroom)



# Wednesday, June 10

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

**0730–0830 Breakfast** (Exhibit Hall B)

## **0830–1000 Oral Wed-AM-1: Action & Event Recognition** (Ballrooms A-B)

Papers in this session are also in Poster Session Wed-PM.

**Chairs:** Ivan Laptev (*INRIA*)  
Shih-Fu Chang (*Columbia Univ.*)

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

1. How Many Bits Does it Take for a Stimulus to Be Salient?,  
*Sayed Hossein Khatoonabadi, Nuno Vasconcelos, Ivan V. Bajic, Yufeng Shan*
2. Deeply Learned Attributes for Crowded Scene Understanding, *Jing Shao, Kai Kang, Chen Change Loy, Xiaogang Wang*
3. Joint Inference of Groups, Events and Human Roles in Aerial Videos, *Tianmin Shu, Dan Xie, Brandon Rothrock, Sinisa Todorovic, Song Chun Zhu*
4. Modeling Video Evolution for Action Recognition, *Basura Fernando, Efstratios Gavves, José Oramas M., Amir Ghodrati, Tinne Tuytelaars*
5. Space-Time Tree Ensemble for Action Recognition, *Shugao Ma, Leonid Sigal, Stan Sclaroff*
6. Social Saliency Prediction, *Hyun Soo Park, Jianbo Shi*

## **0830–1000 Oral Wed-AM-2: Computational Photography** (Rooms 302-306)

Papers in this session are also in Poster Session Wed-PM.

**Chairs:** James Hays (*Brown Univ.*)  
Shai Avidan (*Tel-Aviv Univ.*)

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

1. Visual Vibrometry: Estimating Material Properties From Small Motion in Video, *Abe Davis, Katherine L. Bouman, Justin G. Chen, Michael Rubinstein, Frédo Durand, William T. Freeman*

2. Recovering Inner Slices of Translucent Objects by Multi-Frequency Illumination, *Kenichiro Tanaka, Yasuhiro Mukaigawa, Hiroyuki Kubo, Yasuyuki Matsushita, Yasushi Yagi*
3. Fast Bilateral-Space Stereo for Synthetic Defocus, *Jonathan T. Barron, Andrew Adams, YiChang Shih, Carlos Hernández*
4. Simultaneous Video Defogging and Stereo Reconstruction, *Zhuwen Li, Ping Tan, Robby T. Tan, Danping Zou, Steven Zhiying Zhou, Loong-Fah Cheong*
5. One-Day Outdoor Photometric Stereo via Skylight Estimation, *Jiyoung Jung, Joon-Young Lee, In So Kweon*

**1000–1030 Break** (Exhibit Hall B)

## **1030–1225 Plenary Session** (Ballrooms A-B)

- **Plenary Talk:** What's Wrong with Deep Learning? *Yann LeCun (Facebook and NYU)*

**Abstract:** Deep learning methods have had a profound impact on a number of areas in recent years, including natural image understanding and speech recognition. Other areas seem on the verge of being similarly impacted, notably natural language processing, biomedical image analysis, and the analysis of sequential signals in a variety of application domains. But deep learning systems, as they exist today, have many limitations.

First, they lack mechanisms for reasoning, search, and inference. Complex and/or ambiguous inputs require deliberate reasoning to arrive at a consistent interpretation. Producing structured outputs, such as a long text, or a label map for image segmentation, require sophisticated search and inference algorithms to satisfy complex sets of constraints.

Second, they lack short-term memory. Many tasks in natural language understanding, such as question-answering, require a way to temporarily store isolated facts. Correctly interpreting events in a video and being able to answer questions about it requires remembering abstract representations of what happens in the video. Deep learning systems, including recurrent nets, are notoriously inefficient at storing temporary memories. Lastly, they lack the ability to perform unsupervised learning. Animals and humans learn most of the structure of the perceptual world in an unsupervised manner. Future



progress in computer vision will require breakthroughs in unsupervised learning, particularly for video understanding, but what principles should unsupervised learning be based on?

- **Plenary Talk:** Reverse Engineering the Human Visual System, *Jack Gallant (UC Berkeley)*

**Abstract:** The human brain is the most sophisticated image processing system known, capable of impressive feats of recognition and discrimination under challenging natural conditions. Reverse-engineering the brain might enable us to design artificial systems with the same capabilities. My laboratory uses a data-driven system identification approach to tackle this reverse-engineering problem. Our approach consists of four broad stages. First, we use functional MRI to measure brain activity while people watch naturalistic movies. We divide these data into two parts, one use to fit models and one for testing model predictions. Second, we use a system identification framework (based on multiple linearizing feature spaces) to model activity measured at each point in the brain. Third, we inspect the most accurate models to understand how the brain represents low-, mid- and high-level information in the movies. Finally, we use the estimated models to decode brain activity, reconstructing the structural and semantic content in the movies. Any effort to reverse-engineer the brain is inevitably limited by the spatial and temporal resolution of brain measurements, and at this time the resolution of human brain measurements is relatively poor. Still, as measurement technology progresses this framework could inform development of biologically-inspired computer vision systems, and it could aid in development of practical new brain reading technologies.

**1230–1400 Lunch** (Exhibit Hall B)



## 1400-1530 Oral Wed-PM-1: Learning & Matching Local Features; Visualization (Ballrooms A-B)

Papers in this session are also in Poster Session Wed-PM.

**Chairs:** Vincent Lepetit (*Graz Univ. of Technology*)  
Piotr Dollar (*Facebook AI Research*)

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

1. Domain-Size Pooling in Local Descriptors: DSP-SIFT, *Jingming Dong, Stefano Soatto*
2. Learning Deep Representations for Ground-to-Aerial Geolocalization, *Tsung-Yi Lin, Yin Cui, Serge Belongie, James Hays*
3. Understanding Deep Image Representations by Inverting Them, *Aravindh Mahendran, Andrea Vedaldi*
4. Situational Object Boundary Detection, *Jasper R. R. Uijlings, Vittorio Ferrari*
5. Fast 2D Border Ownership Assignment, *Ching Teo, Cornelia Fermüller, Yiannis Aloimonos*
6. A Flexible Tensor Block Coordinate Ascent Scheme for Hypergraph Matching, *Quynh Nguyen, Antoine Gautier, Matthias Hein*

## 1400-1530 Oral Wed-PM-2: Image & Video Restoration/Processing (Rooms 302-306)

Papers in this session are also in Poster Session Wed-PM.

**Chairs:** Michal Irani (*Weizmann Institute of Science*)  
Eli Shechtman (*Adobe Research*)

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

1. Generalized Video Deblurring for Dynamic Scenes, *Tae Hyun Kim, Kyoung Mu Lee*
2. Approximate Nearest Neighbor Fields in Video, *Nir Ben-Zrihem, Lihi Zelnik-Manor*
3. Single Image Super-Resolution From Transformed Self-Exemplars, *Jia-Bin Huang, Abhishek Singh, Narendra Ahuja*
4.  $\zeta_0$ TV: A New Method for Image Restoration in the Presence of Impulse Noise, *Ganzhao Yuan, Bernard Ghanem*
5. On Learning Optimized Reaction Diffusion Processes for Effective Image Restoration, *Yunjin Chen, Wei Yu, Thomas Pock*

6. Fast and Flexible Convolutional Sparse Coding, *Felix Heide, Wolfgang Heidrich, Gordon Wetzstein*

## 1530-1615 Break (Exhibit Hall B)

## 1530-1800 Exhibits (Exhibit Hall A)

- Same as Tuesday morning Exhibits (see pg. 11)

## 1530-1800 Demos (Room 108)

- Real-Time Articulated Hand Tracking and its Application to 3D Virtual Sculpting, *Hwasup Lim, Sungkuk Chun, Junha Roh, Sang Chul Ahn (Korea Institute of Science and Technology)*
- Real-Time 3D Pose Recognition From a Chest-mounted RGB-D Camera, *Gregory Rogez, James S. Supancic, Deva Ramana (Univ. of California, Irvine)*
- Patient Activity Recognition in a Hospital Room Using Kinect V2, *Liang Liu, Melody Yin, Sanjay Mehrotra (North-western Univ.)*

## 1530-1800 Poster Wed-PM (Exhibit Hall A)

1. 3D All The Way: Semantic Segmentation of Urban Scenes From Start to End in 3D, *Andelo Martinović, Jan Knopp, Hayko Riemenschneider, Luc Van Gool*
2. Fast Bilateral-Space Stereo for Synthetic Defocus, *Jonathan T. Barron, Andrew Adams, YiChang Shih, Carlos Hernández*
3. Large-Scale and Drift-Free Surface Reconstruction Using Online Subvolume Registration, *Nicola Fioraio, Jonathan Taylor, Andrew Fitzgibbon, Luigi Di Stefano, Shahram Izadi*
4. Fast Randomized Singular Value Thresholding for Nuclear Norm Minimization, *Tae-Hyun Oh, Yasuyuki Matsushita, Yu-Wing Tai, In So Kweon*
5. LMI-Based 2D-3D Registration: From Uncalibrated Images to Euclidean Scene, *Danda Pani Paudel, Adlane Haded, Cédric Demonceaux, Pascal Vasseur*
6. Clique-Graph Matching by Preserving Global & Local Structure, *Wei-Zhi Nie, An-An Liu, Zan Gao, Yu-Ting Su*
7. Appearance-Based Gaze Estimation in the Wild, *Xucong Zhang, Yusuke Sugano, Mario Fritz, Andreas Bulling*
8. One-Day Outdoor Photometric Stereo via Skylight Estimation, *Jiyoung Jung, Joon-Young Lee, In So Kweon*



9. A New Retraction for Accelerating the Riemannian Three-Factor Low-Rank Matrix Completion Algorithm, *Zhizhong Li, Deli Zhao, Zhouchen Lin, Edward Y. Chang*
10. Heteroscedastic Max-Min Distance Analysis, *Bing Su, Xiaoqing Ding, Changsong Liu, Ying Wu*
11. Sparse Composite Quantization, *Ting Zhang, Guo-Jun Qi, Jinhui Tang, Jingdong Wang*
12. Sparse Representation Classification With Manifold Constraints Transfer, *Baochang Zhang, Alessandro Perina, Vittorio Murino, Alessio Del Bue*
13. CIDEr: Consensus-Based Image Description Evaluation, *Ramakrishna Vedantam, C. Lawrence Zitnick, Devi Parikh*
14. Joint Inference of Groups, Events and Human Roles in Aerial Videos, *Tianmin Shu, Dan Xie, Brandon Rothrock, Sinisa Todorovic, Song Chun Zhu*
15. Photometric Stereo With Near Point Lighting: A Solution by Mesh Deformation, *Wuyuan Xie, Chengkai Dai, Charlie C. L. Wang*
16. Efficient Label Collection for Unlabeled Image Datasets, *Maggie Wigness, Bruce A. Draper, J. Ross Beveridge*
17. Separating Objects and Clutter in Indoor Scenes, *Salman H. Khan, Xuming He, Mohammed Bennis, Ferdous Sohel, Roberto Togneri*
18. FaLRR: A Fast Low Rank Representation Solver, *Shijie Xiao, Wen Li, Dong Xu, Dacheng Tao*
19. Simulating Makeup Through Physics-Based Manipulation of Intrinsic Image Layers, *Chen Li, Kun Zhou, Stephen Lin*
20. Correlation Filters With Limited Boundaries, *Hamed Kiani Galoogahi, Terence Sim, Simon Lucey*
21. Shape-Based Automatic Detection of a Large Number of 3D Facial Landmarks, *Syed Zulqarnain Gilani, Faisal Shafait, Ajmal Mian*
22. Material Classification With Thermal Imagery, *Philipp Saponaro, Scott Sorensen, Abhishek Kolagunda, Chandra Kamthamattu*
23. Deeply Learned Attributes for Crowded Scene Understanding, *Jing Shao, Kai Kang, Chen Change Loy, Xiaogang Wang*
24. Learning To Look Up: Realtime Monocular Gaze Correction Using Machine Learning, *Daniil Kononenko, Victor Lempitsky*
25. Background Subtraction via Generalized Fused Lasso Foreground Modeling, *Bo Xin, Yuan Tian, Yizhou Wang, Wen Gao*
26. Mirror, Mirror on the Wall, Tell Me, Is the Error Small?, *Heng Yang, Ioannis Patras*
27. Beyond Short Snippets: Deep Networks for Video Classification, *Joe Yue-Hei Ng, Matthew Hausknecht, Sudheendra Vijayanarasimhan, Oriol Vinyals, Rajat Monga, George Toderici*
28. segDeepM: Exploiting Segmentation and Context in Deep Neural Networks for Object Detection, *Yukun Zhu, Raquel Urtasun, Ruslan Salakhutdinov, Sanja Fidler*
29. Situational Object Boundary Detection, *Jasper R. R. Uijlings, Vittorio Ferrari*
30. Real-Time 3D Head Pose and Facial Landmark Estimation From Depth Images Using Triangular Surface Patch Features, *Chavdar Papazov, Tim K. Marks, Michael Jones*
31. Aligning 3D Models to RGB-D Images of Cluttered Scenes, *Saurabh Gupta, Pablo Arbeláez, Ross Girshick, Jitendra Malik*
32. A Stable Multi-Scale Kernel for Topological Machine Learning, *Jan Reininghaus, Stefan Huber, Ulrich Bauer, Roland Kwitt*
33. The Treasure Beneath Convolutional Layers: Cross-Convolutional-Layer Pooling for Image Classification, *Lingqiao Liu, Chunhua Shen, Anton van den Hengel*
34. Face Video Retrieval With Image Query via Hashing Across Euclidean Space and Riemannian Manifold, *Yan Li, Ruiping Wang, Zhiwu Huang, Shiguang Shan, Xilin Chen*
35. EgoSampling: Fast-Forward and Stereo for Egocentric Videos, *Yair Poleg, Tavi Halperin, Chetan Arora, Shmuel Peleg*
36. Social Saliency Prediction, *Hyun Soo Park, Jianbo Shi*
37. Beyond Principal Components: Deep Boltzmann Machines for Face Modeling, *Chi Nhan Duong, Khoa Luu, Kha Gia Quach, Tien D. Bui*
38. Statistical Inference Models for Image Datasets With Systematic Variations, *Won Hwa Kim, Barbara B. Bendlin, Moo K. Chung, Sterling C. Johnson, Vikas Singh*
39. Beyond Frontal Faces: Improving Person Recognition Using Multiple Cues, *Ning Zhang, Manohar Paluri, Yaniv Taigman, Rob Fergus, Lubomir Bourdev*



40. Superpixel-Based Video Object Segmentation Using Perceptual Organization and Location Prior, *Daniela Giordano, Francesca Murabito, Simone Palazzo, Concetto Spampinato*
41. Robust Image Filtering Using Joint Static and Dynamic Guidance, *Bumsub Ham, Minsu Cho, Jean Ponce*
42. Solving Multiple Square Jigsaw Puzzles With Missing Pieces, *Genady Paikin, Ayellet Tal*
43. A Dynamic Convolutional Layer for Short Range Weather Prediction, *Benjamin Klein, Lior Wolf, Yehuda Afek*
44. SWIFT: Sparse Withdrawal of Inliers in a First Trial, *Maryam Jaber, Marianna Pensky, Hassan Foroosh*
45. VIP: Finding Important People in Images, *Clint Solomon Mathialagan, Andrew C. Gallagher, Dhruv Batra*
46. Dataset Fingerprints: Exploring Image Collections Through Data Mining, *Konstantinos Rematas, Basura Fernando, Frank Dellaert, Tinne Tuytelaars*
47. Transport-Based Single Frame Super Resolution of Very Low Resolution Face Images, *Soheil Kolouri, Gustavo K. Rohde*
48. 3D Reconstruction in the Presence of Glasses by Acoustic and Stereo Fusion, *Mao Ye, Yu Zhang, Ruigang Yang, Dinesh Manocha*
49. Deep Sparse Representation for Robust Image Registration, *Yeqing Li, Chen Chen, Fei Yang, Junzhou Huang*
50. Real-Time Part-Based Visual Tracking via Adaptive Correlation Filters, *Ting Liu, Gang Wang, Qingxiong Yang*
51. Beyond Spatial Pooling: Fine-Grained Representation Learning in Multiple Domains, *Chi Li, Austin Reiter, Gregory D. Hager*
52.  $\mathcal{H}^2$ -Search for Structured Prediction in Computer Vision, *Michael Lam, Janardhan Rao Doppa, Sinisa Todorovic, Thomas G. Dietterich*
53. Revisiting Kernelized Locality-Sensitive Hashing for Improved Large-Scale Image Retrieval, *Ke Jiang, Qichao Que, Brian Kulis*
54. High-Speed Hyperspectral Video Acquisition With a Dual-Camera Architecture, *Lizhi Wang, Zhiwei Xiong, Dahua Gao, Guangming Shi, Wenjun Zeng, Feng Wu*
55. More About VLAD: A Leap From Euclidean to Riemannian Manifolds, *Masoud Faraki, Mehrtash T. Harandi, Fatih Porikli*
56. Camera Intrinsic Blur Kernel Estimation: A Reliable Framework, *Ali Moseleh, Paul Green, Emmanuel Onzon, Isabelle Bégin, J.M. Pierre Langlois*
57. Classifier Learning With Hidden Information, *Ziheng Wang, Qiang Ji*
58. Single Target Tracking Using Adaptive Clustered Decision Trees and Dynamic Multi-Level Appearance Models, *Jingjing Xiao, Rustam Stolkin, Aleš Leonardis*
59. Simultaneous Video Defogging and Stereo Reconstruction, *Zhuwen Li, Ping Tan, Robby T. Tan, Danping Zou, Steven Zhiying Zhou, Loong-Fah Cheong*
60. Face Alignment by Coarse-to-Fine Shape Searching, *Shizhan Zhu, Cheng Li, Chen Change Loy, Xiaoou Tang*
61. Learning Deep Representations for Ground-to-Aerial Geolocalization, *Tsung-Yi Lin, Yin Cui, Serge Belongie, James Hays*
62. Unsupervised Simultaneous Orthogonal Basis Clustering Feature Selection, *Dongyoon Han, Junmo Kim*
63. Space-Time Tree Ensemble for Action Recognition, *Shugao Ma, Leonid Sigal, Stan Sclaroff*
64. Subgraph Decomposition for Multi-Target Tracking, *Siyu Tang, Bjørn Andres, Miya Haylo Andriluka, Bernt Schiele*
65. Understanding Image Structure via Hierarchical Shape Parsing, *Xian-Ming Liu, Rongrong Ji, Changhu Wang, Wei Liu, Bineng Zhong, Thomas S. Huang*
66. Coarse-To-Fine Region Selection and Matching, *Yanchao Yang, Zhaojin Lu, Ganesh Sundaramoorthi*
67. Label Consistent Quadratic Surrogate Model for Visual Saliency Prediction, *Yan Luo, Yongkang Wong, Qi Zhao*
68. Subgraph Matching Using Compactness Prior for Robust Feature Correspondence, *Yumin Suh, Kamil Adamczewski, Kyoung Mu Lee*
69. Pedestrian Detection Aided by Deep Learning Semantic Tasks, *Yonglong Tian, Ping Luo, Xiaogang Wang, Xiaoou Tang*
70. Multihypothesis Trajectory Analysis for Robust Visual Tracking, *Dae-Youn Lee, Jae-Young Sim, Chang-Su Kim*
71. Domain-Size Pooling in Local Descriptors: DSP-SIFT, *Jingming Dong, Stefano Soatto*



72. Object Detection by Labeling Superpixels, *Junjie Yan, Yinan Yu, Xiangyu Zhu, Zhen Lei, Stan Z. Li*
73. Fast 2D Border Ownership Assignment, *Ching Teo, Cornelia Fermüller, Yiannis Aloimonos*
74. From Single Image Query to Detailed 3D Reconstruction, *Johannes L. Schönberger, Filip Radenović, Ondrej Chum, Jan-Michael Frahm*
75. Fast and Flexible Convolutional Sparse Coding, *Felix Heide, Wolfgang Heidrich, Gordon Wetzstein*
76. Iteratively Reweighted Graph Cut for Multi-Label MRFs With Non-Convex Priors, *Thalaiyasingam Ajanthan, Richard Hartley, Mathieu Salzmann, Hongdong Li*
77. Pairwise Geometric Matching for Large-Scale Object Retrieval, *Xinchao Li, Martha Larson, Alan Hanjalic*
78. Deep Convolutional Neural Fields for Depth Estimation From a Single Image, *Fayao Liu, Chunhua Shen, Guosheng Lin*
79. Data-Driven Sparsity-Based Restoration of JPEG-Compressed Images in Dual Transform-Pixel Domain, *Xianming Liu, Xiaolin Wu, Jiantao Zhou, Debin Zhao*
80. TVSum: Summarizing Web Videos Using Titles, *Yale Song, Jordi Vallmitjana, Amanda Stent, Alejandro Jaimes*
81. Understanding Deep Image Representations by Inverting Them, *Aravindh Mahendran, Andrea Vedaldi*
82. Single Image Super-Resolution From Transformed Self-Exemplars, *Jia-Bin Huang, Abhishek Singh, Narendra Ahuja*
83. Constrained Planar Cuts - Object Partitioning for Point Clouds, *Markus Schoeler, Jeremie Papon, Florentin Wörgötter*
84. A Weighted Sparse Coding Framework for Saliency Detection, *Nianyi Li, Bilin Sun, Jingyi Yu*
85. Handling Motion Blur in Multi-Frame Super-Resolution, *Ziyang Ma, Renjie Liao, Xin Tao, Li Xu, Jiaya Jia, Enhua Wu*
86. Approximate Nearest Neighbor Fields in Video, *Nir Ben-Zrihem, Lili Zelnik-Manor*
87. Inverting RANSAC: Global Model Detection via Inlier Rate Estimation, *Roei Litman, Simon Korman, Alexander Bronstein, Shai Avidan*
88. Robust Multi-Image Based Blind Face Hallucination, *Yonggang Jin, Christos-Savvas Bouganis*
89. On Learning Optimized Reaction Diffusion Processes for Effective Image Restoration, *Yunjin Chen, Wei Yu, Thomas Pock*
90. A Flexible Tensor Block Coordinate Ascent Scheme for Hypergraph Matching, *Quynh Nguyen, Antoine Gautier, Matthias Hein*
91. TILDE: A Temporally Invariant Learned DETector, *Yannick Verdie, Kwang Yi, Pascal Fua, Vincent Lepetit*
92. A Maximum Entropy Feature Descriptor for Age Invariant Face Recognition, *Dihong Gong, Zhifeng Li, Dacheng Tao, Jianzhuang Liu, Xuelong Li*
93. Sense Discovery via Co-Clustering on Images and Text, *Xinlei Chen, Alan Ritter, Abhinav Gupta, Tom Mitchell*
94. An Approximate Shading Model for Object Relighting, *Zicheng Liao, Kevin Karsch, David Forsyth*
95. Deep Domain Adaptation for Describing People Based on Fine-Grained Clothing Attributes, *Qiang Chen, Junshi Huang, Rogerio Feris, Lisa M. Brown, Jian Dong, Shuicheng Yan*
96. A Convolutional Neural Network Cascade for Face Detection, *Haoxiang Li, Zhe Lin, Xiaohui Shen, Jonathan Brandt, Gang Hua*
97. Visual Vibrometry: Estimating Material Properties From Small Motion in Video, *Abe Davis, Katherine L. Bouman, Justin G. Chen, Michael Rubinstein, Frédo Durand, William T. Freeman*
98. Jointly Learning Heterogeneous Features for RGB-D Activity Recognition, *Jian-Fang Hu, Wei-Shi Zheng, Jianhuang Lai, Jianguo Zhang*
99. Convolutional Neural Networks at Constrained Time Cost, *Kaiming He, Jian Sun*
100. Fine-Grained Histopathological Image Analysis via Robust Segmentation and Large-Scale Retrieval, *Xiaofan Zhang, Hai Su, Lin Yang, Shaoting Zhang*
101. L<sub>0</sub>TV: A New Method for Image Restoration in the Presence of Impulse Noise, *Ganzhao Yuan, Bernard Ghanem*
102. Modeling Video Evolution for Action Recognition, *Basura Fernando, Efstratios Gafves, José Oramas M., Amir Ghodrati, Tinne Tuytelaars*
103. Long-Term Correlation Tracking, *Chao Ma, Xiaokang Yang, Chongyang Zhang, Ming-Hsuan Yang*



104. Joint Tracking and Segmentation of Multiple Targets,  
*Anton Milan, Laura Leal-Taixé, Konrad Schindler, Ian Reid*
105. RGBD-Fusion: Real-Time High Precision Depth Recovery,  
*Roy Or - El, Guy Rosman, Aaron Wetzler, Ron Kimmel,  
Alfred M. Bruckstein*
106. Modeling Deformable Gradient Compositions for Single-  
Image Super-Resolution, *Yu Zhu, Yanning Zhang, Boyan  
Bonev, Alan L. Yuille*
107. Generalized Video Deblurring for Dynamic Scenes, *Tae  
Hyun Kim, Kyoung Mu Lee*
108. Active Pictorial Structures, *Epameinondas Antonakos,  
Joan Alabort-i-Medina, Stefanos Zafeiriou*
109. Ego-Surfing First-Person Videos, *Ryo Yonetani, Kris M.  
Kitani, Yoichi Sato*
110. Visual Saliency Based on Multiscale Deep Features,  
*Guanbin Li, Yizhou Yu*
111. Recovering Inner Slices of Translucent Objects by Multi-  
Frequency Illumination, *Kenichiro Tanaka, Yasuhiro  
Mukaigawa, Hiroyuki Kubo, Yasuyuki Matsushita, Yasushi  
Yagi*
112. Local High-Order Regularization on Data Manifolds,  
*Kwang In Kim, James Tompkin, Hanspeter Pfister, Christian  
Theobalt*
113. Fine-Grained Classification of Pedestrians in Video:  
Benchmark and State of the Art, *David Hall, Pietro Perona*
114. Curriculum Learning of Multiple Tasks, *Anastasia Pentina,  
Viktoria Sharmanska, Christoph H. Lampert*
115. How Many Bits Does it Take for a Stimulus to Be Salient?,  
*Sayed Hossein Khatoonabadi, Nuno Vasconcelos, Ivan V.  
Bajic, Yufeng Shan*
116. Discrete Optimization of Ray Potentials for Semantic 3D  
Reconstruction, *Nikolay Savinov, Lubor Ladický, Christian  
Häne, Marc Pollefeys*
117. SOLD: Sub-Optimal Low-rank Decomposition for  
Efficient Video Segmentation, *Chenglong Li, Liang Lin,  
Wangmeng Zuo, Shuicheng Yan, Jin Tang*
118. On the Appearance of Translucent Edges, *Ioannis  
Gkioulekas, Bruce Walter, Edward H. Adelson, Kavita Bala,  
Todd Zickler*
119. On Pairwise Costs for Network Flow Multi-Object  
Tracking, *Vishesh Chari, Simon Lacoste-Julien, Ivan Laptev,  
Josef Sivic*
120. Fine-Grained Recognition Without Part Annotations,  
*Jonathan Krause, Hailin Jin, Jianchao Yang, Li Fei-Fei*
121. Robust Reconstruction of Indoor Scenes, *Sungjoon Choi,  
Qian-Yi Zhou, Vladlen Koltun*



# Thursday, June 11

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

**0730–0830 Breakfast** (Exhibit Hall B)

**1200–1400 Lunch** (Exhibit Hall B)

## Deep Vision: Deep Learning in Computer Vision

**Organizers:** Jose M. Alvarez  
Yann LeCun  
Fatih Porikli  
Yi Li

**Location:** Room 102

**Schedule:** Full Day

0845 Opening Remarks

0900 **Invited Talk:** Deep Learning: Progress in Theory and Attention Mechanisms, *Yoshua Bengio (Univ. of Montreal)*

**1000 Morning Break**

1030 **Invited Talk:** DeepID-Net: Deformable Deep Convolutional Neural Networks for Object Detection, *Xiaogang Wang (Chinese Univ. of Hong Kong)*

1115 **Invited Talk:** Biologically-Based Error Driven Learning in Thalamocortical Circuits, *Randall O'Reilly (Univ. of Colorado, Boulder)*

**1200 Poster Spotlight**

**1230 Lunch**

1330 **Invited Talk:** Learning and Understanding Visual Representations, *Andrea Vedaldi (Univ. of Oxford)*

1415 **Invited Talk:** TBA, *Rahul Sukthankar (Google)*

**1500 Poster Session**

- Self-Tuned Deep Super Resolution, *Zhangyang Wang, Yingzhen Yang, Zhaowen Wang, Shiyu Chang, Wei Han, Jianchao Yang, Thomas Huang*
- Channel-Max, Channel-Drop and Stochastic Max-Pooling, *Yuchi Huang, Xiuyu Sun, Ming Lu, Ming Xu*
- Convolutional Recurrent Neural Networks: Learning Spatial Dependencies for Image Representation, *Zhen Zuo, Bing Shuai, Gang Wang, Xiao Liu, Xingxing Wang, Bing Wang, Yushi Chen*
- Deep Learning of Binary Hash Codes for Fast Image Retrieval, *Kevin Lin, Huei-Fang Yang, Jen-Hao Hsiao, Chu-Song Chen*
- From Generic to Specific Deep Representations for Visual Recognition, *Hossein Azizpour, Ali Sharif Razavian, Josephine Sullivan, Atsuto Maki, Stefan Carlsson*
- Subset Feature Learning for Fine-Grained Category Classification, *ZongYuan Ge, Christopher McCool, Conrad Sanderson, Peter Corke*
- Exploiting Local Features From Deep Networks for Image Retrieval, *Joe Yue-Hei Ng, Fan Yang, Larry S. Davis*
- Object Level Deep Feature Pooling for Compact Image Representation, *Konda Reddy Mopuri, R. Venkatesh Babu*
- Multi-Scale Pyramid Pooling for Deep Convolutional Representation, *Donggeun Yoo, Sunggyun Park, Joon-Young Lee, In So Kweon*
- Color Constancy Using CNNs, *Simone Bianco, Claudio Cusano, Raimondo Schettini*
- Learning to Count With Deep Object Features, *Santi Seguí, Oriol Pujol, Jordi Vitrià*

**1530 Afternoon Break**

**1600 Best Paper Award**

1605 **Invited Talk:** Deep Neural Networks and GPUs, *Julie Bernauer (NVIDIA)*

1650 **Invited Talk:** Deep Semantic Learning: Teach Machines to Understand Text, Image, and Knowledge Graph, *Xiaodong He (Microsoft Research)*



## Semantics for Visual Reconstruction, Localization & Mapping

**Organizers:** Ian Reid  
Silvio Savarese  
Stephen Gould

**Location:** Room 103

**Schedule:** Full Day

- 0900 Welcome and Introduction
- 0915 **Invited Talk:** Semantically-Referenced Navigation,  
*Martial Hebert (CMU)*
- 0950 **Invited Talk:** Exploiting the Web for Reconstruction,  
Recognition and Self-Localization, *Raquel Urtasun*  
(Univ. of Toronto)
- 1025 Semantically Aware Bag-of-Words for Localization,  
*Arsalan Mousavian, Jana Kosecka*
- 1040 Morning Break**
- 1100 SynthCam: Semantic Understanding With Synthetic  
Indoor Scenes, *Ankur Handa, Viorica Patraucean, Vijay*  
*Badrinarayanan, Simon Stent, Roberto Cipolla*
- 1115 A Whole-Room 3D Context Model for Panoramic Scene  
Understanding, *Yinda Zhang, Shuran Song, Ping Tan,*  
*Jianxiang Xiao*
- 1130 **Invited Talk:** Challenges for Life-Long Visual Mapping  
and Navigation, *John Leonard (MIT)*
- 1205 **Invited Talk:** RoboBrain: Large-Scale Knowledge  
Engine for Robots, *Ashutosh Saxena (Cornell Univ.)*
- 1240 Lunch**
- 1340 **Invited Talk:** What Do We Hope to Achieve With Visual  
SLAM? *Richard Newcombe (Univ. of Washington)*
- 1405 Semantic Alignment of LiDAR Data at City Scale, *Fisher*  
*Yu, Jianxiang Xiao, Thomas Funkhouser*
- 1420 What is needed for Multi-view Semantic  
Segmentation? *Hayko Riemenschneider, et al.*
- 1450 Afternoon Break**
- 1510 Semantic Segmentation of Large-scale Urban 3D data  
with Low Annotation Cost, *Jinglu Wang, Showei Li,*  
*Jongbo Liu, Honghui Zhang, Tian Fang, Siyu Zhu, Punze*  
*Zhang, Shengnan Cai, Long Quan*

1525 **Invited Talk:** Semantic 3D Reconstruction, *Marc*  
*Pollefeys (ETH Zurich)*

1600 **Panel Discussion**

1630 Closing Remarks

## Group & Crowd Behavior Analysis & Understanding

**Organizers:** Vittorio Murino  
Marco Cristani  
Silvio Savarese  
Shishir Shah

**Location:** Room 107

**Schedule:** Full Day

- 0845 **Welcome**
- 0900 **Invited Talk:** TBA, *Shaogang Gong (Queen Mary Univ.)*
- 1000 Morning Break**
- 1030 Walking and Talking: A Bilinear Approach to Multi-Label  
Action Recognition, *Sameh Khamis, Larry S. Davis*
- 1055 Discovering Human Interactions in Videos With Limited  
Data Labeling, *Mehran Khodabandeh, Arash Vahdat,*  
*Guang-Tong Zhou, Hossein Hajimirsadeghi, Mehrsan*  
*Javan Roshtkhari, Greg Mori, Stephen Se*
- 1120 MuseumVisitors: A Dataset for Pedestrian and Group  
Detection, Gaze Estimation and Behavior  
Understanding, *Federico Bartoli, Giuseppe Lisanti,*  
*Lorenzo Seidenari, Svebor Karaman, Alberto Del Bimbo*
- 1145 Subject Centric Group Feature for Person Re-  
Identification, *Li Wei, Shishir K. Shah*
- 1210 The GRODE Metrics: Exploring the Performance of  
Group Detection Approaches, *Francesco Setti, Marco*  
*Cristani*
- 1235 Lunch**
- 1400 **Invited Talk:** TBA, *Chiara Bassetto (Inst. of Cognitive*  
*Science and Technology)*
- 1500 Afternoon Break**
- 1530 Learning to Identify Leaders in Crowd, *Francesco*  
*Solera, Simone Calderara, Rita Cucchiara*



- 1555 A Comparison of Crowd Commotion Measures From Generative Models, *Sadegh Mohammadi, Hamed Kiani, Alessandro Perina, Vittorio Murino*
- 1620 Invited Talk: Structured Models for Looking at People: Towards Sub-Category and Interaction Recognition, *Greg Mori (Simon Fraser Univ.)*
- 1720 Real-Time Anomaly Detection and Localization in Crowded Scenes, *Mohammad Sabokrou, Mahmood Fathy, Mojtaba Hoseini, Reinhard Klette*
- 1745 Dominant Flow Extraction and Analysis in Traffic Surveillance Videos, *Srinivas S S Kruthiventj, R. Venkatesh Babu*

## Language & Vision

**Organizers:** Andrei Barbu

Georgios Evangelopoulos  
Daniel Harari  
Krystian Mikolajczyk  
Siddharth Narayanaswamy  
Caiming Xiong  
Yibiao Zhao

**Location:** Room 201

**Schedule:** Full Day

- 0850 Opening Remarks
- 0900 **Invited Talk:** *Mirella Lapata (Univ. of Edinburgh)*
- 0930 **Invited Talk:** *Linda Smith (Indiana Univ.)*
- 1000 Morning Break**
- 1015 **Invited Talk:** *Kristen Grauman (Univ. of Texas at Austin)*
- 1045 **Invited Talk:** *Jason J. Corso (Univ. of Michigan)*
- 1115 **Invited Talk:** *Stefanie Tellex (Brown Univ.)*
- 1145 Image Annotation Challenge Summary
- 1215 Lunch**
- 1300 **Poster Session**
- 1400 **Invited Talk:** *Jeffrey Mark Siskind (Purdue Univ.)*
- 1430 **Invited Talk:** *Joyce Chai (Michigan State Univ.)*
- 1500 **Invited Talk:** *Song-Chun Zhu (UCLA)*

## 1530 Afternoon Break

- 1545 **Invited Talk:** *Tony Cohn (Univ. of Leeds)*
- 1615 **Invited Talk:** *Fei-Fei Li (Stanford Univ.)*
- 1645 **Invited Talk:** *Tomaso Poggio (MIT)*
- 1715 Discussion & Wrap-Up

## THUMOS Challenge

**Organizers:** Ivan Laptev  
Mubarak Shah  
Rahul Sukthankar  
Alexander Gorban  
Haroon Idrees  
Yu-Gang Jiang  
Amir R Zamir

**Location:** Room 206

**Schedule:** Full Day

- 1030 Opening Remarks, *Mubarak Shah (Univ. of Central Florida)*

### S1: Action Classification (1045-1215)

- 1045 Overview and Results of Classification Challenge, *Haroon Idrees (Univ. of Central Florida)*
- 1115 Challenge Winner Presentation
- 1135 Challenge Top Performer
- 1155 Challenge Top Performer

### **1215 Lunch**

### S2: Action Detection (1430-1530)

- 1430 Overview and Results of Detection Challenge, *Amir R. Zamir (Stanford Univ.)*
- 1450 Challenge Winner Presentation
- 1510 Challenge Top Performer

## 1530 Afternoon Break

- 1600 Summary of the 2015 Challenge and Discussion, *Alexander Gorban (Google Research), Haroon Idrees (UCF), Amir R. Zamir (Stanford Univ.), Ivan Laptev (INRIA), Mubarak Shah (UCF)*



## OpenCV Vision Challenge

**Organizers:** Gary Bradski  
Vadim Pisarevsky  
Vincent Rabaud  
Grace Vesom

**Location:** Room 101

**Schedule:** Half Day — Morning

This is a short workshop one hour before lunch to announce and describe winners of two separate contests:

- 1100 **People's Choice: Best Paper** — tally votes and present results for paper you'd most like to see in OpenCV.
- 1130 **State of the Art Vision Challenge** — present results of OpenCV vision challenge to meet or exceed the state of the art in various areas.

## Vision Meets Cognition: Functionality, Physics, Intentionality & Causality

**Organizers:** Yibiao Zhao  
Tao Gao  
Bo Zheng  
Peter Battaglia  
Lap-Fai Yu

**Location:** Room 104

**Schedule:** Half Day — Morning

- 0800 Welcome
- 0805 **Invited Talk:** *Antonis Argyros (Univ. of Crete)*
- 0835 **Invited Talk:** *Yiannis Aloimonos (Univ. of Maryland)*
- 0905 **Invited Talk:** *Brian Scholl (Yale Univ.)*
- 0935 **Invited Talk:** *Felix Warneken (Harvard Univ.)*
- 1005 **Morning Break & Poster Session**
- 1030 **Invited Talk:** *Josh Tenenbaum (MIT)*
- 1100 **Invited Talk:** *Song-Chun Zhu (UCLA)*
- 1130 **Invited Talk:** *Koray Kavukcuoglu (Google)*
- 1200 **Poster Session**

## Perception Beyond the Visible Spectrum

**Organizers:** Riad I Hammoud  
Guoliang Fan  
Erik Blasch

**Location:** Room 105

**Schedule:** Half Day — Morning

- 0800 Welcome
- 0805 **Spotlight Presentations** (3 min/presenter, 15 posters)
- 0850 **Keynote Talk:** 3D Assisted Common Operational Framework, *Guna Seetharaman (Air Force Research Laboratory)*

### S1: Evaluation and Database (0915-1000)

- 0915 Comparison of Infrared and Visible Imagery for Object Tracking: Toward Trackers With Superior IR Performance, *Erhan Gundogdu, Huseyin Ozkan, H. Seckin Demir, Hamza Ergezer, Erdem Akagündüz, S. Kubilay Pakin*
- 0930 VAIS: A Dataset for Recognizing Maritime Imagery in the Visible and Infrared Spectrums, *Mabel M. Zhang, Jean Choi, Kostas Daniilidis, Michael T. Wolf, Christopher Kanan*
- 0945 A Comparison of Stereo and Multiview 3-D Reconstruction Using Cross-Sensor Satellite Imagery, *Ozge C. Ozcanli, Yi Dong, Joseph L. Mundy, Helen Webb, Riad Hammoud, Victor Tom*

### **1000 Morning Break** (and Poster Setup)

### S2: Poster Session (1015-1100)

- Online Multimodal Video Registration Based on Shape Matching, *Pierre-Luc St-Charles, Guillaume-Alexandre Bilodeau, Robert Bergevin*
- Automated Feature Weighting and Random Pixel Sampling in *k*-Means Clustering for Terahertz Image Segmentation, *Mohamed Walid Ayeche, Djemel Ziou*
- A Model-Based Approach to Finding Tracks in SAR CCD Images, *Tu-Thach Quach, Rebecca Malinas, Mark W. Koch*



- Efficient Person Re-Identification by Hybrid Spatiogram and Covariance Descriptor, *Mingyong Zeng, Zemin Wu, Chang Tian, Lei Zhang, Lei Hu*
- Articulated Gaussian Kernel Correlation for Human Pose Estimation, *Meng Ding, Guoliang Fan*
- Automation of Dormant Pruning in Specialty Crop Production: An Adaptive Framework for Automatic Reconstruction and Modeling of Apple Trees, *Noha M. Elfiky, Shayan A. Akbar, Jianxin Sun, Johnny Park, Avinash Kak*
- A Cloud Infrastructure for Target Detection and Tracking Using Audio and Video Fusion , *Kui Liu, Bingwei Liu, Erik Blasch, Dan Shen, Zhonghai Wang, Haibin Ling, Genshe Chen*
- Make My Day - High-Fidelity Color Denoising With Near-Infrared, *Hiroto Honda, Radu Timofte, Luc Van Gool*
- USDOT Number Localization and Recognition From Vehicle Side-View NIR Images, *Orhan Bulan, Safwan Wshah, Ramesh Palghat, Vladimir Kozitsky, Aaron Burry*
- Robust Object Recognition in RGB-D Egocentric Videos Based on Sparse Affine Hull Kernel, *Shaohua Wan, J.K. Aggarwal*
- Improving Superpixel Boundaries Using Information Beyond the Visual Spectrum, *Keith Sullivan, Wallace Lawson, Donald Sofge*
- A Simply Integrated Dual-Sensor Based Non-Intrusive Iris Image Acquisition System, *Jang-Hee Yoo, Byung Jun Kang*
- Heterogeneous Structure Fusion for Target Recognition in Infrared Imagery, *Guangfeng Lin, Guoliang Fan, Liangjiang Yu, Xiaobing Kang, Erhu Zhang*
- Non-Rigid Articulated Point Set Registration With Local Structure Preservation, *Song Ge, Guoliang Fan*
- Sonar Automatic Target Recognition for Underwater UXO Remediation, *Jason C. Isaacs*

- 1140 Road Segmentation Using Multipass Single-Pol Synthetic Aperture Radar Imagery, *Mark W. Koch, Mary M. Moya, James G. Chow, Jeremy Goold, Rebecca Malinas*
- 1155 **Keynote Talk:** Towards a Unified Understanding of Image Quality: Quantifying Spatial, Temporal, and Spectral Information for Computer Vision, *John M. Irvine, (Charles Stark Draper Laboratory)*
- 1220 Concluding Remarks

1100 **Keynote Talk:** Analyzing the Flight Behavior of Bats in Thermal Infrared Video, *Margrit Betke (Boston Univ.)*

## **S3: Face Recognition & Object Segmentation (1125-1155)**

- 1125 NIR-VIS Heterogeneous Face Recognition via Cross-Spectral Joint Dictionary Learning and Reconstruction, *Felix Juefei-Xu, Dipan K. Pal, Marios Savvides*

## **Vision Industry & Entrepreneur Workshop**

**Organizers:** Sek Chai

Arnab Dhua  
Himanshu Arora  
Samson Timoner  
Kevin Ding  
Terry Boulton  
Boaz Super  
Ramesh Raskar

**Location:** Room 202

**Schedule:** Half Day — Morning

0815 **Poster and Demo Setup**

### **S1: Session 1 (0830–1015)**

- 0830 **Welcome:** *Sek Chai (SRI International)*
- 0840 **Introductions:** *Samson Timoner (Entrepreneur) & Kevin Ding (Entrepreneur)*
- 0850 **Venture Pitch Contest:** *Sek Chai (SRI International)*
- 0950 **Industry Session Spotlights:** *Arnab Dhua (A9.com)*

**1015 Morning Break**



## S2: Session 2 (1030–1145)

- 1030 **Venture Pitch Feedback:** *Samson Timoner*  
(*Entrepreneur*)
- 1030 **Industry Session — Demos, Posters, and Recruiting:**
- ClipMine: AI Assisted Collaborative Tagging of Videos for Enhanced Viewing, *Omar Javed, Ivana Savic, Saad Ali, Zia Syed*
  - Changing the Game: Novel Computer Vision Algorithms for Automated Sports Analytics, *Mehrsan Javan*
  - Brains4Drones: Embedded Vision Increases the Operational Safety of Unmanned Systems, *Goksel Dedeoglu, Susan Rossbach*
  - Conversant Decision Sciences, *Boaz Super*
  - Large Scale Text Detection in Video, *Amy Zhang, Carolina Galleguillos, Ana Murillo*
  - SRI International: World-Changing Solutions, *Sek Chai*
  - Collaborative Computer Vision R&D at Kitware, *Anthony Hoogs, Heather James*
  - Multi-GPU Training For Large Scale Visual Recognition, *Yi Li, Meng Wang, Wei Xia*
  - Visual Search Technologies at A9, *Himanshu Arora, Arnab Dhua*
  - LensBricks Smart Sensing Technology, *Rajeswari Kannan, Pranav Mishra*
  - Face Recognition Research at Systems & Technology Research, *Jeffrey Byrne*
  - Algolux - Picture Clarity through Computational Imaging, *Paul Green, Daniel Nahmias-Léonard*
  - Ingrain: Native Advertising in Online Video, *Aamer Zaheer, Murtaza Taj, Ali Rehan, Abdul Rehman, Faraz Hassan*
  - MIT Lincoln Laboratory Video Analytics, *Jason Thornton*
  - Computer Vision as a Service, *Ahmad Pahlavan Tafti, Zeyun Yu*

## S3: Session 3 (1145–1230)

- 1145 **Venture Pitch Contest Results:** *Samson Timoner*  
(*Entrepreneur*) & *Arnab Dhua* (A9.com)
- 1220 **Beyond VIEW 2015:** *Sek Chai* (SRI International)

## Medical Computer Vision

**Organizers:** Le Lu

Yefeng Zheng  
Bjoern Menze  
Georg Langs  
Leo Grady

**Location:** Room 203

**Schedule:** Half Day — Morning

The schedule of these invited talks is to be determined.

- **Invited Talk:** Surgical Gesture Segmentation and Recognition From Kinematic and Video Data, *René Vidal, (Johns Hopkins Univ.)*
- **Invited Talk:** Using Deep Learning to Detect Small Intestine Disorders, *Petia Radeva (Univ. de Barcelona)*
- **Invited Talk:** Visceral: Open Datasets and Challenges, *Georg Langs (Medical Univ. of Vienna)*
- **Invited Talk:** Towards Per-Voxel Large-Scale Radiology Image Parsing and Understanding, *Ronald M. Summers (National Institutes of Health)*
- **Invited Talk:** TBA, *Ramesh Raskar (MIT)*
- **Invited Talk:** Using Machine Learning Techniques to Reconstruct Complex Curvilinear Structures, *Pascal Fua (EPFL)*
- **Invited Talk:** Personalized Blood Flow Simulation: Changing the Treatment Paradigm for Cardiovascular Disease, *Leo Grady (Heartflow, Inc.)*
- **Invited Talk:** Deep Mining in Text/Image on a Hospital Scale PACS Database: Early Findings, *Hoo-chang Shin (National Institutes of Health)*
- **Invited Talk:** Big Data — Small Training Set: Addressing Medical Image Analysis Bottlenecks, *Tammy Riklin Raviv (MIT)*
- **Invited Talk:** Hierarchical Probabilistic Graphical Models for the Detection and Segmentation of Multiple Sclerosis Lesions in Multicentre Clinical Trial Datasets, *Tal Arbel (McGill Univ.)*
- **Invited Talk:** Large Scale Chest Radiograph Categorization With Deep CNN and No Medical Training, *Hayit Greenspan (Tel-Aviv Univ.)*
- **Invited Talk:** Computational Histopathology — Unlocking Tissue Content in Precision Medicine, *Christophe Chef'd'Hotel (Roche Tissue Diagnostics)*



## Visual Place Recognition in Changing Environments

**Organizers:** Niko Sünderhauf

Michael Milford

Peter Corke

Torsten Sattler

**Location:** Room 204

**Schedule:** Half Day — Morning

0830 Welcome & Introduction

0835 **Visual Place Recognition:** Tutorial & Overview

0900 **Place Recognition Challenge:** Presentation of Results

0915 **Poster Spotlights** (3-5 minutes each)

1000 **Morning Break & Poster Session**

1115 **Invited Talk:** TBD

1145 Open Discussion & Conclusions

## BioImage Computing

**Organizers:** Florian Jug

Dagmar Kainmüller

Carsten Rother

Pavel Tomancak

**Location:** Room 101

**Schedule:** Half Day — Afternoon

### Invited Speaker Session 1 (1330-1530)

1330 **Invited Talk:** We Want You! Challenging Computer Vision Problems in Modern Biology, *Pavel Tomancak (MPI-CBG Dresden)*

1400 **Invited Talk:** Joint Segmentation and Tracking, *Fred Hamprecht (HCI Heidelberg)*

1430 **Invited Talk:** TBA, *Charless Fowlkes (UC Irvine)*

1500 **Invited Talk:** TBA, *Sean Megason (Harvard Medical School)*

### Poster Session & Afternoon Break (1530-1600)

- **Full Paper:** From Photography to Microbiology: Eigen-biome Models for Skin Appearance, *Parneet Kaur, Kristin J. Dana, Gabriela Oana Cula*
- **Full Paper:** Fast Registration of Segmented Images by Normal Sampling, *Jan Kybic, Martin Dolejš, Jiří Borovec*
- **Full Paper:** Deep Neural Networks for Anatomical Brain Segmentation, *Alexander de Brébisson, Giovanni Montana*
- **Abstract:** Scale Invariant Particle Detection Using CNN of Log-Polar Images, *Kenshiro Nishida, Kazuhiro Hotta*
- **Abstract:** Efficient 3D Macromolecular Reconstruction With Electron Cryomicroscopy, *Marcus A. Brubaker, Ali Punjani, David J. Fleet*
- **Abstract:** Segmentation of Overlapping Cervical Cells With Nearest Nucleus Labeling and Graph Cuts, *Hansang Lee, Junmo Kim*

### Invited Speaker Session 2 (1600-1800)

- 1600 **Invited Talk:** Mapping Behavior to Neurons Using Computer Vision and Thermogenetics, *Kristin Branson (HHMI Janelia Farm Research Campus)*
- 1630 **Invited Talk:** Artifacts in Electron Microscopy - Road Blocks and Solutions Towards Complete Connectivity Reconstruction of the Drosophila Brain, *Stephan Saalfeld (HHMI Janelia Farm Research Campus)*
- 1700 **Invited Talk:** Candidate-Based Discrete Optimization for Biomedical Image Computing, *Jan Funke (ETHZ)*
- 1730 **Invited Talk:** Domain Adaptation for Microscopy Imaging, *Pascal Fua (EPFL)*
- 1800 **Awards & Closing**



## Biometrics

**Organizers:** Bir Bhanu

Ross Beveridge

Ajay Kumar

**Location:** Room 104

**Schedule:** Half Day — Afternoon

### S1: Face Recognition I (1330-1430)

**Chair:** Arun Ross (*Michigan State Univ.*)

1330 TAEF: A Cross-Distance/Environment Face Recognition Method, *Chun-Ting Huang, Zhengning Wang, C.-C. Jay Kuo*

1345 Perspective Distortion Modeling, Learning and Compensation, *Joachim Valente, Stefano Soatto*

1400 Locality-Constrained Discriminative Learning and Coding, *Shuyang Wang, Yun Fu*

1415 A Preliminary Investigation on the Sensitivity of COTS Face Recognition Systems to Forensic Analyst-Style Face Processing for Occlusions, *Felix Juefei-Xu, Dipan K. Pal, Karanhaar Singh, Marios Savvides*

### S2: Iris Recognition (1430-1515)

**Chair:** Bruce Draper (*Colorado State Univ.*)

1430 Exploratory Analysis of an Operational Iris Recognition Dataset From a CBSA Border-Crossing Application, *Estefan Ortiz, Kevin W. Bowyer*

1445 Evaluation of Combined Visible/NIR Camera for Iris Authentication on Smartphones, *Shejin Thavalengal, Petronel Bigioi, Peter Corcoran*

1500 A Preliminary Study on Identifying Sensors From Iris Images, *Nathan Kalka, Nick Bartlow, Bojan Cukic, Arun Ross*

### **1515 Afternoon Break**

### S3: Face Recognition II (1535-1630)

**Chair:** Kevin Bowyer (*Univ. of Notre Dame*)

1535 The Emperor's New Masks: On Demographic Differences and Disguises, *Katherine L. Gibson, Jonathan M. Smith*

1550 Latent Max-Margin Metric Learning for Comparing Video Face Tubes, *Gaurav Sharma, Patrick Pérez*

1605 Unsupervised Learning of Overcomplete Face Descriptors, *Juha Ylioinas, Juho Kannala, Abdenour Hadid, Matti Pietikäinen*

### S4: Biometric Security and Identification (1620-1730)

**Chair:** Walter Scheirer (*Harvard Univ.*)

1620 Person Identification From Action Styles, *Igor Kviatkovsky, Ilan Shimshoni, Ehud Rivlin*

1630 A Facial Features Detector Integrating Holistic Facial Information and Part-Based Model, *Eslam Mostafa, Asem A. Ali, Ahmed Shalaby, Aly Farag*

1640 Genetic Algorithm Attack on Minutiae-Based Fingerprint Authentication and Protected Template Fingerprint Systems, *Andras Rozsa, Albert E. Glock, Jr., Terrance E. Boulton*

1650 Electromyograph and Keystroke Dynamics for Spoof-Resistant Biometric Authentication, *Shreyas Venugopalan, Felix Juefei-Xu, Benjamin Cowley, Marios Savvides*

1700 A Multiple Server Scheme for Fingerprint Fuzzy Vaults, *Jesse Hartloff, Matthew Morse, Bingsheng Zhang, Thomas Eftland, Jennifer Cordaro, Jim Schuler, Sergey Tulyakov, Atri Rudra, Venu Govindaraju*

1710 Pore-Based Ridge Reconstruction for Fingerprint Recognition, *Mauricio Pamplona Segundo, Rubisley de Paula Lemes*

1720 **Awards, Valedictory, and Closing Remarks**



## Large Scale Visual Commerce

**Organizers:** Robinson Piramuthu  
Serge J Belongie  
Tsuhan Chen

**Location:** Room 105

**Schedule:** Half Day — Afternoon

1300 Welcome Message

1305 **Invited Talk:** TBA, *Kristen Grauman (Univ. of Texas, Austin)*

1325 **Invited Talk:** Reading Between the Lines, *Devi Parikh (Virginia Tech)*

TBD **Invited Talk:** Group Testing for Image Search With High-dimensional Descriptors, *Herve Jegou (INRIA)*

TBD **Invited Talk:** Selling Nostalgia, *Alyosha Efros (UC, Berkeley)*

TBD **Invited Talk:** A Picture is Worth a Thousand Ratings: Building Image-Based Models of Opinions and Purchases, *Julian McAuley (UC, San Diego)*

TBD **Invited Talk:** Fashion Analytics and Systems, *Shuicheng Yan (NUS)*

### 1515 Afternoon Break

TBD **Invited Talk:** TBA, *Ali Farhadi (Univ. of Washington)*

TBD **Invited Talk:** TBA, *David Forsyth (UIUC)*

TBD **Invited Talk:** TBA, *Anthony Hoogs (Kitware, Inc.)*

TBD **Invited Talk:** TBA, *Fei-Fei Li (Stanford University)*

1700 **Panel Discussions:** Moderator — Serge J. Belongie  
**Panelists:**

- Charles C. Fowlkes (*Univ. of California, Irvine*)
- Florent Perronnin (*Xerox Research Center Europe*)
- Fatih Porikli (*MERL*)
- Matthew Turk (*Univ. of California, Santa Barbara*)
- Luc Vincent (*Google*)
- Jay Yagnik (*Google*)

1800 **Social Event** (Finger food provided)

## Large-Scale Video Search & Mining

**Organizers:** Junsong Yuan  
Shih-Fu Chang  
John Smith

**Location:** Room 109

**Schedule:** Half Day — Afternoon

1330 **Invited Talk:** Emergent Representations in Deep Scene Networks, *Aude Oliva (MIT)*

1400 **Invited Talk:** Multimedia Event Detection, *Mubarak Shah (Univ. of Central Florida)*

1430 **Invited Talk:** Data-Driven Tag Refinement and Localization in Web Video, *Alberto del Bimbo (Univ. of Florence)*

1500 **Invited Talk:** Semantic Video Entity Linking Based on Visual Content and Metadata, *Jiebo Luo (Univ. of Rochester)*

### 1530 Afternoon Break

1600 **Invited Talk:** Hierarchical Context Modeling for Video Event Recognition, *Qiang Ji (Rensselaer Polytechnic Inst.)*

1630 **Invited Talk:** Improving Image Classification by Large-Scale Video Mining, *Rahul Sukthankar (Google)*

### 1700 Industry Panel Discussion:

Moderator — Shih-Fu Chang (*Columbia Univ.*)

#### Panelists:

- R. Manmatha (*Ag*)
- John R. Smith (*IBM*)
- Rahul Sukthankar (*Google*)
- Rong Yan (*Snapchat*)



## Performance Metrics for Correspondence Problems

**Organizers:** Daniel Kondermann

Michael Goesele

Katrin Honauer

Michael Waechter

Bernd Jaehne

**Location:** Room 202

**Schedule:** Half Day — Afternoon

1330 Welcome

1345 **Invited Talk:** *Daniel Scharstein (Middlebury College)*

1415 **Invited Talk:** *Robert Haralick (City Univ. of New York)*

1445 **Invited Talk:** To Be Decided

**1515 Afternoon Break & Poster Session**

1600 **Invited Talk:** To Be Decided

1630 **Invited Talk:** *Andreas Geiger, Jonas Wulff, Michael Black (Max Planck Inst. for Intelligent Systems)*

1700 **Invited Talk:** *Björn Fröhlich, Uwe Franke (Daimler AG)*

1730 **Panel Discussion** (Speakers & Audience)

1830 Closing Remarks

## Women in Computer Vision

**Organizers:** Judy Hoffman

Adriana Kovashka

Olga Russakovsky

Brigit Schroeder

Ning Zhang

**Location:** Room 203

**Schedule:** Half Day — Afternoon

1330 Introduction

1340 **Invited Talk:** *TBA, Fei-Fei Li (Stanford Univ.)*

1410 Oral Presentation #1

1420 Oral Presentation #2

1430 **Invited Talk:** *TBA, Mei Han (Google)*

**1450 Afternoon Break & Poster Session**

1620 Oral Presentation #3

1630 Oral Presentation #4

1640 **Invited Talk:** *TBA, Devi Parikh (Virginia Tech)*

1700 Closing Remarks

## The Future of Datasets in Vision

**Organizers:** Jason Corso

Julia Hockenmaier

Kate Saenko

**Location:** Room 204

**Schedule:** Half Day — Afternoon

**1200 Working Lunch** (*for all speakers, presentation of NSF proposal*)

1300 **Invited Talk:** *TBA, Fei-Fei Li (Stanford Univ.)*

1330 **Invited Talk:** *OntoNotes: Towards Integrated Linguistic Analysis, Sameer Pradhan (Harvard Univ.)*

1400 **Spotlights**

**1430 Afternoon Break & Poster Session**

1530 **Invited Talk:** *UCF Action Data Sets, Mubarak Shah (Univ. of Central Florida)*

1600 **Invited Talk:** *The Dataset Virus: The Infection of Ideas Through Data, Larry Zitnick (Microsoft Research)*

1630 **Discussion:** *Input from community on whether/how to federate datasets*

1730 **Working Dinner** (*by invitation; for speakers/presenters and NSF proposal PIs*)



# Friday, June 12

---

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

**0730–0830 Breakfast** (Exhibit Hall B)

**1200–1400 Lunch** (Exhibit Hall B)

## ChaLearn Looking at People Workshop & Competitions

**Organizers:** Sergio Escalera  
Jordi González  
Xavier Baró  
Isabelle Guyon

**Location:** Room 101

**Schedule:** Full Day

0900 Presentation of the Workshop

0915 **Invited Speaker:** TBA, *Tinne Tuytelaars (KU Leuven)*

1000 Morning Break

### S1: Challenge Results & Award Ceremony (1030-1050)

1030 ChaLearn Looking at People 2015 Challenges: Action Spotting and Cultural Event Recognition, *Xavier Baró, Jordi González, Junior Fabian, Miguel A. Bautista, Marc Oliu, Hugo Jair Escalante, Isabelle Guyon, Sergio Escalera*

### S2: Action & Gesture Recognition (1050-1145)

1050 Exploring Fisher Vector and Deep Networks for Action Spotting, *Zhe Wang, Limin Wang, Wenbin Du, Yu Qiao*

1105 Applying Action Attribute Class Validation to Improve Human Activity Recognition, *David Tahmouh*

1120 Keep it Accurate and Diverse: Enhancing Action Recognition Performance by Ensemble Learning, *Mohammad Bagheri, Qigang Gao, Sergio Escalera, Albert Clapes, Kamal Nasrollahi, Michael B. Holte, Thomas B. Moeslund*

1145 **Invited Speaker:** TBA, *Alexandre Alahi (Stanford)*

**1230 Lunch**

1330 **Invited Speaker:** Best of Both Worlds: Human-Machine Collaboration for Object Annotation, *Olga Russakovsky (Stanford)*

### S3: Cultural Event Recognition (1400-1500)

1400 Object-Scene Convolutional Neural Networks for Event Recognition in Images, *Limin Wang, Zhe Wang, Wenbin Du, Yu Qiao*

1415 Cultural Event Recognition With Visual ConvNets and Temporal Models, *Amaia Salvador, Matthias Zeppelzauer, Daniel Manchón-Vizuet, Andrea Calafell, Xavier Giró-i-Nieto*

1430 Cultural Event Recognition by Subregion Classification With Convolutional Neural Network, *Sungheon Park, Nojun Kwak*

1445 Recognizing Cultural Events in Images: A Study of Image Categorization Models, *Heeyoung Kwon, Kiwon Yun, Minh Hoai, Dimitris Samaras*

1500 **Invited Speaker:** People in Action and Interaction, *Stan Sclaroff (Boston Univ.)*

**1530 Afternoon Break**

### S4: Human Pose & Looking at People (1600-1645)

1600 Articulated Pose Estimation With Tiny Synthetic Videos, *Dennis Park, Deva Ramanan*

1615 A Semantic Occlusion Model for Human Pose Estimation From a Single Depth Image, *Umer Rafi, Juergen Gall, Bastian Leibe*

1630 A New Retexturing Method for Virtual Fitting Room Using Kinect 2 Camera, *Andres Traumann, Gholamreza Anbarjafari, Sergio Escalera*

### S5: Face Analysis (1645-1715)

1645 Pain Recognition Using Spatiotemporal Oriented Energy of Facial Muscles, *Ramin Irani, Kamal Nasrollahi, Thomas B. Moeslund*

1700 Spatiotemporal Analysis of RGB-D-T Facial Images for Multimodal Pain Level Recognition, *Ramin Irani, Kamal Nasrollahi, Marc O. Simon, Ciprian A. Corneanu, Sergio*



*Escalera, Chris Bahnsen, Dennis H. Lundtoft, Thomas B. Moeslund, Tanja L. Pedersen, Maria-Louise Klitgaard, Laura Petrini*

- 1715 **Invited Speaker:** Actions Recognition From Videos: Some Recent Results, *Cordelia Schmid (INRIA Grenoble)*
- 1800 **The COST Action iV&L Net IC1307:** New Challenges and Opportunities for Computer Vision, *Jordi González (UAB-CVC)*
- 1830 **Conclusions and Upcoming Events**

## BigVision: Large Scale Visual Recognition & Retrieval

**Organizers:** Olga Russakovsky  
Jason Corso  
Jia Deng  
Yuanqing Lin  
Samy Bengio  
Fei-Fei Li

**Location:** Room 102

**Schedule:** Full Day

- 0900 **Introduction**
- 0915 **Invited Talk:** TBA, *Antonio Torralba (MIT)*
- 1000 **Morning Break**
- 1030 **Invited Talk:** Learning Deep Structured Models, *Raquel Urtasun (Univ. of Toronto)*
- 1115 **Invited Talk:** TBA, *Christopher Ré (Stanford Univ.)*
- 1200 **Poster Spotlights**
- 1230 **Lunch & Posters**
- 1400 **Invited Talk:** TBA, *Abhinav Gupta (Carnegie Mellon Univ.)*
- 1445 **Invited Talk:** TBA, *Jianxiong Xiao (Princeton Univ.)*
- 1530 **Awards Presentation**
- 1535 **Afternoon Break**
- 1600 **Invited Talk:** Scalable Approaches for Large Scale Vision, *Christian Szegedy (Google)*

## Observing & Understanding Hands in Action

**Organizers:** Tae-Kyun Kim  
Gregory Rogez

**Location:** Room 104

**Schedule:** Full Day

- 0840 **Opening Remarks**
- 0850 **Invited Talk:** Capturing Interacting Hands and Objects, *Jürgen Gall (Univ. of Bonn)*
- 0920 **Invited Talk:** TBD

1000 **Morning Break**

### S1: Oral Session 1 (1030-1130)

- 1030 **Hand Gesture Recognition With 3D Convolutional Neural Networks**, *Pavlo Molchanov, Shalini Gupta, Kihwan Kim, Jan Kautz*
- 1050 **Hierarchical Particle Filtering for 3D Hand Tracking**, *Alexandros Makris, Nikolaos Kyriazis, Antonis A. Argyros*
- 1110 **On-the-Fly Hand Detection Training With Application in Egocentric Action Recognition**, *Jayant Kumar, Qun Li, Survi Kyal, Edgar A. Bernal, Raja Bala*
- 1130 **Invited Talk:** High Performance Hand Pose Estimation, *Yichen Wei (MSR Asia)*
- 1200 **Spotlights** (Extended Abstracts)
- 1230 **Lunch**
- 1400 **Invited Talk:** Rigid-Body Dynamics for Articulated Mesh Tracking, *Leonid Keselman (Intel)*
- 1430 **Invited Talk:** Real-Time Capture of Hands in Motion, *Christian Theobalt (MPII & Saarland Univ.)*
- 1500 **Poster Session** (Extended Abstracts):
- Classifying Object and Task Properties in Unstructured Environments, *Thomas Feix, Ian M. Bullock, Aaron M. Dollar*
  - Tracking Hands of Interacting People in Egocentric Video, *Sven Bambach, Stefan Lee, David J. Crandall, John M. Franchak, Chen Yu*



- Latent Regression Forest: Structured Estimation of 3D Articulated Hand Posture, *Danhang Tang, Hyung Jin Chang, Alykhan Tejani, Tae-Kyun Kim*
- Estimation of 3D Hand Position and Gestures on Unmodified Wearable Devices, *Jie Song, Gábor Sörös, Fabrizio Pece, Otmar Hilliges*
- Grasp Type Recognition for Human Action Intention and Manipulation, *Yezhou Yang, Cornelia Fermüller, Yi Li, Yiannis Aloimonos*
- Hand Grasp Recognition From Egocentric Videos, *Minjie Cai, Kris Kitani, Yoichi Sato*
- Depth-Based Hand Pose Estimation: Data, Method and Challenges, *James S. Supančič III, Grégory Rogez, Yi Yang, Jamie Shotton, Deva Ramanan*
- Robust Articulated-ICP for Real-Time Hand Tracking, *Andrea Tagliasacchi, Matthias Schröder, Anastasia Tkach, Sofien Bouaziz, Mario Botsch, Mark Pauly*
- The Yale Human Grasping Data Set, *Ian M. Bullock, Thomas Feix, Aaron M. Dollar*
- Classifying Dexterous Manipulation in Human and Robotic Systems, *Ian M. Bullock, Raymond R. Ma, Aaron M. Dollar*
- How Do We Use Our Hands? Discovering a Diverse Set of Common Grasps, *De-An Huang, Minghuang Ma, Wei-Chiu Ma, Kris M. Kitani*
- First-Person Pose Recognition Using Egocentric Workspaces, *Grégory Rogez, James S. Supančič III, Deva Ramanan*
- Hand Parsing for Fine-Grained Recognition of Human Grasps in Monocular Images, *Akanksha Saran, Damien Teney, Kris M. Kitani*
- Hand Gesture Recognition for Driver Vehicle Interaction, *Yannick Jacob, Sotiris Manitsaris, Fabien Moutarde, Gautam Lele, Laetitia Pradere*
- Spatio-Temporal Hough Forest for Efficient Detection-Localisation-Recognition of Fingerwriting in Egocentric Camera, *Hyung Jin Chang, Guillermo Garcia-Hernando, Danhang Tang, Tae-Kyun Kim*
- Driver Hands in Action: Introducing Vision for Intelligent Vehicles and Applications (VIVA) Challenge, *Eshed Ohn-Bar, Mohan Trivedi*

## **S2: Oral Session 2 (1600-1700)**

- 1600 ICPiK: Inverse Kinematics Based Articulated-ICP, *Shachar Fleishman, Mark Kliger, Alon Lerner, Gershon Kutliroff*
- 1620 Mining Discriminative States of Hands and Objects to Recognize Egocentric Actions With a Wearable RGBD Camera, *Shaohua Wan, J.K. Aggarwal*
- 1640 American Sign Language Alphabet Recognition Using Microsoft Kinect, *Cao Dong, Ming C. Lev, Zhaozheng Yin*
- 1700 **Invited Talk:** Learning an Efficient Model of Hand Shape Variation From Depth Images, *Jonathan Taylor (Microsoft Research Cambridge)*
- 1730 **Awards & Closing Remarks**

## **Analysis and Modeling of Faces & Gestures**

**Organizers:** Thomas S. Huang  
Dimitris N. Metaxas  
Yun Raymond Fu  
Mohammad Soleymani  
Walter J. Scheirer

**Location:** Room 105

**Schedule:** Full Day

0845 Opening remarks

### **S1: Morning Session (0850-1150)**

**Chair:** Walter J. Scheirer (Harvard Univ.)

- 0850 Towards Robust Cascaded Regression for Face Alignment in the Wild, *Chengchao Qu, Hua Gao, Eduardo Monari, Jürgen Beyerer, Jean-Philippe Thiran*
- 0910 **Invited Talk:** TBA, *Kevin W. Bowyer (Univ. of Notre Dame)*

**1000 Morning Break**

**1530 Afternoon Break**



- 1030 Multi-Observation Face Recognition in Videos Based on Label Propagation, *Bogdan Raducanu, Alireza Bosaghzadeh, Fadi Dornaika*
- 1050 Exemplar Hidden Markov Models for Classification of Facial Expressions in Videos, *Karan Sikka, Abhinav Dhall, Marian Bartlett*
- 1110 Using Hankel Matrices for Dynamics-Based Facial Emotion Recognition and Pain Detection, *Liliana Lo Presti, Marco La Cascia*
- 1130 Age and Gender Classification Using Convolutional Neural Networks, *Gil Levi, Tal Hassner*
- 1150 **Invited Talk:** Re-Rendering Tiny Movements in People, *Bill Freeman (MIT)*

## 1240 Lunch

### 52: Afternoon Session (1400-1700)

**Chair:** *Yun Raymond Fu (Northeastern Univ.)*

- 1400 FAME: Face Association Through Model Evolution, *Eren Golge, Pinar Duygulu-Sahin*
- 1420 Head Pose Estimation in the Wild Using Approximate View Manifolds, *Kalaivani Sundararajan, Damon L. Woodard*
- 1440 **Invited Talk:** Capturing Facial Muscular Interactions for Robust Facial Action Recognition, *Qiang Ji (Rensselaer Polytechnic Inst.)*
- 1530 **Afternoon Break**
- 1600 Mixture of Parts Revisited: Expressive Part Interactions for Pose Estimation, *Anoop R. Katti, Anurag Mittal*
- 1620 Towards Privacy-Preserving Activity Recognition Using Extremely Low Temporal and Spatial Resolution Cameras, *Ji Dai, Jonathan Wu, Behrouz Saghaei, Janusz Konrad, Prakash Ishwar*
- 1640 Sparse Coding Trees With Application to Emotion Classification, *Kevin Chen, Marcus Z. Comiter, H. T. Kung, Brad McDanel*
- 1700 **Best Paper Announcement & Conclusion**

## 3D From a Single Image

**Organizers:** Joao Carreira  
Niloy Mitra  
Hao Su  
Sara Vicente

**Location:** Room 107

**Schedule:** Full Day

- 0855 Welcome and Introduction
- 0900 3D Object Class Detection in the Wild, *Bojan Pepik, Michael Stark, Peter Gehler, Tobias Ritschel, Bernt Schiele*
- 0915 **Invited Talk:** TBA, *Leonidas Guibas (Stanford Univ.)*
- 0955 **Invited Talk:** TBA, *Alan Yuille (UCLA)*
- 1035 **Morning Break**
- 1105 **Invited Talk:** TBA, *Vladlen Koltun (Intel)*
- 1145 **Invited Talk:** TBA, *Silvio Savarese (Stanford Univ.)*
- 1225 **Lunch**
- 1400 **Invited Talk:** TBA, *Lourdes Agapito (Univ. College London)*
- 1440 **Invited Talk:** TBA, *Jitendra Malik (UC Berkeley)*
- 1520 **Invited Talk:** TBA, *Cristian Sminchisescu (Lund Univ.)*
- 1600 **Afternoon Break**
- 1630 **Invited Talk:** TBA, *Jianxiong Xiao (Princeton Univ.)*
- 1710 **Invited Talk:** TBA, *Qixing Huang (Toyota Technological Inst. at Chicago)*
- 1750 Closing Remarks



## Embedded Vision

**Organizers:** Goksel Dedeoglu  
Stefano Mattoccia

**Location:** Room 108

**Schedule:** Full Day

0900 Opening Remarks

0915 **Keynote Talk:** Embedded Vision for Autonomous Micro-Aerial Vehicles, *Marc Pollefeys (ETH Zurich)*

**1000 Morning Break**

1030 On-Board Real-Time Tracking of Pedestrians on a UAV, *Floris De Smedt, Dries Hulens, Toon Goedemé*

1045 Guidance: A Visual Sensing Platform For Robotic Applications, *Guyue Zhou, Lu Fang, Ketan Tang, Honghui Zhang, Kai Wang, Kang Yang*

1100 **Invited Talk:** Low-Power Embedded Vision using Approximate Computing, *Joseph Bates (Singular Computing)*

**1130 Poster Session:**

- Off-the-Shelf Sensor Integration for Mono-SLAM on Smart Devices, *Philipp Tiefenbacher, Timo Schulze, Gerhard Rigoll*
- Retrieving Gray-Level Information From a Binary Sensor and Its Application to Gesture Detection, *Orazio Gallo, Iuri Frosio, Leonardo Gasparini, Kari Pulli, Massimo Gottardi*
- Recursive Edge-Aware Filters for Stereo Matching, *Cevahir Çiğla*
- A Real-Time High Dynamic Range HD Video Camera, *Rajesh Narasimha, Umit Batur*
- FPGA Acceleration for Feature Based Processing Applications, *Gooitzen Van der Wal, David Zhang, Indu Kandaswamy, Jim Marakowitz, Kevin Kaighn, Joe Zhang, Sek Chai*

**1230 Lunch**

1400 **Invited Talk:** Perception for Autonomous Navigation in Dynamic Environments on Land and Sea, *Larry Matthies (JPL NASA)*

1445 Locally Non-Rigid Registration for Mobile HDR Photography, *Orazio Gallo, Alejandro Troccoli, Jun Hu, Kari Pulli, Jan Kautz*

1500 Real-Time Embedded Age and Gender Classification in Unconstrained Video, *Ramin Azarmehr, Robert Laganière, Won-Sook Lee, Christina Xu, Daniel Laroche*

1515 FPGA-Based Pedestrian Detection Under Strong Distortions, *Daniele Tasson, Alessio Montagnini, Roberto Marzotto, Michela Farenzena, Marco Cristani*

**1530 Afternoon Break**

1600 **Invited Talk:** Visual Single and General Object Tracking: Where Are We Today? *Roman Pfugfelder (Austrian Institute of Technology)*

1630 **Awards**

1640 Closing Remarks

## Large-Scale Scene Understanding Challenge

**Organizers:** Yinda Zhang  
Fisher Yu  
Shuran Song  
Pingmei Xu  
Jianxiong Xiao

**Location:** Room 103

**Schedule:** Half Day — Morning

0830 Introduction

0920 **Winner:** Scene Classification Challenge

**1000 Morning Break**

1030 **Winner:** Room Layout Estimation Challenge

1110 **Winner:** Saliency Prediction Challenge

1150 **Winner:** Image Captioning Challenge



## Computational Cameras & Displays

**Organizers:** Kathrin Berkner  
Oliver Cossairt  
Jingyi Yu

**Location:** Room 109

**Schedule:** Half Day — Morning

0810 Welcome and Introductory Remarks

### S1: Papers Session 1 (0815-0845)

- 0815 SparkleVision: Seeing the World Through Random Specular Microfacets, *Zhengdong Zhang, Phillip Isola, Edward H. Adelson*
- 0830 Efficient 3D Kernel Estimation for Non-Uniform Camera Shake Removal Using Perpendicular Camera System, *Tao Yue, Jinli Suo, Qionghai Dai*
- 0845 **Keynote Talk:** In Defense of Compressive Imaging, *Aswin Sankaranarayanan (Carnegie Mellon Univ.)*

### S2: Papers Session 2 (0930-1015)

- 0930 Reconstruction-Free Inference on Compressive Measurements, *Suhas Lohit, Kuldeep Kulkarni, Pavan Turaga, Jian Wang, Aswin C. Sankaranarayanan*
- 0945 Dense Sampling of 3D Color Transfer Functions Using HDR Photography, *Marcel Heinz, Guido Brunnett*
- 1000 Fresnel Lens Imaging With Post-Capture Image Processing, *Artem Nikonov, Roman Skidanov, Vladimir Fursov, Maksim Petrov, Sergey Bibikov, Yuri Yuzfovich*
- 1015 **Morning Break**
- 1015 **Keynote Talk:** Transport-Aware Cameras: The Future of Indoor and Outdoor Structured-Light Imaging? *Kyros Kutulakos (Univ. of Toronto)*

### S3: Papers Session 3 (1115-1215)

- 1115 Video Stitching With Spatial-Temporal Content-Preserving Warping, *Wei Jiang, Jinwei Gu*
- 1130 Video Compressive Sensing With On-Chip Programmable Subsampling, *Leonidas Spinoulas, Kuan He, Oliver Cossairt, Aggelos Katsaggelos*

- 1145 Fast Single-Frequency Time-of-Flight Range Imaging, *Ryan Crabb, Roberto Manduchi*
- 1200 High Speed Sequential Illumination With Electronic Rolling Shutter Cameras, *Matis Hudon, Paul Kerbiriou, Arno Schubert, Kadi Bouatouch*

1215 **Best Paper Award & Concluding Remarks**

## Fine-Grained Visual Categorization

**Organizers:** Ryan Farrell  
Subhransu Maji  
Anelia Angelova  
Grant Van Horn

**Location:** Room 110

**Schedule:** Half Day — Morning

- 0755 Welcome
- 0800 **Invited Talk:** Invariant Image Representations for Fine-Grained Categorization, *Florent Perronnin (Facebook AI Research)*
- 0830 **Invited Talk:** Visipedia Circa 2015: Vision Systems Composed of People and Machines, *Pietro Perona (California Institute of Technology)*
- 0900 **Invited Talk:** eButterfly: Bringing Analog and Digital Together for Science and Conservation, *Kent McFarland (Vermont Center for Ecotudies)*
- 0930 **Poster Session and Morning Break:**
- Weakly Supervised Fine-Grained Image Categorization, *Yu Zhang, Jianxin Wu, Jianfei Cai*
  - Fine-Grained Recognition for Biodiversity Analysis, *Erik Rodner, Marcel Simon, Gunnar Brehm, Stephanie Pietsch, J. Wolfgang Wägele, Joachim Denzler*
  - Fine-Grained Object Categorization via Localized Attributes, *Babak Saleh, Ahmed Elgammal*
  - Deep Filter Banks for Texture Recognition and Segmentation, *Mircea Cimpoi, Subhransu Maji, Andrea Vedaldi*



- Hierarchical Transfer of Semantic Attributes, *Ziad Al-Halah, Rainer Stiefel*
- Understanding Objects in Detail With Fine-Grained Attributes, *Andrea Vedaldi, Siddarth Mahendran, Stavros Tsogkas, Subhransu Maji, Ross Girshick, Juho Kannala, Esa Rathu, Iasonas Kokkinos, Matthew Blaschko, David Weiss, Ben Taskar, Karen Simonyan, Naomi Saphra, Sammy Mohamed*
- Learning Regions and Descriptors for Fine-Grained Recognition, *Dequan Wang*
- Fine-Grained Visual Comparisons With Local Learning, *Aron Yu, Kristen Grauman*
- Fast Bird Part Localization for Fine-Grained Categorization, *Yaser Souri, Shohreh Kasaei*
- WHOI-Plankton - A Large Scale Fine Grained Visual Recognition Benchmark Dataset for Plankton Classification, *Eric Orenstein, Oscar Beijbom, Emily Peacock, Heidi Sosik*
- Similarity Comparisons for Interactive Fine-Grained Categorization, *Catherine Wah, Grant Van Horn, Steve Branson, Subhransu Maji, Pietro Perona, Serge Belongie*
- A Dataset For Vehicle Make And Model Recognition, *Amine Ben Khalifa, Hichem Frigui*
- Fine-Grained Recognition of Physical Objects on Mobile Phones: From Categorization to Identification, *Maurits Diephuis, Slava Voloshynovskiy, Taras Holotyak*

1030 **Invited Talk:** TBA, *Alexei (Alyosha) Efros (UC Berkeley)*

1100 **Invited Talk:** TBA, *Ken Nakayama (Harvard Univ.)*

1130 **Panel Discussion**

1230 **Closing Remarks**

## Looking From Above: When Earth Observation Meets Vision

**Organizers:** Konrad Schindler

Josiane Zerubia

Devis Tuia

Gabriele Moser

Jan Dirk Wegner

**Location:** Room 111

**Schedule:** Half Day — Morning

0820 Welcome

0835 **Invited Talk:** Monitoring Earth Climate Variables with Statistical Inference, *Gustau Camps-Valls (Univ. de Valencia)*

0920 Semantic Segmentation of Urban Scenes by Learning Local Class Interactions, *Michele Volpi, Vittorio Ferrari*

0940 Active Learning Approach to Detecting Standing Dead Trees From ALS Point Clouds Combined With Aerial Infrared Imagery, *Przemyslaw Polewski, Wei Yao, Marco Heurich, Peter Krzystek, Uwe Stilla*

### 1000 Morning Break & Poster Session

- Universality of Wavelet-Based Non-Homogeneous Hidden Markov Chain Model Features for Hyperspectral Signatures, *Siwei Feng, Marco F. Duarte, Mario Parente*
- A Semi-Supervised Approach for Ice-Water Classification Using Dual-Polarization SAR Satellite Imagery, *Fan Li, David A. Clausi, Lei Wang, Linlin Xu*
- Effective Semantic Pixel Labelling With Convolutional Networks and Conditional Random Fields, *Sakrapee Paisitkriangkrai, Jamie Sherrah, Pranam Janney, Anton Van-Den Hengel*
- Do Deep Features Generalize From Everyday Objects to Remote Sensing and Aerial Scenes Domains?, *Otávio A. B. Penatti, Keiller Nogueira, Jefersson A. dos Santos*
- Matching Persistent Scatterers to Optical Oblique Images, *Lukas Schack, Uwe Soergel*
- Simultaneous Registration and Change Detection in Multitemporal, Very High Resolution Remote Sensing Data, *Maria Vakalopoulou, Konstantinos Karantzalos, Nikos Komodakis, Nikos Paragios*



- On the Location Dependence of Convolutional Neural Network Features, *Scott Workman, Nathan Jacobs*
- Oil Spill Candidate Detection From SAR Imagery Using a Thresholding-Guided Stochastic Fully-Connected Conditional Random Field Model, *Linlin Xu, M. Javad Shafiee, Alex Wong, Fan Li, Lei Wang, David Clausi*
- Detection of Incomplete Enclosures of Rectangular Shape in Remotely Sensed Images, *Igor Zingman, Dietmar Saupe, Karsten Lambers*

1130 **Invited Talk:** Large-Scale Aerial 3D Reconstruction for Google Earth and Google Maps, *Aleksey Golovinskiy (Google)*

1215 **Closing Remarks & Best Papers Awards**

## Scene Understanding

**Organizers:** James Hays  
Aditya Khosla  
Silvio Savarese  
Jianxiong Xiao

**Location:** Room 103

**Schedule:** Half Day — Afternoon

1330 **Invited Talk:** *Jitendra Malik (Univ. of California, Berkeley)*

1400 **Invited Talk:** *Josh Tenenbaum (MIT)*

1430 **Invited Talk:** *Ashutosh Saxena (Cornell Univ.)*

1500 **Invited Talk:** *Martial Hebert (Carnegie Mellon Univ.)*

1530 Poster Session & Coffee Break

1645 **Invited Talk:** *Rob Fergus (New York Univ.; Facebook AI Research)*

1715 **Invited Talk:** *Antonio Torralba (MIT)*

1745 **Invited Talk:** *Abhinav Gupta (Carnegie Mellon Univ.)*

## Multi-Sensor Fusion for Dynamic Scene Understanding

**Organizers:** Kim Boyer  
Christian Heipke  
Alper Yilmaz  
Clément Mallet  
Michael Ying Yang  
Yury Vizilter

**Location:** Room 110

**Schedule:** Half Day — Afternoon

1330 Welcome and Overview

1340 Geometric Inpainting of 3D Structures, *Pratyush Sahay, A. N. Rajagopalan*

1400 Real-Time Non-Rigid Multi-Frame Depth Video Super-Resolution, *Kassem Al Ismaeil, Djamila Aouada, Thomas Solignac, Bruno Mirbach, Björn Ottersten*

1420 **Keynote Talk:** TBA, *Jason J. Corso (Univ. of Michigan)*

**1515 Afternoon Break**

1550 Accurate Localization by Fusing Images and GPS Signals, *Kumar Vishal, C. V. Jawahar, Vishes Chari*

1610 Exploiting Global Priors for RGB-D Saliency Detection, *Jianqiang Ren, Xiaojin Gong, Lu Yu, Wenhui Zhou, Michael Ying Yang*

1630 Sparse Re-Id: Block Sparsity for Person Re-Identification, *Srikrishna Karanam, Yang Li, Richard J. Radke*



## Computer Vision in Vehicle Technology

**Organizers:** Jose M Alvarez  
David Vazquez  
Sebastian Ramos  
Tomas Pajdla  
Antonio M Lopez

**Location:** Room 111

**Schedule:** Half Day — Afternoon

1300 Opening Remarks

1305 **Invited Talk:** Realizing Self-Driving Cars, *Andreas Wendel (Google [X])*

1350 **Invited Talk:** 3D Reconstruction on Planets, *Tomas Pajdla (CTU Prague)*

1435 **Spotlights Talks**

**1447 Posters Session & Afternoon Break**

- Seamless Change Detection and Mosaicing for Aerial Imagery, *Nimisha T.M, A. N. Rajagopalan, Rangarajan Aravind*
- Absolute Geo-Localization Thanks to Hidden Markov Model and Exemplar-Based Metric Learning, *Cédric Le Barz, Nicolas Thome, Matthieu Cord, Stéphane Herbin, Martial Sanfourche*
- Sequence Searching With Deep-Learnt Depth for Condition- and Viewpoint-Invariant Route-Based Place Recognition, *Michael Milford, Chunhua Shen, Stephanie Lowry, Niko Suenderhauf, Sareh Shirazi, Guosheng Lin, Fayao Liu, Edward Pepperell, Cesar Lerma, Ben Upcroft, Ian Reid*
- Robust and Fast Detection of Moving Vehicles in Aerial Videos Using Sliding Windows, *Michael Teutsch, Wolfgang Krüger*
- Driver Cell Phone Usage Detection on Strategic Highway Research Program (SHRP2) Face View Videos, *Ke-shav Seshadri, Felix Juefei-Xu, Dipan K. Pal, Marios Savvides, Craig P. Thor*

1525 **Best Paper Award**

1530 **Invited Talk:** Progress Toward Autonomous Aerial Mobility Applications in Planetary Exploration, *Larry Matthies (JPL NASA)*

1615 **Invited Talk:** Vision-Based Robots to Explore the Ocean, *Rafael Garcia (Univ. of Girona)*

1700 **Invited Talk:** Vehicle Localization and Mapping with Monocular Wide-Field-of-View Cameras, *Marc Pollefeys (ETH Zurich)*

1745 Closing Remarks



## Silver Donors

---



**IBM Research**

## Bronze Donors

---





## Platinum Donors

---



## Gold Donors

---

