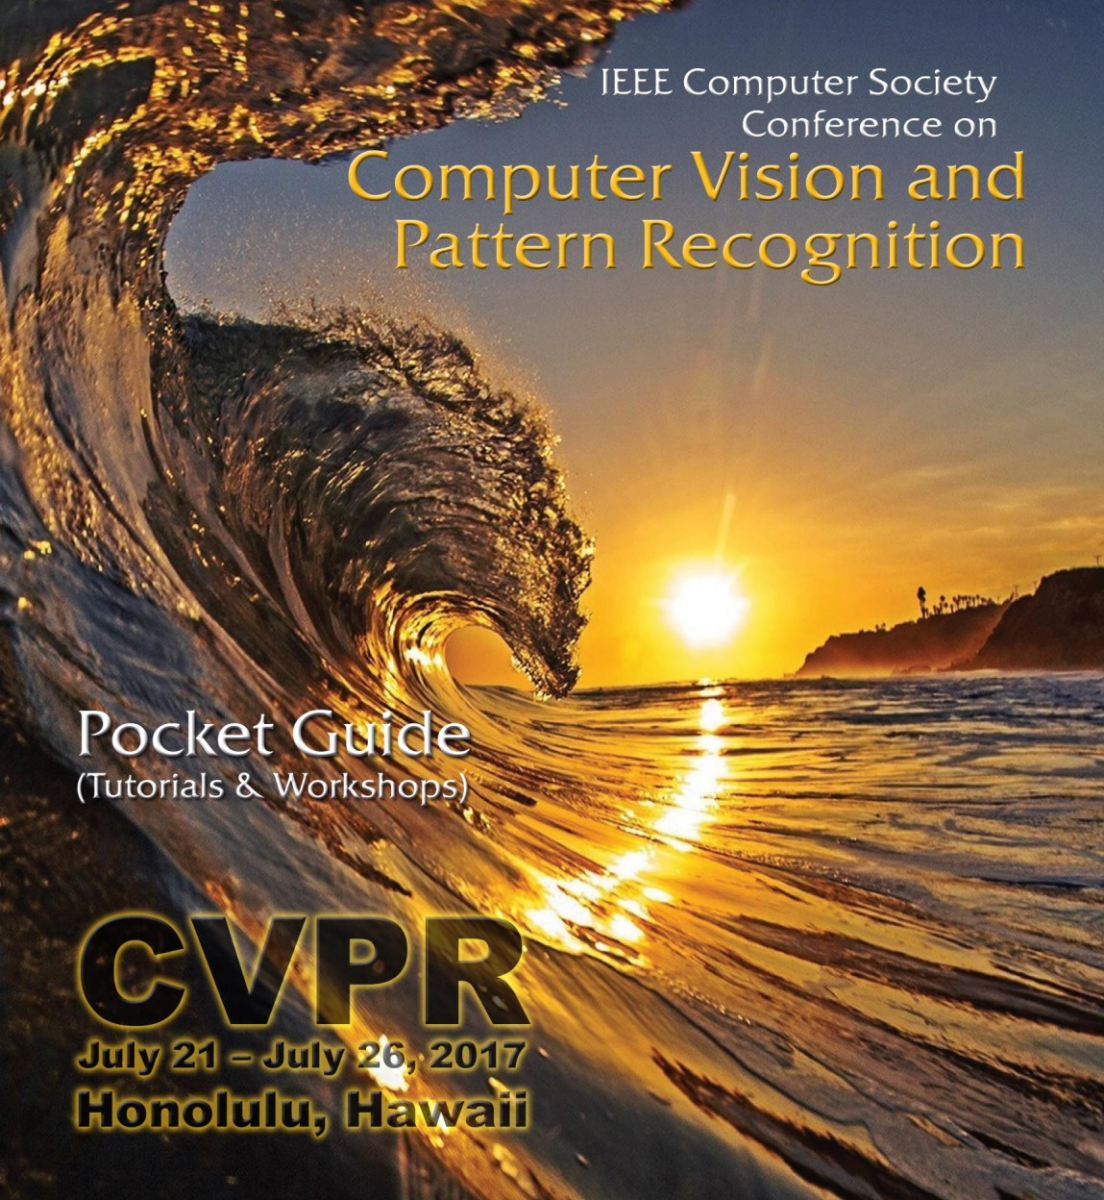
IEEE Computer Society
Conference on
Computer Vision and
Pattern Recognition
Pocket Guide
(Tutorials & Workshops)
CVPR
July 21 – July 26, 2017
Honolulu, Hawaii

HAWAI'I CONVENTION CENTER

Level 3*

CVPR 2017
July 21-26
HONOLULU

LEGEND

- Information desk
- Business center
- 808 Café
- First aid
- Defibrillator
- Escalator (2nd FL Parking)
- Escalator (3rd & 4th FL)
- Elevator
- Restroom (Men)
- Restroom (Women)
- Art Display
- ATM
- Vending area
- Water fountain
- Bottle filler
- Smoking area
- LCD board
- Parking
- Entrance
- Automatic entry door

CVPR 2017 Sponsors

IEEE computer society

The Computer Vision Foundation

Complimentary Wi-Fi Area
Connect to : Hawaii Free Wi-Fi

*The other levels are in the main conference pocket guide.

Message from the General and Program Chairs

Aloha!

Welcome to Honolulu, Hawaii and the 30th IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). In addition to the main four-day program of presentations, interactive sessions, plenary talks, demos, exhibitions, and social functions, CVPR 2017 has a number of co-located events, including 44 workshops and 20 tutorials. As the fields of computer vision, pattern recognition, machine learning and artificial intelligence continue to break new ground and scale new heights, so does our conference. This year CVPR 2017 received a record 2680 valid submissions to the main conference, of which 2620 were fully reviewed (the others were administratively rejected for technical or ethical reasons or were withdrawn before review).

The number of papers reviewed for CVPR 2017 was 40% larger than the number reviewed for last year's edition of the conference, a growth rate that posed significant organizational challenges.

To select papers for the program from these submissions, we invited 85 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. Additionally, we recruited a record number of 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 reviewing process accepted 783 papers (29% of valid submissions). 71 of these were accepted as oral presentations (2.65% of valid submissions) and 144 were accepted as spotlight oral presentations, for a total of 8% of valid submissions with live presentations. Continuing the successful innovation from CVPR 2016, the inclusion of spotlight oral presentations has allowed us to increase the number of works presented from the podium.

All accepted papers will appear in the interactive poster sessions where we hope that lively discussions will ensue. The total number of papers presented at CVPR 2017 is 22% larger than the number presented at last year's edition.

The review process was similar to previous years. Each paper was reviewed by at least three reviewers and considered by at least three ACs before a decision was made. Borderline papers and candidate orals and spotlights were discussed in groups of three non-conflicted ACs with common areas of expertise. Oral and spotlight recommendations were made by panels of ACs after extensive discussion.

The Program Chairs did not submit any papers to CVPR 2017, allowing them to avoid direct conflicts throughout the review process. This year, General Chairs who were allowed to submit papers, did not have any software access to the CMT system beyond that of an author. The double-blind nature of the CVPR review process was maintained throughout.

This year we have expanded the format of the technical program in two important ways. First, there are three parallel oral sessions for the first two days of the conference. While CVPR has grown immensely over the years, it has not increased the number of parallel oral tracks for more than two decades. Expanding from two to three tracks has enabled us to have a higher combined percentage of long and spotlight oral presentations. Second, we have continued the 2016 innovation of four days for the main conference, instead of three, but reduced the third day to a half-day to allow time for relaxation and mental regrouping. In recent post-conference surveys, the community has voted overwhelmingly for a four-day main conference.

Continuing the tradition established in CVPR 2016, we are providing an exciting, "trade-show" like atmosphere to foster maximal visibility and exposure for each onsite exhibitor from promising startups and creative standouts to the biggest industry leaders. Over 110 companies are showcasing their technologies at CVPR 2017 and demonstrating the impact that their hardware and software products are

Message from the General and Program Chairs

having on a wide range of industries. Many of these companies are sponsoring the conference through a wide range of promotional mechanisms, resulting in a sponsorship funding increase of more than 50% from 2016. The conference would not be possible in its current form without the generosity of our corporate partners and their support is deeply appreciated.

Last but not least, we wish to thank all members of the Organizing Committee, the Area Chairs, reviewers, emergency reviewers and authors for the immense amount of hard work and professionalism that has gone into making CVPR 17 one of the most important venues in Computer Science. Our thanks also go to the organizers of previous CVPRs, many of whom provided helpful advice and guidance. The organizers are particularly indebted to Jana Košecká and René Vidal for patiently answering many questions. Critical aspects of the paper review process were handled using Microsoft's CMT system, the Toronto Paper Matching System and Researcher.cc and we would like to thank everyone who works on those project teams. We are particularly grateful to Laurent Charlin from TPMS and Ari Kobren from Researcher.cc for their role in helping us to handle the increased scale of our conference. Once again Eric Mortensen has done the seemingly impossible in pulling together the many aspects of the publication process and we are very grateful for all of his hard work. We also want to thank Nicole Finn and her staff at C to C Events for the crucial organizational support that allowed us to stage this conference.

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

Enjoy CVPR 2017 and Hawaiian hospitality!

Program Chairs: **Yanxi Liu, James M. Rehg, Camillo J. Taylor, Ying Wu**

General Chairs: **Rama Chellappa, Anthony Hoogs, Zhengyou Zhang**

CVPR 2017 Organizing Committee

General Chairs:	Rama Chellappa	**Publications Chair:**	Eric Mortensen
	Zhengyou Zhang	**Website Chair:**	Junsong Yuan
	Anthony Hoogs	**Corporate Relations Chairs:**	Matt Turek
Program Chairs:	Jim Rehg		Shuicheng Yan
	Yanxi Liu		Mei Han
	Ying Wu	**Doctoral Consortium Chairs:**	Adriana Kovashka
	Camillo Taylor		Nathan Jacobs
Workshops Chairs:	Jason Corso	**Technology Chair:**	Terry Boult
	Mei Chen	**Publicity Chair:**	Yihong Gong
Tutorials Chairs:	Robert Pless	**Student Volunteers Chair:**	Amit Roy-Chowdhury
	David Crandall	**Logistics Advisor:**	Ginger Boult
Finance Chair:	Walter Scheirer	**Program Coordination Lead:**	Christopher Funk
Demos Chair:	Le Lu		

Friday, July 21

NOTE: Tutorial locations are in the online proceedings and the mobile app.
Tutorial times with an asterisk (*) indicates a default time—no schedule was provided.

0700–1700 Registration (Main Lobby)

0745–0845 Breakfast (Kamehameha II)

1000–1045 Morning Break (Kamehameha II)

1200–1330 Lunch (Kamehameha II)

1530–1615 Afternoon Break (Kamehameha II)

Large-Scale 3D Modeling From Crowdsourced Data

Organizers: Jan-Michael Frahm
Enrique Dunn
Marc Pollefeys
Jared Heinly
Johannes L. Schönberger

Time: 0900-1730 (Full Day)

Description: Large-scale image-based 3D modeling has been a major goal of computer vision, enabling a wide range of applications including virtual reality, image-based localization, and autonomous navigation. One of the most diverse data sources for modeling is Internet photo collections. In the last decade, the computer vision community has made tremendous progress in large-scale structure-from-motion and multi-view stereo from Internet datasets. However, utilizing this wealth of information for 3D modeling remains a challenging problem due to the ever-increasing amount of image data. In a short period of time, research in large-scale modeling has progressed from modeling using several thousand images, to modeling from city-scale datasets of several million, and recently to reconstructing an Internet-scale dataset comprising 100 million images. This tutorial will present the main underlying technologies enabling these innovations.

Schedule:

0900 Motivation & Basic principles
1000 Morning Break (Kamehameha II)
1045 Basic Principles (continued)
1115 Sparse Modeling
1215 Lunch (Kamehameha II)
1345 Dense Modeling
1500 Large-Scale Modeling Pipeline
1530 Afternoon Break (Kamehameha II)
1615 Large-Scale Modeling Pipeline (continued)
1700 Resources for Large-Scale Modeling
1730 Closing Remarks

Computer Vision for Automated Driving in MATLAB

Organizers: Witek Jachimczyk
Anand Raja
Avi Nehemiah

Time: 0830-1200* (Half Day — Morning)

Description: In recent years, the development of autonomous vehicles has generated an enormous amount of interest. This tutorial introduces you to practical approaches for the design and verification of automated driving systems using new MATLAB features provided in the Automated Driving System Toolbox†. The tutorial concentrates on the design of a monocular vision system, including camera setup, coordinate system transforms, and object detection with deep learning. We show you tools for processing Lidar data, ground truth labeling, and verification of your algorithms. We also cover tracking and sensor fusion between vision and radar sensors. Also included, is a framework for synthetic data generation to help you verify your tracking algorithms.

† https://www.mathworks.com/products/automated-driving.html

DIY A Multiview Camera System: Panoptic Studio Teardown

Organizers: Hanbyul Joo
Tomas Simon
Hyun Soo Park
Shohei Nobuhara
Yaser Sheikh

Time: 0830-1230 (Half Day — Morning)

Description: The tutorial will cover a wide spectrum of multi-camera systems from micro to macro. We will first cover technical hardware issues that are common across systems, such as synchronization, calibration, and data communications, and then we will discuss hardware design factors, such as camera placement, resolution, and framerate, which are strongly related to visual representations. We will also discuss algorithmic challenges associated with the design factors, e.g., matching, tracking, and reconstruction. Three case studies will be conducted: DRZ (submillimeter, hair and eye), MPII (centimeter- meter, single person), and Panoptic Studio (meter, more than 5 people). In particular, we will use the Panoptic Studio as a primary example where we will demonstrate a modular system at the venue. Three distinguished speakers from CMU, MPII, and DRZ are invited. To this end, we will release 3D data and its computational challenges regarding dynamic scene matching, 3D reconstruction, and micro pose/activity recognition.

Schedule:

0830 **Invited Talk:** TBA, *Takeo Kanade (CMU)*
0855 Panoptic Studio: Introduction
0920 **Invited Talk:** TBA, *Christian Theobalt (MPI Informatik)*
0950 Panoptic Studio: System

1020 Morning Break (Kamehameha II)

1050 **Invited Talk:** TBA, *Thabo Beeler, Derek Bradley (Disney Research Zurich)*
1120 Panoptic Studio: Applications
1150 Multiview System Demo

Mathematics of Deep Learning

Organizers: René Vidal
Joan Bruna
Raja Giryes

Time: 0900-1200 (Half Day — Morning)

Description: The past few years have seen a dramatic increase in the performance of recognition systems thanks to the introduction of deep networks for representation learning. However, the mathematical reasons for this success remain elusive. This tutorial will review recent work that aims to provide a mathematical justification for properties of special classes of deep networks, such as invariance, stability, generalization and global optimality of the learned representations.

Schedule:

0900 Introduction, *René Vidal (Johns Hopkins Univ.)*
0915 Global Optimality in Deep Learning, *René Vidal (Johns Hopkins Univ.)*

1000 Morning Break (Kamehameha II)

1030 Structure-Based Theory for Deep Learning, *Raja Giryes (Tel Aviv Univ.)*
1115 Signal Recovery from Scattering Convolutional Networks, *Joan Bruna (NYU)*

OpenCV 3.x New Functionality & Optimizations

Organizers: Vadim Pisarevsky
Yury Gorbachev
Alexander Alekhin

Time: 0830-1200* (Half Day — Morning)

Description: This tutorial provides an in-depth introduction into the new OpenCV functionality developed over the last two years. This functionality includes (1) deep learning, including support for classification, semantic segmentation and object detection networks, (2) advanced tracking algorithms, (3) improved dense optical flow & stereo algorithms, (4) improved text detection, and (5) support for augmented reality through SFM and ArUco modules.

Detailed examples demonstrate the significant performance advancements of OpenCV 3.x through universal and platform-specific intrinsics, enhanced T-API (OpenCL acceleration layer) with OpenCL 2.0 support, and interoperability with OpenVX. A practical hands-on session will allow participants will experiment with the new functionality and acceleration techniques. For this, we will prepare and give out USB drives with pre-configured virtual machines with OpenCV installed.

We will conclude with our view of the future development of OpenCV, what OpenCV 4.0 may look like, some insights on the distributed package model, micro core, high-level performance-portable language for representation of the vision algorithms.

Spectral Methods for 3D Data Analysis

Organizers: Alexander Bronstein
Michael Bronstein
Maks Ovsjanikov

Time: 0830-1200* (Half Day — Morning)

Description: In recent years, the advances in 3D acquisition technology and the availability of affordable 3D scanners (such as Microsoft Kinect or Intel RealSense) have brought problems dealing with 3D shape analysis to the spotlight of computer vision research. In particular, deformable and non-rigid shapes have attracted a growing interest, which has led to rapid development of the field, where state-of-the-art results from very different sciences - theoretical and numerical geometry, optimization, linear algebra, graph theory and machine learning, to mention a few - are applied to find solutions. In this tutorial, we will overview the foundations, tools, numerical methods, and most recent advances in the field. We will use spectral methods as the common denominator, given the usual familiarity of the audience with classical Fourier analysis and signal/image processing methods. New topics will be discussion of relations between image and shape analysis, 3D feature detectors and descriptors and their analogies and differences to methods in image analysis; functional correspondence and relation to sparse coding models; relation to joint approximate diagonalization problems. A special emphasis will be made on practical applications and real-life demos and examples of code will be shown.

Dealing With Reality: Low-Quality Visual Data Processing and Analytics

Organizers: Zhangyang (Atlas) Wang
Haichao Zhang Jiebo Luo

Time: 1400-1700 (Half Day — Afternoon)

Description: While many sophisticated models are developed for visual information processing, very few pay attention to their usability in the presence of (heavy) data quality degradations. For example, video surveillance systems have to rely on cameras of very limited definitions, due to the prohibitive costs of installing high-definition cameras all around, leading to the practical need to recognize objects reliably from very low resolution images. Other quality factors, such as occlusion, motion blur, missing data and label ambiguity, are also ubiquitous in the wild. The tutorial will present a comprehensive and in-depth review, on the recent advances in the robust sensing, processing and understanding of low-quality visual data. As the low data quality appears to be the bottleneck for numerous applications, such as visual recognition, object tracking, medical image processing and 3D vision, our proposed tutorial is expected to be of broad interests to the CVPR community.

Schedule:

1400 Image & Video Restorations With Structural Priors, *Haichao Zhang (Baidu Research)*

1445 Recognition From Very Low-Quality Images and Videos Using Deep Networks, *Zhangyang (Atlas) Wang (Texas A&M Univ.)*

1530 Afternoon Break (Kamehameha II)

1615 Learning Sith Noisy Visual Big Data: Model & Applications, *Jiebo Luo (Univ. of Rochester)*

Deep Learning for Objects and Scenes

Organizers: Bolei Zhou
Xiaogang Wang
Ross Girshick
Kaiming He

Time: 1400-1720 (Half Day — Afternoon)

Description: Deep learning is becoming the driving force for the visual recognition models in computer vision field. The half-day tutorial will focus on providing a high-level summary of the recent work on deep learning models for visual recognition of objects and scenes. The goal is to share some of the lessons and experiences learned by the speakers who specialized in various topics of visual recognition. The tutorial will be composed of four lectures given by each one of the speakers. Each lecture will cover one specific topic of deep learning for visual recognition, from general deep representation architecture and network intepretability, object detection, scene parsing, to video detection. The four lectures will be unified into a coherent tutorial of recent progress in visual recognition, for the wide audiences in computer vision community.

Schedule:

1400 Welcome

1410 Learning Deep Representations for Visual Recognition, *Kaiming He (Facebook AI Research)*

1450 Deep Learning for Object Detection, *Ross Girshick (Facebook AI Research)*

1530 Afternoon Break (Kamehameha II)

1600 Towards Deeper Scene Understanding and Network Intepretability, *Bolei Zhou (MIT)*

1640 Deep Learning for Video Analysis, *Xiaogang Wang (Chinese Univ. of Hong Kong)*

Geometric Deep Learning on Graphs and Manifolds

Organizers: Michael Bronstein
Joan Bruna
Arthur Szlam
Xavier Bresson
Yann LeCun

Time: 1330-1700* (Half Day — Afternoon)

Description: Recently, deep learning methods have achieved unprecedented performance on notoriously hard computer vision problems, leading to an overwhelming trend in the community of departing from axiomatic models towards generic learnable approaches. However, so far deep learning research has mainly focused on Euclidean-structured data (such as acoustic signals, images, or videos). Many applications in computer vision and pattern recognition have to deal with non-Euclidean structured data, such as graphs and manifolds. Such geometric data are becoming increasingly important in computer graphics and 3D vision, sensor networks, biomedicine, and recommendation systems. The adoption of deep learning in these fields has been lagging behind until recently, primarily since the non-Euclidean nature of objects dealt with makes the very definition of basic operations used in deep networks rather elusive. The purpose of this tutorial is to overview the basic ideas of deep learning on graphs and manifolds, existing methods, novel applications, and future research directions.

Notes:

Local Feature Extraction and Learning for Computer Vision

Organizers: Bin Fan
Jiwen Lu
Pascal Fua

Time: 1400-1730 (Half Day — Afternoon)

Description: Local feature is at the core of many computer vision tasks and has played an important role in the development of various computer vision applications. A large number of local features oriented for different tasks have been proposed over the past two decades by using either hand-crafted methods or learning-based techniques. This tutorial will give an extensive introduction of the latest advances on this topic by dividing them into two categories: those were proposed for high matching accuracy, and those were introduced for high efficiency performance. Typical computer vision applications based on these introduced local features will also be discussed.

Schedule:

1400 Introduction and Brief Review of Classical Feature Descriptors, *Pascal Fua (EPFL)*

1430 Modern Descriptors: Towards High Matching Performance, *Bin Fan (Chinese Academy of Sciences)*

1530 Afternoon Break (Kamehameha II)

1600 Modern Descriptors: Towards High Matching Efficiency, *Jiwen Lu (Tsinghua Univ.)*

1710 Summary & Discussion, *Pascal Fua (EPFL), Bin Fan (Chinese Academy of Sciences), Jiwen Lu (Tsinghua Uni.)*

Motion Averaging: A Framework for Efficient and Accurate Large-Scale Camera Estimation in 3D Vision

Organizer: Venu Madhav Govindu

Time: 1400-1800 (Half Day — Afternoon)

Description: In recent years there has been growing interest in large-scale 3D reconstruction using both RGB and depth cameras. The concomitant need for accuracy, efficiency and scalability in camera motion estimation is addressed by the framework of motion averaging. Given many relative motion estimates between pairs of cameras, motion averaging solves for the 3D motions of individual cameras. The efficacy of motion averaging has attracted research interest leading to significant theoretical and algorithmic maturity. Owing to its major advantages over conventional approaches, motion averaging is now utilised in many 3D reconstruction pipelines. This tutorial will provide a comprehensive introduction to motion averaging in 3D vision. An intuitive and systematic understanding of the underlying geometry of matrix Lie groups will be developed. A comparative classification and summarization of various motion averaging methods will be presented. In addition, this tutorial will provide a clear exposition of algorithms and best practices. Along with developing a clear understanding of the state-of-the-art, this tutorial will aim to enable researchers to utilise motion averaging principles in novel contexts of large-scale structure-from-motion as well as dense 3D modeling using depth cameras.

Schedule:

1400 Introduction
Theory & Formulation
Intrinsic Methods

1545 Afternoon Break (Kamehameha II)

1615 Robustness
Extrinsic Methods
Rotation Averaging
Translation Averaging

1715 Hierarchical SfM
Euclidean Motion Averaging
Conclusion

Friday, July 21

NOTE: Workshop locations are in the online proceedings and the mobile app

0700–1700 Registration (Main Lobby)

0745–0845 Breakfast (Kamehameha II)

1000–1045 Morning Break (Kamehameha II)

1200–1330 Lunch (Kamehameha II)

1530–1615 Afternoon Break (Kamehameha II)

Vision Meets Cognition: Functionality, Physics, Intentionality and Causality

Organizers: Yixin Zhu
Chenfanfu Jiang
Lap-Fai (Craig) Yu
Yibiao Zhao
Tao Gao
Ping Wei
Peter Battaglia

Schedule: Full Day

0830 **Welcome**

0840 **Invited Talk:** What Eye-Gaze Tells Us About the Image, *Dimitris Samaras (Stony Brook)*

0905 **Keynote Talk:** TBA

0945 **Invited Talk:** Crack the Autonomous Driving Puzzle From Vision and Cognition Perspective, *Yibiao Zhao (MIT)*

1010 Morning Break (Kamehameha II)

1020 **Invited Talk:** Misconceptions in Artificial Intelligence and the Tasks Forward, *Seng-Beng Ho (Inst. of High Performance Computing)*

1045 **Invited Talk:** Learning to Reason About Images, *Kate Saenko (Boston Univ.)*

1110 **Keynote Talk:** Dark, Beyond Deep, *Song-Chun Zhu (UCLA)*

1150 **Poster Session**

- What Will I Do Next? The Intention From Motion Experiment, *Andrea Zunino, Jacopo Cavazza, Atesh Koul, Andrea Cavallo, Cristina Becchio, Vittorio Murino*
- The Role of Synchronic Causal Conditions in Visual Knowledge Learning, *Seng-Beng Ho*
- Joint 3D Human Motion Capture and Physical Analysis From Monocular Videos, *Petrissa Zell, Bastian Wandt, Bodo Rosenhahn*
- Attention-Based Natural Language Person Retrieval, *Tao Zhou, Muhao Chen, Jie Yu, Demetri Terzopoulos*
- AcFR: Active Face Recognition Using Convolutional Neural Networks, *Masaki Nakada, Han Wang, Demetri Terzopoulos*
- Automated Layout Synthesis and Visualization From Images of Interior or Exterior Spaces, *Tomer Weiss, Masaki Nakada, Demetri Terzopoulos*
- Inferring Hidden Statuses and Actions in Video by Causal Reasoning, *Amy Fire, Song-Chun Zhu*

1230 Lunch (Kamehameha II)

1330 **Invited Talk:** TBA, *Elias Bareinboim (Purdue Univ.)*

1355 **Invited Talk:** TBA, *Joseph Lim (USC)*

1420 **Keynote Talk:** TBA, *Abhinav Gupta (Carnegie Mellon Univ.)*

1500 **Invited Talk:** TBA, *Chaz Firestone (Johns Hopkins Univ.)*

1525 Afternoon Break (Kamehameha II)

1535 **Invited Talk:** "Theory of Mind" From Videos, *Tao Gao (GE Research)*

1600 **Invited Talk:** Synthesizing 3D Shapes via Modeling Multi-View Depth Maps and Silhouettes with Deep Generative Networks, *Amir Arsalan Soltani (MIT)*

1625 **Invited Talk:** Inferring Human Interaction From Motion Trajectories, *Tianmin Shu (UCLA)*

1650 **Invited Talk:** Deep Predictive Model for Autonomous Driving, *Wongun Choi (NEC Lab)*

1715 **Closing Remarks**

1740 **Poster Session** (continued)

Computer Vision in Sports

Organizers: Thomas B. Moeslund
Graham Thomas
Adrian Hilton
Peter Carr
Rikke Gade

Schedule: Full Day

0830 **Welcome**

0840 **Invited Talk:** What Is the Best Representation for Representation Learning on Player Tracking Data? *Patrick Lucey (STATS)*

S1: Video Summarization & Scoring (0925-1025)

0925 Auto-Curation and Personalization of Sports Highlights Through Multimodal Excitement Measures, *Michele Merler, Dhiraj Joshi, Quoc-Bao Nguyen, Stephen Hammer, John Kent, John R. Smith, Rogerio S. Feris*

0945 Singlets: Multi-Resolution Motion Singularities for Soccer Video Abstraction, *Katy Blanc, Diane Lingrand, Frédéric Precioso*

1005 Learning to Score Olympic Events, *Paritosh Parmar, Brendan Tran Morris*

1025 Morning Break (Kamehameha II)

S2: Pose Estimation & Action Recognition (1045-1205)

1045 Hockey Action Recognition via Integrated Stacked Hourglass Network, *Mehrnaz Fani, Helmut Neher, David A. Clausi, Alexander Wong, John Zelek*

1105 Extraction and Classification of Diving Clips From Continuous Video Footage, *Aiden Nibali, Zhen He, Stuart Morgan, Daniel Greenwood*

1125 Accurate and Efficient 3D Human Pose Estimation Algorithm Using Single Depth Images for Pose Analysis in Golf, *Soonchan Park, Ju Yong Chang, Hyuk Jeong, Jae-Ho Lee, Ji-Young Park*

1145 Athlete Pose Estimation by a Global-Local Network, *Jihye Hwang, Sungheon Park, Nojun Kwak*

1205 Lunch (Kamehameha II)

1330 **Invited Talk:** Pose, Transitions and Actions in Sports Video, *Jim Little (Univ. of British Columbia)*

S3: Posters (1415-1600)

1415 **Poster Spotlights**

1435 **Posters**

- Continuous Video to Simple Signals for Swimming Stroke Detection With Convolutional Neural Networks, *Brandon Victor, Zhen He, Stuart Morgan, Dino Miniutti*
- Application of Computer Vision and Vector Space Model for Tactical Movement Classification in Badminton, *Kokum Weeratunga, Anuja Dharmaratne, Khoo Boon How*
- Automatic Tactical Adjustment in Real-Time: Modeling Adversary Formations With Radon-Cumulative Distribution Transform and Canonical Correlation Analysis, *Amir M. Rahimi, Soheil Kolouri, Rajan Bhattacharyya*
- Classification of Puck Possession Events in Ice Hockey, *Moumita Roy Tora, Jianhui Chen, James J. Little*
- Football Action Recognition Using Hierarchical LSTM, *Takamasa Tsunoda, Yasuhiro Komori, Masakazu Matsugu, Tatsuya Harada*
- Ball 3D Trajectory Reconstruction Without Preliminary Temporal and Geometrical Camera Calibration, *Shogo Miyata, Hideo Saito, Kosuke Takahashi, Dan Mikami, Mariko Isogawa, Hideaki Kimata*
- Deep Learning for Domain-Specific Action Recognition in Tennis, *Silvia Vinyes Mora, William J. Knottenbelt*
- Court-Based Volleyball Video Summarization Focusing on Rally Scene, *Takahiro Itazuri, Tsukasa Fukusato, Shugo Yamaguchi, Shigeo Morishima*
- Measuring Energy Expenditure in Sports by Thermal Video Analysis, *Rikke Gade, Ryan Godsk Larsen, Thomas B. Moeslund*

1530 Afternoon Break (Kamehameha II)

1600 **Invited Talk:** Deeper Understanding of Games: How Far We Can Go With the Broadcast Feed, *Mehrsan Javan (SportLogiq)*

1645 **Best Paper Award & Closing Remarks**

Perception Beyond the Visible Spectrum

Organizer Riad Hammoud
Haibin Ling
Maryam Rahnemoonfar
Yi Ding

Schedule: Full Day

0850 **Welcome:** *Riad I. Hammoud (BAE Systems)*

0900 **Keynote Talk:** Practical Deep Learning, *Leo Tam (NVIDIA)*

0945 **Keynote Talk:** The Hype, Hope and Promise of Learning Machines, *Kevin Priddy (US Air Force Research Lab)*

1030 Morning Break (Kamehameha II)

S1: Session 1 (1045-1200)

1045 Infrared Variation Optimized Deep Convolutional Neural Network for Robust Automatic Ground Target Recognition, *Sungho Kim, Woo-Jin Song, So-Hyun Kim*

1100 RGB-D Scene Labeling With Multimodal Recurrent Neural Networks, *Heng Fan, Xue Mei, Danil Prokhorov, Haibin Ling*

1115 Infrared Image Colorization Based on a Triplet DCGAN Architecture, *Patricia L. Suárez, Angel D. Sappa, Boris X. Vintimilla*

1130 An Algorithm for Parallel Reconstruction of Jointly Sparse Tensors With Applications to Hyperspectral Imaging, *Qun Li, Edgar A. Bernal*

1145 Deep Heterogeneous Face Recognition Networks Based on Cross-Modal Distillation and an Equitable Distance Metric, *Christopher Reale, Hyungtae Lee, Heesung Kwon*

1200 Lunch (Kamehameha II)

S2: Session 2 (1330-1520)

1330 Aerial Vehicle Tracking by Adaptive Fusion of Hyperspectral Likelihood Maps, *Burak Uzkent, Aneesh Rangnekar, Matthew Hoffman*

1345 Fully Convolutional Region Proposal Networks for Multispectral Person Detection, *Daniel König, Michael Adam, Christian Jarvers, Georg Layher, Heiko Neumann, Michael Teutsch*

1400 A Logarithmic X-Ray Imaging Model for Baggage Inspection: Simulation and Object Detection, *Domingo Mery, Aggelos K. Katsaggelos*

1415 A Fast Approximate Spectral Unmixing Algorithm Based on Segmentation, *Jing Ke, Yi Guo, Arcot Sowmya*

1430 The First Automatic Method for Mapping the Pothole in Seagrass, *Maryam Rahnemoonfar, Masoud Yari, Abdullah Rahman, Richard Kline*

1445 Face Presentation Attack With Latex Masks in Multispectral Videos, *Akshay Agarwal, Daksha Yadav, Naman Kohli, Richa Singh, Mayank Vatsa, Afzel Noore*

1500 Privacy-Preserving Understanding of Human Body Orientation for Smart Meetings, *Indrani Bhattacharya, Noam Eshed, Richard J. Radke*

1520 Afternoon Break (Kamehameha II)

1545 **Keynote Talk:** Towards Ultimate Plenoptic Imaging, *Jingyi Yu (Univ. of Delaware)*

S3: Session 3 (1630-1715)

1630 Selecting an Optimized COTS Filter Set for Multispectral Plenoptic Sensing, *Timothy Doster, Colin C. Olson, Erin Fleet, Michael Yetzbacher*

1645 A Novel Detection Paradigm and Its Comparison to Statistical and Kernel-Based Anomaly Detection Algorithms for Hyperspectral Imagery, *Colin C. Olson, Timothy Doster*

1700 Learning Spatiotemporal Features for Infrared Action Recognition With 3D Convolutional Neural Networks, *Zhuolin Jiang, Viktor Rozgic, Sancar Adali*

S4: Closing session (1715-1730)

1715 **Best Paper Award** (HW prize from NVIDIA)

1720 **Closing Remarks**

Embedded Vision

Organizers: Martin Humenberger
Swarup Medasani
Ravi Kumar Satzoda
Zoran Nikolic

Schedule: Full Day

0800 **Welcome**

0810 **Invited Talk:** Deep Learning for Autonomous Driving, *Branislav Kisacanin (Nvidia)*

0900 Fast, Accurate Thin-Structure Obstacle Detection for Autonomous Mobile Robots, *Chen Zhou, Jiaolong Yang, Chunshui Zhao, Gang Hua*

0920 Sparse, Quantized, Full Frame CNN for Low Power Embedded Devices, *Manu Mathew, Kumar Desappan, Pramod Kumar Swami, Soyeb Nagori*

0940 Reconstructing Intensity Images From Binary Spatial Gradient Cameras, *Suren Jayasuriya, Orazio Gallo, Jinwei Gu, Timo Aila, Jan Kautz*

1000 Morning Break (Kamehameha II)

1030 **Invited Talk:** Computer Vision for Camera Drones, *Friedrich Fraundorfer (Graz Univ. of Technology)*

1120 Binarized Convolutional Neural Networks With Separable Filters for Efficient Hardware Acceleration, *Jeng-Hau Lin, Tianwei Xing, Ritchie Zhao, Zhiru Zhang, Mani Srivastava, Zhuowen Tu, Rajesh K. Gupta*

1140 Joint Mobile-Cloud Video Stabilization, *Gbolahan S. Adesoye, Oliver Wang*

1200 Embedded Robust Visual Obstacle Detection on Autonomous Lawn Mowers, *Mathias Franzius, Mark Dunn, Nils Einecke, Roman Dirnberger*

1230 Lunch (Kamehameha II)

1330 **Invited Talk:** Bootstrapping AI From Naturalistic Driving, *Stefan Heck (Nauto)*

1420 **Poster Spotlights & Demos**

1440 **Posters**

- Improved Cooperative Stereo Matching for Dynamic Vision Sensors With Ground Truth Evaluation, *Ewa Piatkowska, Jürgen Kogler, Nabil Belbachir, Margrit Gelautz*
- Diagnostic Mechanism and Robustness of Safety Relevant Automotive Deep Convolutional Networks, *Robert Krutsch, Rolf Schlagenhaft*
- Hand Gesture Based Region Marking for Tele-Support Using Wearables, *Archie Gupta, Shreyash Mohatta, Jitender Maurya, Ramakrishna Perla, Ramya Hebbalaguppe, Ehtesham Hassan*
- Even More Confident Predictions With Deep Machine-Learning, *Matteo Poggi, Fabio Tosi, Stefano Mattoccia*
- Low-Complexity Global Motion Estimation for Aerial Vehicles, *Nirmala Ramakrishnan, Alok Prakash, Thambipillai Srikanthan*
- LCDet: Low-Complexity Fully-Convolutional Neural Networks for Object Detection in Embedded Systems, *Subarna Tripathi, Gokce Dane, Byeongkeun Kang, Vasudev Bhaskaran, Truong Nguyen*
- Image-Based Visual Perception and Representation for Collision Avoidance, *Cevahir Cigla, Roland Brockers, Larry Matthies*
- Pruning ConvNets Online for Efficient Specialist Models, *Jia Guo, Miodrag Potkonjak*
- Real-Time Driver Drowsiness Detection for Embedded System Using Model Compression of Deep Neural Networks, *Bhargava Reddy, Ye-Hoon Kim, Sojung Yun, Chanwon Seo, Junik Jang*

1540 Afternoon Break (Kamehameha II)

1600 SqueezeDet: Unified, Small, Low Power Fully Convolutional Neural Networks for Real-Time Object Detection for Autonomous Driving, *Bichen Wu, Forrest Iandola, Peter H. Jin, Kurt Keutzer*

1620 Training Sparse Neural Networks, *Suraj Srinivas, Akshayvarun Subramanya, R. Venkatesh Babu*

1640 SqueezeMap: Fast Pedestrian Detection on a Low-Power Automotive Processor Using Efficient Convolutional Neural Networks, *Rytis Verbickas, Robert Laganiere, Daniel Laroche, Changyun Zhu, Xiaoyin Xu, Ali Ors*

1700 **Closing Remarks & Awards**

Deep Learning for Robotic Vision

Organizers: Anelia Angelova
Gustavo Carneiro
Kevin Murphy
Niko Sünderhauf
Ian Lenz
Vijay Kumar
Jürgen Leitner
Trung T. Pham
Ingmar Posner
Michael Milford
Wolfram Burgard
Ian Reid
Peter Corke

Schedule: Full Day

0800 **Welcome**

0815 **Invited Talk:** *Ashutosh Saxena (Brain of Things)*

0845 **Invited Talk:** *Richard Newcombe (Facebook)*

0915 **Invited Talk:** *Jitendra Malik (Univ. of California Berkeley)*

0945 **Poster Spotlights**

- Learning Robot Activities from First-Person Human Videos Using Convolutional Future Regression, *Jangwon Lee, Michael Ryoo*
- End-to-End Driving in a Realistic Racing Game with Deep Reinforcement Learning, *Etienne Perot, Maximilian Jaritz, Marin Toromanoff, Raoul de Charette*
- Automated risk assessment for scene understanding and domestic robots using RGB-D data and 2.5D CNNs at a patch level, *Rob Dupre, Georgios Tzimiropoulos, Vasileios Argyriou*
- Semantic Instance Segmentation for Autonomous Driving, *Bert De Brabandere, Davy Neven, Luc Van Gool*
- Real-time hand grasp recognition using weakly supervised two-stage convolutional neural networks for understanding manipulation actions, *Ji Woong Kim, Sujeong You, Sang Hoon Ji, Hong Seok Kim*
- Finding Anomalies with Generative Adversarial Networks for a Patrolbot, *Wallace Lawson, Esube Bekele, Keith Sullivan*
- Time-Contrastive Networks: Self-Supervised Learning from Multi-View Observation, *Pierre Sermanet, Corey Lynch, Jasmine Hsu, Sergey Levine*
- Curiosity-driven Exploration by Self-supervised Prediction, *Deepak Pathak, Pulkit Agrawal, Alyosha Efros, Trevor Darrell*

1015 Morning Break (Kamehameha II)

1045 **Poster Session** (includes spotlights above)

- Leveraging Deep Reinforcement Learning for Reaching Robotic Tasks, Kapil Katyal, I-Jeng Wang, Philippe Burlina
- Hand Movement Prediction Based Collision-Free Human-Robot Interaction, Yiwei Wang, Xin Ye, Yezhou Yang, Wenlong Zhang
- 3D Pose Regression Using Convolutional Neural Networks, Siddharth Mahendran, Haider Ali, René Vidal
- Tuning Modular Networks With Weighted Losses for Hand-Eye Coordination, Fangyi Zhang, Jürgen Leitner, Michael Milford, Peter I. Corke
- Episode-Based Active Learning With Bayesian Neural Networks, Feras Dayoub, Niko Sünderhauf, Peter I. Corke
- Detecting and Grouping Identical Objects for Region Proposal and Classification, Wim Abbeloos, Sergio Caccamo, Esra Ataer-Cansizoglu, Yuichi Taguchi, Chen Feng, Teng-Yok Lee

1200 Lunch (Kamehameha II)

1330 **Invited Talk:** *Honglak Lee (Univ. of Michigan/Google Brain)*

1400 **Invited Talk:** *Jianxiong Xiao (Auto X)*

1430 **Invited Talk:** *Dieter Fox (Univ. of Washington)*

1500 **Invited Talk:** *Raquel Urtasun (Univ. of Toronto)*

1530 Afternoon Break (Kamehameha II)

1600 **Invited Talk:** *Abhinav Gupta (Carnegie Melon Univ.)*

1630 **Invited Talk:** *Raia Hadsell (Google Deepmind)*

1700 **Panel Discussion:** All Speakers

1755 **Closing Remarks**

Biometrics

Organizers: Bir Bhanu
Ajay Kumar

Schedule: Full Day

S1: Soft Biometrics (0800-0920)

Chair: Xiaoming Liu *(Michigan State Univ.)*

0800 Age Estimation Guided Convolutional Neural Network for Age-Invariant Face Recognition, *Tianyue Zheng, Weihong Deng, Jiani Hu*

0820 Deep LDA-Pruned Nets for Efficient Facial Gender Classification, *Qing Tian, Tal Arbel, James J. Clark*

0840 Adaptive Deep Metric Learning for Identity-Aware Facial Expression Recognition, *Xiaofeng Liu, B. V. K. Vijaya Kumar, Jane You, Ping Jia*

0900 GaitGAN: Invariant Gait Feature Extraction Using Generative Adversarial Networks, *Shiqi Yu, Haifeng Chen, Edel B. García Reyes, Norman Poh*

S2: Biometrics Aging & Identification (0920-1000)

Chair: Richa Singh *(IIIT Delhi)*

0920 Component Biologically Inspired Features With Moving Segmentation for Age Estimation, *Gee-Sern Jison Hsu, Yi-Tseng Cheng, Choon Ching Ng, Moi Hoon Yap*

0930 Face Recognition Performance Under Aging, *Debayan Deb, Lacey Best-Rowden, Anil K. Jain*

0940 Predicting Face Recognition Performance in Unconstrained Environments, *P. Jonathon Phillips, Amy N. Yates, J. Ross Beveridge, Geof Givens*

0950 Person Re-Identification for Improved Multi-Person Multi-Camera Tracking by Continuous Entity Association, *Neeti Narayan, Nishant Sankaran, Devansh Arpit, Karthik Dantu, Srirangaraj Setlur, Venu Govindaraju*

1000 Morning Break (Kamehameha II) **& Poster Session**

1030 **Invited Talk:** 3D Morphable Face Model as Prior Knowledge for Face Analysis and Recognition, *Josef Kittler (Univ. of Surrey)*

S3: Face Recognition I (1130-1230)

Chair: Gang Hua *(Microsoft Research Asia)*

1130 Toward Open-Set Face Recognition, *Manuel Günther, Steve Cruz, Ethan M. Rudd, Terrance E. Boult*

1150 Investigating Nuisance Factors in Face Recognition With DCNN Representation, *Claudio Ferrari, Giuseppe Lisanti, Stefano Berretti, Alberto Del Bimbo*

1210 IARPA Janus Benchmark-B Face Dataset, *Cameron Whitelam, Emma Taborsky, Austin Blanton, Brianna Maze, Jocelyn Adams, Tim Miller, Nathan Kalka, Anil K. Jain, James A. Duncan, Kristen Allen, Jordan Cheney, Patrick Grother*

S4: Face Recognition II (1230-1300)

Chair: Hugo Proença *(Univ. of Beira Interior)*

1230 Efficient Image Set Classification Using Linear Regression Based Image Reconstruction, *Syed A. A. Shah, Uzair Nadeem, Mohammed Bennamoun, Ferdous Sohel, Roberto Togneri*

1240 Deep Convolutional Neural Network Using Triplets of Faces, Deep Ensemble, and Score-Level Fusion for Face Recognition, *Bong-Nam Kang, Yonghyun Kim, Daijin Kim*

1250 Transfer Learning Based Evolutionary Algorithm for Composite Face Sketch Recognition, *Tarang Chugh, Maneet Singh, Shruti Nagpal, Richa Singh, Mayank Vatsa*

1300 Lunch (Kamehameha II)

S5: Fingerprint & Signature Verification (1400-1500)

Chair: Mayank Vatsa *(IIIT Delhi)*

1400 Analysis, Comparison, and Assessment of Latent Fingerprint Image Preprocessing, *Haiying Guan, Paul Lee, Andrew Dienstfrey, Mary Theofanos, Curtis Lamp, Brian Stanton, Matthew T. Schwarz*

1420 Parsimonious Coding and Verification of Offline Handwritten Signatures, *Elias N. Zois, Ilias Theodorakopoulos, Dimitrios Tsourounis, George Economou*

1440 Robust Verification With Subsurface Fingerprint Recognition Using Full Field Optical Coherence

Tomography, *Kiran B. Raja, Egidijus Auksorius, R. Raghavendra, A. Claude Boccara, Christoph Busch*

S6: Iris Recognition & Spoof Attack Detection (1500-1530)

Chair: Gee-Sern Hsu *(National Taiwan Univ. of Science & Technology)*

1500 Iris Super-Resolution Using Iterative Neighbor Embedding, *Fernando Alonso-Fernandez, Reuben A. Farrugia, Josef Bigun*

1510 Iris Liveness Detection by Relative Distance Comparisons, *Federico Pala, Bir Bhanu*

1520 Face Presentation Attack Detection by Exploring Spectral Signatures, *R. Raghavendra, Kiran B. Raja, Sushma Venkatesh, Christoph Busch*

1530 Afternoon Break (Kamehameha II) **& Poster Session**

1600 **Panel Session:** Key Problems in Biometrics for the Next Five Years, *Chair: Jonathon Phillips (NIST)*

1700 **Awards, Valedictory & Closing Remarks:** *Bir Bhanu*

Diff-CVML: Differential Geometry in Computer Vision and Machine Learning

Organizers: Hassen Drira
Mehrtash Harandi
Sebastian A. Kurtek
Minh Hà Quang
Vittorio Murino
Pavan Turaga

Schedule: Full Day

0900 **Opening Remarks:** Diff-CVML Organizers

0915 **Keynote Talk:** Group Action Induced Distances for Averaging and Clustering Linear Dynamical Systems, *René Vidal (The Johns Hopkins Univ.)*

1000 Morning Break (Kamehameha II)

1030 **Keynote Talk:** Barycentric Subspace Analysis: An Extension of PCA to Manifolds, *Xavier Pennec (INRIA, Sophia Antipolis)*

S1: Oral Session 1 (1115-1155)

1115 The Square Root Velocity Framework for Curves in a Homogeneous Space, *Zhe Su, Eric Klassen, Martin Bauer*

1135 Poisson Disk Sampling on the Grassmannnian: Applications in Subspace Optimization, *Rushil Anirudh, Bhavya Kailkhura, Jayaraman J. Thiagarajan, Peer-Timo Bremer*

1155 Lunch (Kamehameha II)

1330 **Keynote Talk:** Classifying Object Manifolds: Statistical Mechanics and Conic Geometry, *Daniel Lee (Univ. of Pennsylvania)*

S2: Oral Session 2 (1415-1535)

1415 Riemannian Variance Filtering: An Independent Filtering Scheme for Statistical Tests on Manifold-Valued Data, *Ligang Zheng, Hyunwoo J. Kim, Nagesh Adluru, Michael A. Newton, Vikas Singh*

1435 Measuring Glide-Reflection Symmetry in Human Movements Using Elastic Shape Analysis, *Qiao Wang, Chaitanya Potaraju, Pavan Turaga*

1455 Learning Shape Trends: Parameter Estimation in Diffusions on Shape Manifolds, *Valentina Staneva, Laurent Younes*

1515 A Riemannian Framework for Linear and Quadratic Discriminant Analysis on the Tangent Space of Shapes, *Susovan Pal, Roger P. Woods, Suchit Panjiyar, Elizabeth Sowell, Katherine L. Narr, Shantanu H. Joshi*

1535 Afternoon Break (Kamehameha II)

1600 **Keynote Talk:** Stochastic Models and Lie Groups in Robotics, *Gregory Chirikjian (The Johns Hopkins Univ.)*

1645 **Poster Session**

- Signal Classification in Quotient Spaces via Globally Optimal Variational Calculus, *Gregory S. Chirikjian*
- Manifold Guided Label Transfer for Deep Domain Adaptation, *Breton Minnehan, Andreas Savakis*

Computer Vision for Microscopy Image Analysis

Organizers: Peter Bajcsy
Dimitris N. Metaxas

Schedule: Full Day

0945 **Welcome**

1000 **Invited Talk:** Image Analytics for Endomicroscopy and Digital Holographic Microscopy, *Terrence Chen (Siemens)*

1030 **Invited Talk:** Deep Learning Applied to Challenges in Detecting Melanoma Cells in Lymph Node Biopsies, *Charles Law (Kitware)*

1100 Morning Break (Kamehameha II)

1115 CNN Based Yeast Cell Segmentation in Multi-Modal Fluorescent Microscopy Data, *Ali Selman Aydin, Abhinandan Dubey, Daniel Dovrat, Amir Aharoni, Roy Shilkrot*

1130 Classification and Retrieval of Digital Pathology Scans: A New Dataset, *Morteza Babaie, Shivam Kalra, Aditya Sriram, Christopher Mitcheltree, Shujin Zhu, Amin Khatami, Shahryar Rahnamayan, Hamid R. Tizhoosh*

1145 Breast Cancer Histopathological Image Classification: Is Magnification Important? *Vibha Gupta, Arnav Bhavsar*

1200 Delineation of Skin Strata in Reflectance Confocal Microscopy Images With Recurrent Convolutional Networks, *Alican Bozkurt, Trevor Gale, Kivanc Kose, Christi Alessi-Fox, Dana H. Brooks, Milind Rajadhyaksha, Jennifer Dy*

1215 Crowdsourcing for Chromosome Segmentation and Deep Classification, *Monika Sharma, Oindrila Saha, Anand Sriraman, Ramya Hebbalaguppe, Lovekesh Vig, Shirish Karande*

1230 Lunch (Kamehameha II)

1345 **Posters & Demo Session**

- Generative Adversarial Learning for Reducing Manual Annotation in Semantic Segmentation on Large Scale Miscroscopy Images, *Avisek Lahiri*
- Microscopic Blood Smear Segmentation and Classification using Deep ContourAware CNN and Extreme Machine Learning, *Imran Razzak, Saeeda Naz*
- Applying Faster R-CNN for Object Detection on Malaria Images, *Jane Hung, Allen Goodman, Anne Carpenter*
- An Early Experience Toward Developing Computer Aided Diagnosis for Gram-Stained Smears Images, *Johanna Carvajal, Daniel Smith, Kun Zhao, Arnold Wiliem, Paul Finucane,Peter Hobson, Anthony Jennings, Rodney McDougall, Brian Lovell*
- Looking Under the Hood: Deep Neural Network Visualization to Interpret Whole-Slide Image Analysis Outcomes for Colorectal Polyps, *Bruno Korbar,Andrea Olofson,Allen Miraflor,Catherine Nicka,Matthew Suriawinata,Lorenzo Torresani,Arief Suriawinata,Saeed Hassanpour*
- High-magnification Multi-views Based Classification of Breast Fine Needle Aspiration Cytology Cell Samples using Fusion of Decisions from Deep Convolutional Networks, *Hrushikesh Garud, Sri Phani Krishna Karri, Debdoot Sheet, Arindam Ghosh, Ashok Maity, Dr. Jyotirmoy Chatterjee, Manjunatha Mahadevappa, Ajoy Kumar Ray*
- Nuclei Segmentation of Fluorescence Microscopy Images Using Three Dimensional Convolutional Neural Networks, *David Ho, Chichen Fu, Paul Salama, Kenneth Dunn,Edward Delp*
- ssEMnet: Serial-Section Electron Microscopy Image Registration Using a Spatial Transformer Network With Learned Features, *Inwan Yoo, David Hildebrand, Willie Tobin, Wei-chung Lee, Won-Ki Jeong*
- Detection of Elongated Touching Cells With Inhomogeneous Illumination Using a Stack of Conditional Random Fields, *A. Memariani, S. Upadhyay, C. Nikou, B.T. Endres, E. Basseres, K.W. Garey, I.A. Kakadiaris*
- Tools for Exploratory Shape Analysis of Cell Populations, *Gunay Dogan, Javier Bernal, Charles R. Hagwood, Eve Fleisig*

1430 **Invited Talk:** Development of Non-Invasive Release Criteria for Clinical Retinal Cell Therapies, *Nathan Hotaling (NIH)*

1500 **Invited Talk:** Enabling Stem Cell Characterization from Large Microscopy Images, *Peter Bajcsy (NIST)*

1530 DeepXScope: Segmenting Microscopy Images with a Deep Neural Network, *Philip Saponaro, Wayne Treible, Abhishek Kolagunda, Timothy Chaya, Jeffrey Caplan, Chandra Kambhamettu, Randall Wisser*

1545 Transferring Microscopy Image Modalities with Conditional Generative Adversarial Networks, *Liang Han, Zhaozheng Yin*

1600 Fast Neural Cell Detection Using Light-Weight SSD Neural Network, *Jingru Yi, Pengxiang Wu, Daniel Hoeppner, Dimitris Metaxas*

1615 A Level Set Method for Gland Segmentation, *Chen Wang, Hong Bu, Ji Bao, Chunming Li*

1630 Afternoon Break (Kamehameha II)

1645 **Invited Talk:** Human Computation Approaches to Microscopy Image Analysis, *Margrit Betke (Boston Univ.)*

1715 **Invited Talk:** Microscopy Image Analysis: Optics, Algorithms and Community, *Zhaozheng Yin (Missouri S&T)*

1745 **Awards & Closing Remarks**

Traffic Surveillance Workshop and Challenge

Organizers: Pierre-Marc Jodoin
Justin Eichel
Andrew Achkar
Thomas B. Moeslund
Nicolas Saunier
Janusz Konrad
Akshaya Mishra
Shaozi Li
Kalle Åström
Zhiming Luo

Schedule: Full Day

0900 **Welcome**

0915 **Invited Talk:** A Synthetic Collection of Annotated Images of Driving Scenarios, *David Vázquez (Univ. Autònoma de Barcelona)*

1000 Morning Break (Kamehameha II)

1045 A Large and Diverse Dataset for Improved Vehicle Make and Model Recognition, *Faezeh Tafazzoli, Hichem Frigui, Keishin Nishiyama*

1105 Evaluating State-Of-The-Art Object Detector on Challenging Traffic Light Data, *Morten B. Jensen, Kamal Nasrollahi, Thomas B. Moeslund*

1125 Slot Cars: 3D Modelling for Improved Visual Traffic Analytics, *Eduardo R. Corral-Soto, James H. Elder*

1145 A Cost-Effective Framework for Automated Vehicle-Pedestrian Near-Miss Detection Through Onboard Monocular Vision, *Ruimin Ke, Jerome Lutin, Jerry Spears, Yinhai Wang*

1205 Lunch (Kamehameha II)

1330 **Invited Talk:** TBA, *Raquel Urtasun (Univ. of Toronto)*

1415 **Challenge and Datasets Description**

1430 EDeN: Ensemble of Deep Networks for Vehicle Classification, *Rajkumar Theagarajan, Federico Pala, Bir Bhanu*

1450 Vehicle Type Classification Using Bagging and Convolutional Neural Network on Multi View Surveillance Image, *Pyong-Kun Kim, Kil-Taek Lim*

1510 Deep Learning-Based Vehicle Classification Using an Ensemble of Local Expert and Global Networks, *Jong Taek Lee, Yunsu Chung*

1530 Afternoon Break (Kamehameha II)

1600 Efficient Scene Layout Aware Object Detection for Traffic Surveillance, *Tao Wang, Xuming He, Songzhi Su, Yin Guan*

1620 ResNet-Based Vehicle Classification and Localization in Traffic Surveillance Systems, *Heechul Jung, Min-Kook Choi, Jihun Jung, Jin-Hee Lee, Soon Kwon, Woo Young Jung*

1640 **Best Paper Award & Final Remarks**

Visual Odometry and Computer Vision Applications Based on Location Clues

Organizers: Guoyu Lu
Yan Yan
Friedrich Fraundorfer
Nicu Sebe
Chandra Kambhamettu

Schedule: Full Day

0900 **Keynote Talk:** TBA, *Raquel Urtasun (Univ. of Toronto)*

0940 **Keynote Talk:** Dense & Direct Methods for 3D Reconstruction & Visual SLAM, *Daniel Cremers (Technical Univ. of Munich)*

1020 Morning Break (Kamehameha II)

1040 Scene-Text-Detection Method Robust Against Orientation and Discontiguous Components of Characters, *Rei Endo, Yoshihiko Kawai, Hideki Sumiyoshi, Masanori Sano*

1100 Uncertainty Quantification of Lucas Kanade Feature Track and Application to Visual Odometry, *Xue Iuan Wong, Manoranjan Majji*

1120 Cluster-Wise Ratio Tests for Fast Camera Localization, *Raúl Díaz, Charless C. Fowlkes*

1140 Ground Truth Accuracy and Performance of the Matching Pipeline, *Josef Maier, Martin Humenberger, Oliver Zendel, Markus Vincze*

1200 Lunch (Kamehameha II)

1400 **Keynote Talk:** Visual Indoor Localization for Micro Aerial Vehicles, *Horst Bischof (Graz Univ. of Technology)*

1440 **Keynote Talk:** Robust and Efficient Large-Scale 3D Reconstruction From Crowd Sourced Imagery With Dynamic Scene Elements, *Jan-Michael Frahm (Univ. of North Carolina at Chapel Hill)*

1520 Afternoon Break (Kamehameha II)

1540 **Keynote Talk:** TBA, *Davide Scaramuzza (Univ. of Zurich)*

1620 EgoTracker: Pedestrian Tracking With Re-Identification in Egocentric Videos, *Jyoti Nigam, Renu M. Rameshan*

1640 Probabilistic Global Scale Estimation for MonoSLAM Based on Generic Object Detection, *Edgar Sucar, Jean-Bernard Hayet*

1700 **Closing Remarks**

Language and Vision

Organizers: Andrei Barbu
Tao Mei
Siddharth Narayanaswamy
Puneet Kumar Dokania
Quanshi Zhang
Nishant Shukla
Jiebo Luo
Rahul Sukthankar

Schedule: Full Day

0815 **Poster Setup**

0845 **Welcome**

0900 **Invited Talk:** TBA, *Lawson Wong, Stefanie Tellex (Brown Univ.)*

0930 **Invited Talk:** Visual Dialog: Towards AI Agents That Can See, Talk, and Act, *Dhruv Batra (Georgia Tech)*

1000 **Invited Talk:** TBA, *Alan Yuille (Johns Hopkins Univ.)*

1030 Morning Break (Kamehameha II)

1100 **Invited Talk:** TBA, *David Hogg (Univ. of Leeds)*

1130 **Invited Talk:** Human Language as a Code for Thought, *Ev Fedorenko (MIT)*

1200 **Poster Highlights**

1230 Lunch (Kamehameha II)

1400 **Poster Session**

1500 **Invited Talk:** Color Naming Across Languages Reflects Color Use, *Ted Gibson (MIT)*

1530 **Invited Talk:** Towards Theory of AI's Mind, *Devi Parikh (Georgia Tech)*

1600 Afternoon Break (Kamehameha II)

1630 **Invited Talk:** TBA, *Song-Chun Zhu (UCLA)*

1700 **Invited Talk:** Video-To-Text Corpus, *Tao Mei (Microsoft Research Asia)*

1710 **Invited Talk:** Robots That Communicate, *Andrei Barbu (MIT)*

1720 **Vision-Language Panel**

1800 **Closing Remarks**

New Trends in Image Restoration and Enhancement & Example-Based Single Image Super-Resolution Challenge

Organizers: Radu Timofte
Eirikur Agustsson
Ming-Hsuan Yang
Lei Zhang
Luc Van Gool

Schedule: Full Day

0730 **Poster Setup** (all papers have poster panels for the whole day)

0750 **Opening Remarks**

0800 **Invited Talk:** Unsupervised Image-To-Image Translation Networks, *Jan Kautz (NVIDIA)*

0830 **Invited Talk:** Regularization by Denoising – "The Little Engine That Could", *Peyman Milanfar (Google)*

0900 Locally Adaptive Color Correction for Underwater Image Dehazing and Matching, *Codruta O. Ancuti, Cosmin Ancuti, Christophe De Vleeschouwer, Rafael Garcia*

0915 Depth-Stretch: Enhancing Depth Perception Without Depth, *Hagit Hel-Or, Yacov Hel-Or, Renato Keshet*

0930 FAST: A Framework to Accelerate Super-Resolution Processing on Compressed Videos, *Zhengdong Zhang, Vivienne Sze*

0945 Fast External Denoising Using Pre-Learned Transformations, *Shibin Parameswaran, Enming Luo, Charles-Alban Deledalle, Truong Q. Nguyen*

1000 FormResNet: Formatted Residual Learning for Image Restoration, *Jianbo Jiao, Wei-Chih Tu, Shengfeng He, Rynson W. H. Lau*

1015 Morning Break (Kamehameha II) **& Poster Session**

- Exploiting Reflectional and Rotational Invariance in Single Image Superresolution, *Simon Donné, Laurens Meeus, Hiep Quang Luong, Bart Goossens, Wilfried Philips*
- Image Super Resolution Based on Fusing Multiple Convolution Neural Networks, *Haoyu Ren, Mostafa El-Khamy, Jungwon Lee*
- PaletteNet: Image Recolorization With Given Color Palette, *Junho Cho, Sangdoo Yun, Kyoung Mu Lee, Jin Young Choi*
- SRHRF+: Self-Example Enhanced Single Image Super-Resolution Using Hierarchical Random Forests, *Jun-Jie Huang, Tianrui Liu, Pier Luigi Dragotti, Tania Stathaki*
- Image Denoising via CNNs: An Adversarial Approach, *Nithish Divakar, R. Venkatesh Babu*
- Multi-Resolution Data Fusion for Super-Resolution Electron Microscopy, *Suhas Sreehari, S. V. Venkatakrishnan, Katherine L. Bouman, Jeffrey P. Simmons, Lawrence F. Drummy, Charles A. Bouman*
- Fast and Accurate Image Super-Resolution Using a Combined Loss, *Jinchang Xu, Yu Zhao, Yuan Dong, Hongliang Bai*
- Deep Wavelet Prediction for Image Super-Resolution, *Tiantong Guo, Hojjat Seyed Mousavi, Tiep Huu Vu, Vishal Monga*

1100 **Invited Talk:** Near-Infrared for Image Enhancement and Restoration, *Sabine Süsstrunk (EPFL)*

1130 **Invited Talk:** Attention-Aware Face Hallucination via Deep Reinforcement Learning, *Liang Lin (SenseTime)*

1200 Lunch (Kamehameha II)

1330 **Invited Talk:** Neural Networks for Image and Video Super Resolution, *Wenzhe Shi & Christian Ledig (Twitter)*

1400 NTIRE 2017 Challenge on Single Image Super-Resolution: Methods and Results, *Radu Timofte, Eirikur Agustsson, Luc Van Gool, Ming-Hsuan Yang, Lei Zhang*

1415 NTIRE 2017 Challenge on Single Image Super-Resolution: Dataset and Study, *Eirikur Agustsson, Radu Timofte*

1430 **Award Ceremony**

1445 Enhanced Deep Residual Networks for Single Image Super-Resolution, *Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, Kyoung Mu Lee*

1500 Beyond Deep Residual Learning for Image Restoration: Persistent Homology-Guided Manifold Simplification, *Woong Bae, Jaejun Yoo, Jong Chul Ye*

1515 A Deep Convolutional Neural Network With Selection Units for Super-Resolution, *Jae-Seok Choi, Munchurl Kim*

1530 Balanced Two-Stage Residual Networks for Image Super-Resolution, *Yuchen Fan, Honghui Shi, Jiahui Yu, Ding Liu, Wei Han, Haichao Yu, Zhangyang Wang, Xinchao Wang, Thomas S. Huang*

1545 Afternoon Break (Kamehameha II) **& Poster Session**

1630 **Invited Talk:** Image Stylization – From Patches to Neural Networks and Back, *Eli Shechtman (Adobe)*

1700 **Invited Talk:** TBA, *Alexei Efros (UC Berkeley)*

1730 **Poster Session**

1830 **Closing Remarks**

Computer Vision in Vehicle Technology and Autonomous Driving Challenge

Organizers: José M. Álvarez
David Vázquez
Tomas Pajdla
Antonio M. López
Gabriel Villalonga
Markus Enzweiler
Fisher Yu
Dequan Wang
John Leonard
Xiaodi Hou
Jianxiong Xiao

Schedule: Full Day

S1: Computer Vision in Vehicle Technology (0810-1235)

0825 **Opening Remarks**

0830 **Invited Talk:** Visual Flow and Control of Aerial Robotic Vehicles, *Rob Mahoney (Australian National Univ.)*

0925 **Invited Talk:** Autonomous, Agile, Vision-Controlled Drones: From Active to Event Vision, *Davide Scaramuzza (Univ. of Zurich)*

0930 **Poster Spotlights**

1000 Morning Break (Kamehameha II) **& Poster Session**

- DriveAHead – A Large-Scale Driver Head Pose Dataset, *Anke Schwarz, Monica Haurilet, Manuel Martinez, Rainer Stiefelhagen*
- The One Hundred Layers Tiramisu: Fully Convolutional DenseNets for Semantic Segmentation, *Simon Jégou, Michal Drozdzal, David Vazquez, Adriana Romero, Yoshua Bengio*
- Rear-Stitched View Panorama: A Low-Power Embedded Implementation for Smart Rear-View Mirrors on Vehicles, *Janice Pan, Vikram Appia, Jesse Villarreal, Lucas Weaver, Do-Kyoung Kwon*
- End-To-End Ego Lane Estimation Based on Sequential Transfer Learning for Self-Driving Cars, *Jiman Kim, Chanjong Park*
- Robust Hand Detection and Classification in Vehicles and in the Wild, *T. Hoang Ngan Le, Kha Gia Quach, Chenchen Zhu, Chi Nhan Duong, Khoa Luu, Marios Savvides*
- Motion Language of Stereo Image Sequence, *Tomoya Kato, Hayato Itoh, Atsushi Imiya*

1030 **Invited Talk:** TBA, *Raquel Urtasun (Univ. of Toronto, Uber ATG)*

1100 **Invited Talk:** TBA, *Mohan M. Trivedi (UC San Diego)*

1130 **Invited Talk:** What Should I Do Here? Visual Semantic Mapping for Automated Automobiles, *Christoph Stiller (Karlsruhe Inst. of Technology)*

1200 **Invited Talk:** Vision & Visualization Technologies for Planetary Rovers, *Gerhard Paar (Joanneum Research)*

1230 Lunch (Kamehameha II)

S2: Autonomous Driving Challenge (1435-1700)

1325 **Session Opening**

1330 **Invited Talk:** TBA, *Trevor Darrell (UC Berkeley)*

1400 **Invited Talk:** TBA, *Jan Becker (Faraday Future)*

1430 **Invited Talk:** TBA, *Alan Yuille (Johns Hopkins Univ.)*

1500 **Invited Talk:** Autonomous Driving Challenge Winner

1525 Afternoon Break (Kamehameha II)

1600 **Invited Talk:** TBA, *Andreas Geiger (MPI)*

1630 **Invited Talk:** TBD

1700 **Panel Discussion:** *Jesse Levinson (Zoox), Urs Muller (NVIDIA), Kai Ni (LeTV Supercar), Kai Yu (Horizon Robotics), Gansha Wu (UiSee)*

1800 **Closing Remarks**

Open Domain Action Recognition Challenge

Organizers: Mohan S. Kankanhalli
An-An Liu
Yongkang Wong

Schedule: Full Day

0900 **Opening Remarks**

0910 **Keynote Talk:** UCF Action Datasets, *Mubarak Shah (Univ. of Central Florida)*

1000 **Invited Talk:** Weak Learning for Scalable Activity Recognition, *Hildegard Kuhne (The Univ. of Bonn)*

1030 Morning Break (Kamehameha II)

1100 Deep Local Video Feature for Action Recognition, *Zhenzhong Lan, Yi Zhu, Alexander G. Hauptmann, Shawn Newsam*

1120 Video Action Recognition Based on Deeper Convolution Networks With Pair-Wise Frame Motion Concatenation, *Yamin Han, Peng Zhang, Tao Zhuo, Wei Huang, Yanning Zhang*

1140 Hand-Object Interaction Detection With Fully Convolutional Networks, *Matthias Schröder, Helge Ritter*

1200 Object State Recognition for Automatic AR-Based Maintenance Guidance, *Pavel Dvorak, Radovan Josth, Elisabetta Delponte*

1220 Lunch (Kamehameha II)

1300 **Poster Session**

1345 **Overview of ODAR Dataset and Challenge Tasks:** *Yongkang Wong*

1400 **Challenge Presentation 1:** TBD

1410 **Challenge Presentation 2:** TBD

1420 **Challenge Presentation 3:** TBD

1430 **Open Discussion**

1450 When Kernel Methods Meet Feature Learning: Log-Covariance Network for Action Recognition From Skeletal Data, *Jacopo Cavazza, Pietro Morerio, Vittorio Murino*

1510 Fast Simplex-HMM for One-Shot Learning Activity Recognition, *Mario Rodriguez, Carlos Orrite, Carlos Medrano, Dimitrios Makris*

1530 Afternoon Break (Kamehameha II)

1600 **Invited Talk:** On the recognition from human action to social relation, *Qianru Sun (Max Planck Inst. for Informatics)*

1630 **Keynote Talk:** Deep Action Parsing Using Synthesized Training Data, *Ling Shao (Univ. of East Anglia)*

1720 **Closing Remarks**

Computational Cameras and Displays

Organizers: Mohit Gupta
Matthew O'Toole
Aswin C. Sankaranarayanan

Schedule: Full Day

0830 **Welcome**

0845 **Keynote Talk:** *Matthias Hullin (Univ. of Bonn)*

0945 **Poster Spotlights**

1000 Morning Break (Kamehameha II)

S1: Computational Cameras (1030-1100)

1030 Intel RealSense Stereoscopic Depth Cameras, *Leonid Keselman, John Iselin Woodfill, Anders Grunnet-Jepsen, Achintya Bhowmik*

1050 Compressive Light Field Reconstructions Using Deep Learning, *Mayank Gupta, Arjun Jauhari, Kuldeep Kulkarni, Suren Jayasuriya, Alyosha Molnar, Pavan Turaga*

1110 **Invited Talk:** TBA

1130 Lunch (Kamehameha II)

S2: Computational Displays (1300-1500)

1300 **Keynote Talk:** *David Luebke (NVIDIA)*

1400 Generating 5D Light Fields in Scattering Media for Representing 3D Images, *Eri Yuasa, Fumihiko Sakaue, Jun Sato*

1420 The Stereoscopic Zoom, *Sergi Pujades, Frédéric Devernay, Laurent Boiron, Rémi Ronfard*

1440 **Invited Talk:** TBA

1500 Afternoon Break (Kamehameha II) **& Poster Session**

1600 **Keynote Talk:** *Vivek Goyal (Boston Univ.)*

1700 **Closing Remarks**

The Bright and Dark Sides of Computer Vision: Challenges and Opportunities for Privacy and Security

Organizers: David J. Crandall
Jan-Michael Frahm
Apu Kapadia

Schedule: Full Day

0820 **Welcome**

0830 **Invited Talk:** Towards a Visual Privacy Advisor: Understanding and Controlling Privacy in Visual Data, *Mario Fritz (Max Planck)*

S1: Oral Session – Attacks (0920-1020)

0920 Deceiving Google's Cloud Video Intelligence API Built for Summarizing Videos, *Hossein Hosseini, Baicen Xiao, Radha Poovendran*

0940 Simple Black-Box Adversarial Attacks on Deep Neural Networks, *Nina Narodytska, Shiva Kasiviswanathan*

1000 I Know That Person: Generative Full Body and Face De-Identification of People in Images, *Karla Brkić, Ivan Sikirić, Tomislav Hrkać, Zoran Kalafatić*

1020 Morning Break (Kamehameha II)

S2: Oral Session – Protecting Privacy (1045-1205)

1045 Protecting Visual Secrets Using Adversarial Nets, *Nisarg Raval, Ashwin Machanavajjhala, Landon P. Cox*

1105 Cartooning for Enhanced Privacy in Lifelogging and Streaming Videos, *Eman T. Hassan, Rakibul Hasan, Patrick Shaffer, David Crandall, Apu Kapadia*

1125 Blur vs. Block: Investigating the Effectiveness of Privacy-Enhancing Obfuscation for Images, *Yifang Li, Nishant Vishwamitra, Bart P. Knijnenburg, Hongxin Hu, Kelly Caine*

1145 ASePPI: Robust Privacy Protection Against De-Anonymization Attacks, *Natacha Ruchaud, Jean-Luc Dugelay*

1205 Lunch (Kamehameha II)

1330 **Invited Talk:** Machine Learning and Privacy: Friends or Foes? *Vitaly Shmatikov (Cornell Tech)*

S3: Oral Session – Ethics & Education (1420-1520)

1420 Trusting the Computer in Computer Vision: A Privacy-Affirming Framework, *Andrew Tzer-Yeu Chen, Morteza Biglari-Abhari, Kevin I-Kai Wang*

1440 Designing a Moral Compass for the Future of Computer Vision Using Speculative Analysis, *Michael Skirpan, Tom Yeh*

1500 Teaching Computer Vision and Its Societal Effects: A Look at Privacy and Security Issues From the Students' Perspective, *Melissa Cote, Alexandra Branzan Albu*

S4: Spotlights (1520-1540)

- Computer Vision Attacks Against 3D CAPTCHAs, *Simon Woo*
- Filter-Amplifier Network for Detecting Integrated Circuit Packages on Printed Circuit Boards, *Zhenhua Chen, David Crandall*
- From Understanding to Controlling Privacy Against Automatic Person Identification in Social Media, *Seong Joon Oh, Mario Fritz, Bernt Schiele*
- Privacy Risks of Using Camera Assisted Tools for People With Visual Impairments, *Taslima Akter, Tousif Ahmed, Kay Connelly, David Crandall, Apu Kapadia*
- Detection Without Recognition for Redaction, *Shagan Sah, Ram Longman, Ameya Shringi, Robert Loce, Majid Rabbani, Raymond Ptucha*
- Privacy-Preserving Human Activity Recognition From Extreme Low Resolution, *Michael Ryoo*
- Computational Privacy Cameras, *Francesco Pittaluga, Koppal Sanjeev, Aleksandar Zivkovic*
- Privacy-Preserving Visual Learning Using Doubly Permuted Homomorphic Encryption, *Ryo Yonetani, Vishnu Boddeti, Kris Kitani, Yoichi Sato*
- Towards Enhancement of Gender Estimation From Fingerprints, *Emanuela Marasco, Emanuele Plebani, Pegah Karimi, Bojan Cukic*
- Vulnerability of Deep Learning Based Gait Biometric Recognition to Adversarial Perturbations, *Vinay Uday Prabhu, John Whaley*
- Smile in the Face of Adversity Much? A Print Based Spoofing Attack, *Vinay Uday Prabhu, John Whaley*

1540 Afternoon Break (Kamehameha II)

1600 **Poster Session**

S5: Oral Session – Opportunities (1700-1800)

1700 Assisting Users in a World Full of Cameras: A Privacy-Aware Infrastructure for Computer Vision Applications, *Anupam Das, Martin Degeling, Xiaoyou Wang, Junjue Wang, Norman Sadeh, Mahadev Satyanarayanan*

1720 Caught Red-Handed: Toward Practical Video-Based Subsequences Matching in the Presence of Real-World Transformations, *Yi Xu, True Price, Fabian Monrose, Jan-Michael Frahm*

1740 Information Hiding in RGB Images Using an Improved Matrix Pattern Approach, *Amirfarhad Nilizadeh, Wojciech Mazurczyk, Cliff Zou, Gary T. Leavens*

1800 On the Effectiveness of Visible Watermarks, *Tali Dekel, Michael Rubinstein, Ce Liu, William T. Freeman*

1820 **Discussion & Concluding Remarks**

Target Re-Identification and Multi-Target Multi-Camera Tracking

Organizers: Rita Cucchiara
Wen Gao
Shaogang Gong
Thomas S. Huang
Ergys Ristani
Francesco Solera
Qi Tian
Carlo Tomasi
Simone Calderara
Cees G. M. Snoek
Jingdong Wang
Shiliang Zhang

Schedule: Full Day

0845 **Welcome**

0900 **Invited Talk:** Towards Learning Universal Feature Representations for Person Search, *Xiaogang Wang (Chinese Univ. of Hong Kong)*

0930 **Invited Talk:** MOTChallenge: Unifying Detection and Multi-Target Tracking Benchmarks, *Laura Leal-Taixé, Anton Milan (Technical Univ.Munich, Univ.of Adelaide)*

1000 Track-Clustering Error Evaluation for Track-Based Multi-Camera Tracking System Employing Human Re-Identification, *Chih-Wei Wu, Meng-Ting Zhong, Yu Tsao, Shao-Wen Yang, Yen-Kuang Chen, Shao-Yi Chien*

1015 DukeMTMC4ReID: A Large-Scale Multi-Camera Person Re-Identification Dataset, *Mengran Gou, Srikrishna Karanam, Wenqian Liu, Octavia Camps, Richard J. Radke*

1030 Morning Break (Kamehameha II)

1100 **Invited Talk:** Learning to Segment Moving Objects, *Cordelia Schmid (INRIA)*

1130 **Invited Talk:** How Far Are We From Real-World Person Re-Identification, *Ziyan Wu (Siemens)*

1200 Person Re-Identification by Deep Learning Attribute-Complementary Information, *Arne Schumann, Rainer Stiefelhagen*

1215 Towards a Principled Integration of Multi-Camera Re-Identification and Tracking Through Optimal Bayes Filters, *Lucas Beyer, Stefan Breuers, Vitaly Kurin, Bastian Leibe*

1230 Lunch (Kamehameha II)

1400 **Invited Talk:** Asymmetrical Person Re-Identification, *Wei-Shi Zheng (Sun Yat-sen Univ.)*

1430 **Invited Talk:** Practices of Large-Scale Target Re-Identification, *Xian-Sheng Hua (Alibaba)*

1500 **Spotlights**

- Video-Based Person Re-Identification by Deep Feature Guided Pooling, *Youjiao Li, Li Zhuo, Jiafeng Li, Jing Zhang, Xi Liang, Qi Tian*
- A Dataset for Persistent Multi-Target Multi-Camera Tracking in RGB-D, *Ryan Layne, Sion Hannuna, Massimo Camplani, Jake Hall, Timothy M. Hospedales, Tao Xiang, Majid Mirmehdi, Dima Damen*
- Trajectory Ensemble: Multiple Persons Consensus Tracking Across Non-Overlapping Multiple Cameras Over Randomly Dropped Camera Networks, *Yasutomo Kawanishi, Daisuke Deguchi, Ichiro Ide, Hiroshi Murase*
- Deep Spatial-Temporal Fusion Network for Video-Based Person Re-Identification, *Lin Chen, Hua Yang, Ji Zhu, Qin Zhou, Shuang Wu, Zhiyong Gao*

1530 Afternoon Break (Kamehameha II) **& Poster Session**

1615 **Invited Talk:** Multi-Target Tracking in Non-Overlapping Cameras Using Constraint Dominant Sets, *Mubarak Shah (Univ.of Central Florida)*

1645 **Invited Talk:** Leveraging the Network for Person Re-Identification in Camera Networks, *Amit Roy-Chowdhury (Univ.of California Riverside)*

1715 **Panel Discussion**

1755 **Closing Remarks**

EarthVision: Large Scale Computer Vision for Remote Sensing Imagery

Organizers: Devis Tuia
Jan Dirk Wegner
Konrad Schindler
Pietro Perona
Josiane Zerubia
Gabriele Moser
Ryan Mukherjee
Myron Brown
HakJae Kim
Todd Bacastow
Tony Frazier
Giovanni Marchisio

Schedule: Full Day

0830 **Welcome**

0850 **Invited Talk:** Distributed Views Across Media: From Space to Ocean Depths, *Yoav Schechner (Technion)*

0930 **Posters Spotlights**

1000 Morning Break (Kamehameha II)

1030 **Best Paper Awards**

- Robocodes: Towards Generative Street Addresses From Satellite Imagery, *İlke Demir, Forest Hughes, Aman Raj, Kleovoulos Tsourides, Divyaa Ravichandran, Suryanarayana Murthy, Kaunil Dhruv, Sanyam Garg, Jatin Malhotra, Barrett Doo, Grace Kermani, Ramesh Raskar*
- Temporal Vegetation Modelling Using Long Short-Term Memory Networks for Crop Identification From Medium-Resolution Multi-Spectral Satellite Images, *Marc Rußwurm, Marco Körner*

1120 **Invited Talk:** Remote Sensing in an HD Map World, *Mark Tabb (Here Maps)*

1200 Lunch (Kamehameha II)

1330 **Invited Talk:** Benchmark Datasets and Open Competitions, *Todd Stavish, Patrick Hagerty (CosmiQ)*

1410 **Posters**

- Super-Resolution of Multispectral Multiresolution Images From a Single Sensor, *Charis Lanaras, José Bioucas-Dias, Emmanuel Baltsavias, Konrad Schindler*
- On the Role of Representations for Reasoning in Large-Scale Urban Scenes, *Randi Cabezas, Maroš Bláha, Sue Zheng, Guy Rosman, Konrad Schindler, John W. Fisher III*
- Monitoring Ethiopian Wheat Fungus With Satellite Imagery and Deep Feature Learning, *Reid Pryzant, Stefano Ermon, David Lobell*
- Filmy Cloud Removal on Satellite Imagery With Multispectral Conditional Generative Adversarial Nets, *Kenji Enomoto, Ken Sakurada, Weimin Wang, Hiroshi Fukui, Masashi Matsuoka, Ryosuke Nakamura, Nobuo Kawaguchi*
- Automatic 3D Reconstruction From Multi-Date Satellite Images, *Gabriele Facciolo, Carlo de Franchis, Enric Meinhardt-Llopis*
- Joint Learning From Earth Observation and OpenStreetMap Data to Get Faster Better Semantic Maps, *Nicolas Audebert, Bertrand Le Saux, Sébastien Lefèvre*
- Dense Semantic Labeling of Very-High-Resolution Aerial Imagery and LiDAR With Fully-Convolutional Neural Networks and Higher-Order CRFs, *Yansong Liu, Sankaranarayanan Piramanayagam, Sildomar T. Monteiro, Eli Saber*
- Nonrigid Registration of Hyperspectral and Color Images With Vastly Different Spatial and Spectral Resolutions for Spectral Unmixing and Pansharpening, *Yuan Zhou, Anand Rangarajan, Paul D. Gader*
- Earth Observation Using SAR and Social Media Images, *Yuanyuan Wang, Xiao Xiang Zhu*

1530 Afternoon Break (Kamehameha II)

1600 **Panel Discussion:** The Future of Geospatial Data Sets for Computer Vision and Machine Learning

1640 **Closing Remarks**

Visual Understanding of Humans in Crowds & Look Into Person Challenge

Organizers: Xiaodan Liang
Shenghua Gao
Xiaohui Shen
Wei-Shi Zheng
Wanli Ouyang

Schedule: Full Day

0830 **Opening Remarks & Welcome**

0840 **The Look Into Person (LIP) Challenge:** Introduction & Results

0900 **Oral Talk:** Winner of LIP Challenge

0915 **Invited Talk:** *Alan Yuille (Johns Hopkins Univ.)*

1000 Morning Break (Kamehameha II) **& Poster Session**

- Human Activity Recognition Using Combinatorial Deep Belief Networks, *Shreyank N. Gowda*
- Self-Supervised Neural Aggregation Networks for Human Parsing, *Jian Zhao, Jianshu Li, Xuecheng Nie, Fang Zhao, Yunpeng Chen, Zhecan Wang, Jiashi Feng, Shuicheng Yan*

1030 **Invited Talk:** *Trevor Darrell (Univ. of California, Berkeley)*

1115 **Invited Talk:** *Xiaogang Wang (Chinese Univ. of Hong Kong)*

1200 Lunch (Kamehameha II)

1300 **Invited Talk:** *Yaser Sheikh (Carnegie Mellon Univ.)*

1345 **Invited Talk:** *Shuicheng Yan (National Univ. of Singapore, Qihoo/360)*

1430 **Invited Talk:** *Abhinav Gupta (Carnegie Mellon Univ.)*

1515 Afternoon Break (Kamehameha II) **& Poster Session**

1545 **Invited Talk:** *Shaogang Sean Gong (Queen Mary Univ. of London)*

1630 **Oral Talk:** TBA

1645 **Oral Talk:** TBA

1700 **Awards & Future Plans**

Brave New Ideas for Motion and Spatio-Temporal Representations

Organizers: Stratis Gavves
Basura Fernando
Chenliang Xu
Yan Yan
Hakan Bilen
Xuming He
Michael Ying Yang
Jan van Gemert

Schedule: Full Day

0845 **Welcome**

0900 **Invited Talk:** *William T. Freeman (MIT)*

0945 Unsupervised Human Action Detection by Action Matching, *Basura Fernando, Sareh Shirazi, Stephen Gould*

1000 Morning Break (Kamehameha II)

1015 **Invited Talk:** *René Vidal (Johns Hopkins Univ.)*

1100 **Invited Talk:** *Alan Yuille (Johns Hopkins Univ.)*

1145 Lunch (Kamehameha II) **& Poster Session**

1330 **Invited Talk:** *Jiebo Luo (Univ. of Rochester)*

1415 RATM: Recurrent Attentive Tracking Model, *Samira Ebrahimi Kahou, Vincent Michalski, Roland Memisevic, Christopher Pal, Pascal Vincent*

1430 Afternoon Break (Kamehameha II)

1445 **Invited Talk:** *Cees Snoek (Univ. of Amsterdam)*

1530 Interpretable 3D Human Action Analysis With Temporal Convolutional Networks, *Tae Soo Kim, Austin Reiter*

1545 Learning Dynamic GMM for Attention Distribution on Single-Face Videos, *Yun Ren, Zulin Wang, Mai Xu, Haoyu Dong, Shengxi Li*

1600 Optical Acceleration for Motion Description in Videos, *Anitha Edison, Jiji C. V.*

1615 **Summary & Closing Remarks**

1630 **Poster Session**

Fine-Grained Visual Categorization

Organizers: Ryan Farrell
Subhransu Maji
Yang Song
Grant Van Horn
Oisin Mac Aodha
Yin Cui

Schedule: Full Day

0900 **Opening Remarks**

0915 **Invited Talk:** TBA, *Erik Rodner (Carl Zeiss AG)*

0945 **Poster Spotlights**

1015 Morning Break (Kamehameha II) **& Posters**

- Fine-Grained Semantic Part Segmentation Using Encoder-Multiple Decoders CNN, *Hiroaki Aizawa, Kunihito Kato, Takayoshi Yamashita*
- Reasoning About Fine-Grained Attribute Phrases Using Reference Games, *Jong-Chyi Su, Chenyun Wu, Huaizu Jiang, Subhransu Maji*
- Gaze Embeddings for Zero-Shot Image Classification, *Nour Karessli, Zeynep Akata, Bernt Schiele, Andreas Bulling*
- Low-Rank Bilinear Pooling for Fine-Grained Classification, *Shu Kong, Charless Fowlkes*
- BoxCars: 3D Boxes for Improved Fine-Grained Vehicle Recognition, *Jakub Sochor, Jakub Špaňhel, Adam Herout*
- Face Verification With Rank-1 Counts for Fewer False Positives, *SouYoung Jin, Erik Learned-Miller*
- Fine-Grained Image Classification via Combining Vision and Language, *Xiangteng He, Yuxin Peng*
- Weakly Supervised Learning of Part Selection Model With Spatial Constraints for Fine-Grained Image Classification, *Xiangteng He, Yuxin Peng*
- Deep Convolutional Networks With Non-convolutional Recurrent Memory for Fine-Grained Structured Prediction, *Joel Ruben Antony Moniz, Christopher Pal*
- Improved Bilinear Pooling With CNNs, *Tsung-Yu Lin, Subhransu Maji*
- How Ethical Considerations Can Improve the Design of Computer Vision Technology as Applied to Species Recognition, *Andrew Robinson*
- Large-Scale Automatic Species Identification, *Jeff Mo, Eibe Frank, Varvara Vetrova*
- The iMaterialist Challenge 2017 Dataset, *Yin Cui, Xiao Zhang, Yang Song, Yuan Li, Hartwig Adam, Serge Belongie*
- The iNaturalist Challenge 2017 Dataset, *Grant Van Horn, Oisin Mac Aodha, Yang Song, Alex Shepard, Hartwig Adam, Pietro Perona, Serge Belongie*

1100 **Invited Talk:** Using Deep Learning and Google Street View to Estimate the Demographic Makeup of the US, *Timnit Gebru (Stanford Univ.)*

1130 **Invited Talk**: Advancing Medicine With Intelligent Pathology, *Aditya Khosla (PathAI)*

1200 Lunch (Kamehameha II)

1400 **Invited Talk:** TBA, *Alex Shepard (iNaturalist.org)*

1430 **Invited Talk:** TBA, *Hartwig Adam (Google)*

1500 **Competitions Overview:** iNaturalist & iMatarialist

1530 Afternoon Break (Kamehameha II) **& Posters** (cont.)

1615 **Finalist Presentations**

1700 **Panel Discussion**

1745 **Awards & Concluding Remarks**

Visual Understanding for Interaction

Organizers: Joseph J. Lim
Phillip Isola
Abhinav Gupta

Schedule: Full Day

0850 **Welcome**

0900 **Invited Talk:** *Raquel Urtasun (Univ. of Toronto, Uber*)

0930 **Invited Talk:** *Ali Farhadi (Univ. of Washington)*

1000 Morning Break (Kamehameha II)

1045 **Invited Talk:** *Dieter Fox (Univ. of Washington)*

1115 **Invited Talk:** *Antonio Torralba (MIT)*

1200 Lunch (Kamehameha II)

1330 **Poster Session**

1430 **Invited Talk:** *Josef Sivic (INRIA)*

1500 **Invited Talk:** TBA

1530 Afternoon Break (Kamehameha II)

1615 **Invited Talk:** TBA

1645 **Panel Session**

Wednesday, July 26

NOTE: Tutorial locations are in the online proceedings and the mobile app.
Tutorial times with an asterisk (*) indicates a default time—no schedule was provided.

0800–1300 Registration (Main Lobby)

0745–0845 Breakfast (Kamehameha II)

1000–1045 Morning Break (Kamehameha II)

1200–1330 Lunch (Kamehameha II)

1530–1615 Afternoon Break (Kamehameha II)

Geometric and Semantic 3D Reconstruction

Organizers: Christian Häne
Sudeep Pillai
Srikumar Ramalingam
Sudipta Sinha

Time: 0900-1715 (Full Day)

Description: 3D scene reconstruction from images is a fundamental topic in computer vision which has witnessed rapid progress in the last two decades. Reconstruction techniques typically involve solving parameter estimation problems that are inherently ill-posed, and therefore require appropriate regularization to handle noise and ambiguities present in the input data. While traditional methods mostly relied on low-level geometric priors, nowadays, mid and high-level scene information in the form of structured and semantics scene priors are increasingly being used. In this tutorial, we will discuss techniques for single image reconstruction, dense stereo correspondence in images and video, multi-view stereo, volumetric 3D reconstruction, mesh-based reconstruction and depth-map fusion approaches. This includes semantic depth map fusion techniques which utilize semantic information to guide the regularization and eventually output a semantically segmented 3D model. We will cover both batch 3D reconstruction methods as well as outline and demo recent capabilities in Visual SLAM (Simultaneous Localization and Mapping) that enable real-time 3D reconstructions. Additionally, we will discuss recent successful applications of deep learning to geometric tasks such as dense stereo matching, 3D shape prediction and camera pose estimation.

Schedule:

0900 Feature-Based and Deep Learning Techniques for Single-View Problem, *Srikumar Ramalingam (Utah)*
- Depth Estimation
- Semantic Segmentation
- Semantic Boundary Labeling

1015 Morning Break (Kamehameha II)

1045 Geometric Visual-SLAM: Feature-Based & Direct Methods, *Sudeep Pillai (MIT)*
- Sparse, Dense, & Semi-Dense Methods
- Stereo & RGB-D vSLAM
- Semantic SLAM

1200 Lunch (Kamehameha II)

1330 Semi-Global Stereo Matching (SGM) and Variants, *Sudipta Sinha (Microsoft Research)*
- Discrete and Continuous Optimization in Stereo
- Deep Learning in Stereo
- Efficient Scene Flow Estimation From Stereoscopic Video

1445 Volumetric Reconstruction, Depth Map Fusion, *Christian Häne (UC Berkeley)*
- Semantic 3D Reconstruction
- 3D Object Shape Priors

1530 Afternoon Break (Kamehameha II)

1615 3D Prediction Using ConvNets, *Christian Häne (UC Berkeley)*

1645 Various Visual-SLAM demos (time permitting), *Sudeep Pillai (MIT)*

Large-Scale Visual Place Recognition and Image-Based Localization

Organizers: Torsten Sattler
Akihiko Torii
Alex Kendall
Giorgos Tolias

Time: 0845-1230 (Half Day — Morning)

Description: Place recognition and image-based localization are the problems of determining which place is visible in an image and of estimating the position and orientation from which an image was taken, respectively. They play important roles in many (geometric) Computer Vision applications such as Structure-from-Motion, Simultaneous Localization and Mapping, or Augmented Reality and solving them is vital in the context of autonomous vehicles such as self-driving cars or drones. This tutorial offers a detailed introduction to both problems, covering both classical feature-based approaches and recently proposed methods employing machine learning. The tutorial focuses on scalable and efficient solutions to the two problems. It also discusses related problems such as image retrieval and camera pose estimation.

Schedule:

0845 Introduction, *Alex Kendall (Univ. of Cambridge), Torsten Sattler (ETH Zurich), Giorgos Tolias (CTU in Prague), Akihiko Torii (Tokyo Inst. of Technology)*

0855 Feature-Based Place Recognition, *Akihiko Torii (Tokyo Inst. of Technology)*

0940 Learning-Based Place Recognition, *Giorgos Tolias (CTU in Prague)*

1025 Morning Break (Kamehameha II)

1045 Feature-Based Visual Localization, *Torsten Sattler (ETH Zurich)*

1130 Learning-Based Visual Localization, *Alex Kendall (Univ. of Cambridge)*

1215 **Panel Discussion:** The Future of Place Recognition and Visual Localization

Computer Vision on Microsoft HoloLens

Organizers: Marc Pollefeys, Georg Klein, Drew Steedly, Pawel Olszta, Sudipta Sinha

Time: 0900-1200 (Half Day — Morning)

Description: The HoloLens is the world's leading augmented reality headset, but it's also a potent computer vision research device. Application code can access audio and video streams and surface meshes, all in a world coordinate space maintained by HoloLens' highly accurate head-tracking. This half day tutorial will dive into the vision processing on the HoloLens and show how to build applications that access the sensor streams for processing on- and off-device. We will also be demonstrating the latest updates to the HoloLens software which offer increasing research access to the full power of the device.

Schedule:

0900 Introduction to HoloLens, *Marc Pollefeys, Drew Steedly, Georg Klein*

1015 Morning Break (Kamehameha II)

1045 Computer Vision on HoloLens, *George Petre, Pawel Olszta, Sudipta Sinha*

Notes:

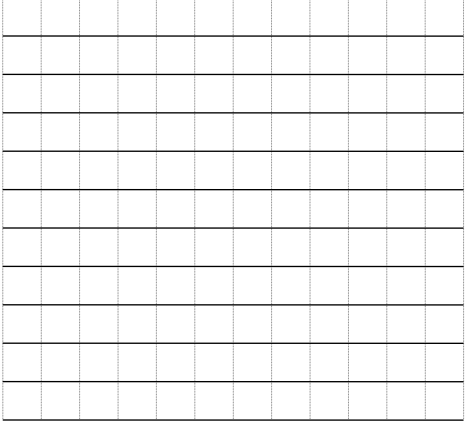

Towards Next Generation Deep Learning Framework: An Introduction to MXNet

Organizers: Naiyan Wang
Mu Li

Time: 0830-1200* (Half Day — Morning)

Description: Deep learning continues to push the state of the art in computer vision. One of the key reasons for this progress is the availability of highly flexible and developer-friendly deep learning frameworks. During this session, you'll learn how to use Apache MXNet to help speed your development and leave able to quickly spin up AWS GPU clusters to train at record speeds. Topics covered: (1) A walk-through on setting up MXNet on both your laptop and AWS, (2) A peek under the MXNet hood and a comparison with other deep learning frameworks (3) Hands on with Apache MXNet on computer vision applications.

Anomaly Detection: A Novel Framework

Organizer: Josef Kittler

Time: 0830-1200* (Half Day — Morning)

Description: Anomaly refers to events or situations, which deviate from normality (usual observation, order, form or rule). Anomaly detection has numerous applications in machine perception. The notion of anomaly has many nuances, which are impossible to discern using simply the classical mathematical concept of anomaly as an outlier of a probability distribution. This tutorial presents an anomaly detection architecture, which comprises several distinct mechanisms to detect anomalous events and facilitates their characterisation. In addition to the conventional process of outlier detection, the mechanisms include classifier incongruence detection, data quality assessment, classifier confidence gauging, and model-drift detection. The outputs of these processes feed into a reasoning engine, which draws conclusions about the presence of anomaly and its nature. The advocated approach to anomaly detection is illustrated on a number of applications.

Zero-Shot Learning for Computer Vision

Organizers: Thomas Mensink
Efstratios Gavves
Zeynep Akata
Cees G.M. Snoek

Time: 1330-1700 (Half Day — Afternoon)

Description: We live in the age of Big Data, featuring huge image and video datasets. Despite their size, however, we cannot guarantee sufficient annotations for all possible concepts. Moreover while annotations are easy to obtain for common object concepts, such as ball or helicopter, this is not straightforward for more exotic concepts like a "lagerphone" (a percussion musical instrument): not only the available images do not suffice, but often the annotations can be made only be experts. In the absence of annotations we promote zero-shot learning, where the combination of a) existing classifiers and b) semantic, cross-concept mappings between these classifiers allows for building novel classifiers without resorting to any visual examples. From a more philosophical point-of-view zero-shot learning relates to the ability to "learn new things" and to "reason over what is learned". While a DeepNet can reason (almost) perfectly over the 1,000 concepts it is trained on, it can not reason over any new concept, nor explain novel concepts in terms of what is already known. In this tutorial we focus on zero-shot learning for Computer Vision.

Schedule:

1330 Introduction, *Efstratios Gavves*
1340 Classification, *Zeynep Akata*
1430 Localization, *Efstratios Gavves*
1500 Retrieval, *Cees G.M. Snoek*

1530 Afternoon Break (Kamehameha II)

1600 Open Problems, *Zeynep Akata & Efstratios Gavves*
1640 Conclusion, *Efstratios Gavves*

Scalable Deep Learning With Microsoft Cognitive Toolkit

Organizers: Emad Barsoum
Sayan Pathak and Cha Zhang;

Time: 1330-1700* (Half Day — Afternoon)

Description: We will introduce Microsoft's Cognitive Toolkit, also known as CNTK, to the computer vision community. CNTK was originally developed by the speech team at Microsoft Research, and it was the key to Microsoft Research's recent breakthrough in reaching human parity in conversational speech recognition. In February 2016, CNTK was released on GitHub with MIT license, and today's 2.0 version contains both C++ and Python APIs, and runs natively on both Windows and Linux. While CNTK is little known to the vision community, it is highly regarded by other communities, achieving at least 5-10x speed up on recurrent neural networks compared with any other toolkits, on either CPU or GPU. In this tutorial, we will explain how this is done with symbolic loops, automatic batching and innovative data parallel training algorithms. Furthermore, we will use many computer vision related examples to demonstrate the flexibility, extensibility and scalability of CNTK.

Theory and Application of Generative Adversarial Network

Organizers: Ming-Yu Liu
Jan Kautz
Julie Bernauer

Time: 1330-1700* (Half Day — Afternoon)

Description: Generative adversarial network framework has recently emerged as a promising generative modeling approach. It consists of a generative network and a discriminative network. Through the competition between the two networks, it learns to model the data distribution. In addition to modeling the image/video distribution in computer vision problems, the framework finds use in defining visual concept using examples. To a large extent, it eliminates the need of hand-crafting objective functions for various computer vision problems. In this tutorial, we will present a comprehensive overview of generative adversarial network research. We will cover several recent theoretical studies as well as training techniques and will also cover multiple vision applications of generative adversarial networks.

3D Deep Learning

Organizers: Leonidas Guibas
Michael Bronstein
Evangelos Kalogerakis
Qixing Huang
Jimei Yang;Hao Su
Charles Qi

Time: 1330-1700* (Half Day — Afternoon)

Description: Understanding 3D data has been attracting increasing attention recently due to its importance for many vision systems, such as self-driving cars, autonomous robots, augmented reality and medical image processing. This tutorial covers deep learning algorithms for 3D geometric data. Different from 2D images that have a dominant representation as arrays, 3D geometric data have multiple popular representations, ranging from point cloud, meshes, volumetric field to multi-view images, each fitting their own application scenarios. Each type of data format has its own properties that pose challenges to deep architecture design while also provides the opportunity for novel and efficient solutions. In this course, we will introduce the major advance of deep learning for each of the 3D representation types. We systematically introduce topics such as the characteristics of each representation type, how to encode them as neural network input and output, and what are the keys in the design of corresponding network structures. Through the course, the audience will learn the big picture of cutting-edge techniques as well as open problems in the field. For schedule and course material, please check http://3ddl.stanford.edu.

Wednesday, July 26

NOTE: Workshop locations are in the online proceedings and the mobile app.

0800–1300 Registration (Main Lobby)

0745–0845 Breakfast (Kamehameha II)

1000–1045 Morning Break (Kamehameha II)

1200–1330 Lunch (Kamehameha II)

1530–1615 Afternoon Break (Kamehameha II)

PASCAL in Detail Challenge

Organizers: Sanja Fidler
Iasonas Kokkinos
Roozbeh Mottaghi
George Papandreou
Raquel Urtasun
Andrea Vedaldi
Alan L. Yuille

Schedule: Full Day

0900 **Introduction & Welcome**
0930 **Single-Track Challenges & Benchmarks**
1030 **Winning Entries Presentations**
1200 **Keynote Talk:** *Andrew Zisserman (Univ. of Oxford)*
1300 Lunch (Kamehameha II)
1400 **Keynote Talk:** *Larry Zitnick (Facebook AI Research)*
1400 **Multi-Track Challenges**
1430 **Winning Entries Presentations**
1530 Afternoon Break (Kamehameha II)
1600 **Keynote Talk:** *Kaiming He (Facebook AI Research)*
1700 **Panel Discussion & Closing Remarks**

Beyond ImageNet Large Scale Visual Recognition Challenge

Organizers: Olga Russakovsky
Eunbyung Park
Wei Liu
Jia Deng
Fei-Fei Li
Alex Berg

Schedule: Full Day

0900 **Welcome**
0905 ImageNet: Where Are We Going? And Where Have We Been? *Fei-Fei Li (Stanford Univ.)*
0935 **Invited Talk:** *Raquel Urtasun (Uber ATG, Univ. of Toronto)*
1010 **Poster Session:** Recent work from winners of ILSVRC 2010-2016

- Aggregated Residual Transformations for Deep Neural Networks, *Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, Kaiming He*
- Unsupervised Imitation Learning, *Pierre Sermanet, Corey Lynch, Kelvin Xu, Jasmine Hsu, Sergey Levine*
- RED-Net: A Recurrent Encoder-Decoder Network for Video-Based Face Alignment, *Xi Peng, Rogerio S. Feris, Xiaoyu Wang, Dimitris N. Metaxas*
- SEP-Net: Simple and Effective Pattern Networks, *Zhe Li, Xiaoyu Wang, Xutao Lv, Tianbao Yang*
- UberNet: Training a 'Universal' Convolutional Neural Network for Low-, Mid-, and High-Level Vision Using Diverse Datasets and Limited Memory, *Iasonas Kokkinos*
- Brain MRI Diagnostic Using Deep Convolutional Network, *Shao Jie, Zhang Jie, Wu Jinsong, Zhang Zheng*
- Modeling Context and Deformation in Object Detection, *Wanli Ouyang, Hongsheng Li, Junjie Yan, Xingyu Zeng, Kai Kang, Tong, Xiao, Kun Wang, Hongyang Li, Zhe Wang, Yucong, Zhou, Bin Yang, Xuanyi Dong, Ping Luo, Shi Qiu, Yonglong Tian, Shuo Yang, Yuanjun Xiong, Chen Qian, Zhenyao Zhu, Ruohui Wang, Yubin Deng, Buyu Li, Xin Zhu, Xihui Liu, Chen-Change Loy, Shengen Yan, Dahua Lin, Xiaogang Wang, Xiaoou Tang*

- SNPE: Snapdragon Neural Processing Engine Powering Deep Learning on Mobile Devices, *Koen van de Sande, Cees Snoek, Daniel Fontijne*
- ...

1130 **Overview:** ILSVRC 2017 -- "Taster" Amazon Bin Image Dataset Challenge and Low Power Image Recongition Challenge (LPRIC)

1145 **Challenge Spotlights:** Session 1

1215 Lunch (Kamehameha II)

1400 **Invited Talk:** *Jitendra Malik (UC Berkeley)*

1430 **Awards**

1445 **Challenge Spotlights:** Session 2

1515 **Invited Talk:** *Larry Zitnick (Facebook AI Research)*

1545 **Poster Session:** Winners of ILSVRC 2017 -- includes "Taster" Amazon Bin Image Dataset Challenge (ABIDC) and Low Power Image Recongition Challenge (LPRIC)

1700 **Closing Remarks**

Medical Computer Vision

Organizers: Le Lu
Tal Arbel
Leo Grady
Bjoern Menze
Georg Langs

Schedule: Full Day

0730 **Invited Talk:** Keeping Human in the Loop of Analyzing Big Data of Microscopy Images, *Zhaozheng Yin (Missouri Univ. S&T)*

0800 **Invited Talk:** Dermatologist-Level Classification of Skin Cancer With Deep Neural Networks, *Andre Esteva (Stanford Univ.)*

0830 **Invited Talk:** Domain Adaptation and Active Learning for Microscopy Imaging, *Pascal Fua (EPFL)*

0900 **Invited Talk:** Exploring and Harvesting the Healing Power of Big Data, *Jiebo Luo (Univ. of Rochester)*

0930 **Invited Talk:** Behavioral Imaging and the Study of Autism, *James Rehg (Georgia Tech)*

1000 Morning Break (Kamehameha II)

1015 **Invited Talk:** Medical Imaging at Google, *Martin Stumpe, Jonathan Krause (Google Research)*

1045 **Invited Talk:** Large Scale Biomedical Image and Genetic Data Analytics for Precision Medicine, *Wiro Niessen (Erasmus MC, TU Delft)*

1115 **Invited Talk:** Big Data, Weak Label and True Clinical Impacts for Radiology Imaging Diagnosis, *Xiaosong Wang (NIH)*

1145 **Invited Talk:** Deep Learning Based Survival Prediction From Big Image-Omics Data, *Junzhou huang (UT Arlington)*

1215 Lunch (Kamehameha II)

1330 **Invited Talk:** Learning to Interpret Ultrasound Video, *Alison Noble (Oxford Univ.)*

1400 **Invited Talk:** Quantitative Imaging Biomarkers: To Learn or Not to Learn? *Marleen de Bruijne (Erasmus MC Rotterdam)*

1430 **Invited Talk:** From Surgical Activities to Radiologic Reports: Computer Vision and Deep Learning for Medicine, *Allan Yuille, Greg Hager (Johns Hopkins Univ.)*

1500 **Invited Talk:** Learning Clinically Useful Information From Medical Images, *Daniel Rueckert (Imperial College London)*

1530 Afternoon Break (Kamehameha II)

1545 **Invited Talk:** Integrative Clinical Datanomics Through Computational Intelligence, *Nikos Paragios (Ecole Centrale Paris)*

1615 **Invited Talk:** Advances in Deep Neural Architectures for Image Segmentation, *Chris Pal (École Polytechnique de Montréal)*

1645 **Invited Talk:** Deep Learning and Its Application for Microscopic Image Analysis, *Lin Yang (UFL)*

1715 **Invited Talk:** Deep Learning in Medical Imaging: For Efficient and Enhanced-Value Radiology Reporting, *Hayit Greenspan (Tel-Aviv Univ.)*

ChaLearn: Explainable Computer Vision Workshop and Job Candidate Screening Competition

Organizers: Sergio Escalera
Hugo Jair Escalante
Xavier Baró
Isabelle Guyon
Meysam Madadi
Evelyne Viegas
Stephane Ayache
Julio Jacques
Umut Guclu
Yagmur Gucluturk
Marcel van Gerven
Rob van Lier

Schedule: Full Day

0845 **Presentation of Workshop and Challenge Summary**, *Sergio Escalera*

S1: Job Candidate Screening Coopetition (900-1000)

0900 Multi-Modal Score Fusion and Decision Trees for Explainable Automatic Job Candidate Screening From Video CVs, *Heysem Kaya, Furkan Gürpınar, Albert Ali Salah*

0920 Personality Traits and Job Candidate Screening via Analyzing Facial Videos, *Salah Eddine Bekhouche, Fadi Dornaika, Abdelkrim Ouafi, Abdelmalik Taleb-Ahmed*

0940 Human-Explainable Features for Job Candidate Screening Prediction, *Achmadnoer Sukma Wicaksana, Cynthia C. S. Liem*

1000 Morning Break (Kamehameha II)

1030 **Invited Talk:** Deep Learning for Perception, Action, and Explanation, *Trevor Darrel*

1115 **Invited Talk:** TBA, *Antonio Torralba*

1200 Lunch (Kamehameha II)

S2: Explainable Computer Vision I (1400-1520)

1400 Explaining Distributed Neural Activations via Unsupervised Learning, *Soheil Kolouri, Charles E. Martin, Heiko Hoffmann*

1420 Automated Screening of Job Candidate Based on Multimodal Video Processing, *Jelena Gorbova, Iiris Lüsi, Andre Litvin, Gholamreza Anbarjafari*

1440 Explaining the Unexplained: A CLass-Enhanced Attentive Response (CLEAR) Approach to Understanding Deep Neural Networks, *Devinder Kumar, Alexander Wong, Graham W. Taylor*

1500 Grad-CAM: Visual Explanations From Deep Networks via Gradient-Based Localization, *Ramprasaath Ramasamy Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, Dhruv Batra*

1520 Afternoon Break (Kamehameha II)

1600 **Invited Talk:** Structured Models for Human Action Recognition, *Cordelia Schmid*

S3: Explainable Computer Vision I (1645-1745)

1645 It Takes Two to Tango: Towards Theory of AI's Mind. , *Arjun Chandrasekharan, Deshraj Yadav, Prithvijit Chattopadhyay, Viraj Prabhu, Devi Parikh*

1705 Decoding the Deep: Exploring Class Hierarchies of Deep Representations Using Multiresolution Matrix Factorization, *Vamsi K. Ithapu*

1725 Interpreting CNN Models for Apparent Personality Trait Regression, *Carles Ventura, David Masip, Agata Lapedriza*

1745 **Closing Remarks**

Notes:

Light Fields for Computer Vision

Organizers: Bastian Goldlücke
Ole Johannsen
Katrin Honauer
Jingyi Yu

Schedule: Full Day

0850 **Welcome**

S1: Keynotes, Oral, & Posters (900-1215)

0900 **Keynote Talk:** Light Fields - From Shape Recovery to Sparse Reconstruction, *Ravi Ramamoorthi (UC San Diego)*

0945 **Keynote Talk:** Light Field for Cinema, *Thomas Nonn (Lytro)*

1030 Linearizing the Plenoptic Space, *Grégoire Nieto, Frédéric Devernay, James Crowley*

1050 **Poster Session**

- Full BRDF Reconstruction Using CNNs From Partial Photometric Stereo-Light Field Data, *Doris Antensteiner, Svorad Štolc*
- Surface Normal Reconstruction From Specular Information in Light Field Data, *Marcel Gutsche, Hendrik Schilling, Maximilian Diebold, Christoph Garbe*
- Dataset and Pipeline for Multi-View Light-Field Video, *Neus Sabater, Guillaume Boisson, Benoit Vandame, Paul Kerbiriou, Frederic Babon, Matthieu Hog, Remy Gendrot, Tristan Langlois, Olivier Bureller, Arno Schubert, Valerie Allié*
- Light Field Convergency: Implicit Photometric Consistency on Transparent Surface, *Yuta Ideguchi, Yuki Uranishi, Shunsuke Yoshimoto, Yoshihiro Kuroda, Osamu Oshiro*
- Optimizing the Lens Selection Process for Multi-Focus Plenoptic Cameras and Numerical Evaluation, *Luca Palmieri, Reinhard Koch*
- Depth from a Light Field Image with Learning-based Matching Costs, *Hae-Gon Jeon, Jaesik Park, Gyeongmin Choe, Jinsun Park, Yunsu Bok, Yu-Wing Tai, In So Kweon*
- Spinning Parallelogram Operator for Light Field Depth Estimation, *Shuo Zhang, Hao Sheng, Zhang Xiong*
- RM3DE: Multi-Resolution Depth Field Estimation, *Federica Battisti, Michele Brizzi, Marco Carli, Alessandro Neri*
- Occlusion-Model Guided Anti-Occlusion Depth Estimation in Light Field, *Hao Zhu, Qing Wang, Jingyi Yu*
- Dense Depth-map Estimation and Geometry Inference from Light Fields via Global Optimization, *Lipeng Si, Qing Wang*
- Zero Crossing Depth Reconstruction with Second Order Total Variation, *Maximilian Diebold, Hendrik Schilling, Marcel Gutsche*
- OBER - Occluded Bilateral EPI Regularization for Light Field Depth Estimation, *Hendrik Schilling, Maximilian Diebold*

1215 Lunch (Kamehameha II)

S2: Keynote & Oral Presentations (1330-1455)

1330 **Keynote Talk:** EPI Imaging on Contours for Range with Structure, *Harlyn Baker (Consulting Scientist at Stanford Univ.)*

1415 Underwater Image Dehazing with a Light Field Camera, *Katherine Skinner*

1435 Richardson-Lucy Deblurring for Moving Light Field Cameras, *Donald Dansereau*

1455 Afternoon Break

S3: Benchmark, Challenge, & Panel Discussion (1515-1745)

1515 Keynote Talk: Common Visual Data Foundation: Enabling Community-driven Research in Computer Vision, *Tsung-Yi Lin (Cornell Tech)*

1540 Benchmark: Challenge Evaluation and State-of-the-Art, *Katrin Honauer, Ole Johannsen*

1600 **Outstanding Challenge Participant 1:** *TBA*

1610 **Outstanding Challenge Participant 2:** *TBA*

1620 **Outstanding Challenge Participant 3:** *TBA*

1630 **Outstanding Challenge Participant 4:** *TBA*

1640 Short Break (to set up discussion)

1645 **Panel Discussion:** *Tsung-Yi Lin, Thomas Nonn, …*

1730 **Closing Remarks**

Media Forensics

Organizers: Kevin W. Bowyer
Shih-Fu Chang
Larry S. Davis

Schedule: Full Day

0850 **Welcome**

0900 **Keynote Talk:** Photo Forensics From JPEG Coding Artifacts, *Hany Farid (Dartmouth College)*

S1: Face Image Forensics (950-1010)

0950 Position Determines Perspective: Investigating Perspective Distortion for Image Forensics of Faces, *Bo Peng; Wei Wang; Jing Dong; Tieniu Tan*

1010 Morning Break (Kamehameha II)

S2: Neural Networks & Camera Identification (1040-1200)

1040 Transferable Deep-CNN Features for Detecting Digital and Print-Scanned Morphed Face Images, *R. Raghavendra; Kiran B. Raja; Sushma Venkatesh; Christoph Busch*

1100 Two-Stream Neural Networks for Tampered Face Detection, *Peng Zhou; Xintong Han; Vlad I. Morariu; Larry S. Davis*

1120 A Counter-Forensic Method for CNN-Based Camera Model Identification, *David Güera; Yu Wang; Luca Bondi; Paolo Bestagini; Stefano Tubaro; Edward J. Delp*

1140 Camera Source Identification Using Discrete Cosine Transform Residue Features and Ensemble Classifier, *Aniket Roy; Rajat Subhra Chakraborty; Udaya Sameer; Ruchira Naskar*

1200 Lunch (Kamehameha II)

1330 **Invited Talk:** The 2017 Nimble Challenge Evaluation: Results and Future Directions, *Jonathan Fiscus (NIST)*

S3: Tampering Detection (1400-1520)

1400 Tampering Detection and Localization Through Clustering of Camera-Based CNN Features, *Luca Bondi; Silvia Lameri; David Güera; Paolo Bestagini; Edward J. Delp; Stefano Tubaro*

1420 Localization of JPEG Double Compression Through Multi-Domain Convolutional Neural Networks, *Irene Amerini; Tiberio Uricchio; Lamberto Ballan; Roberto Caldelli*

1440 Detection of Metadata Tampering Through Discrepancy Between Image Content and Metadata Using Multi-Task Deep Learning, *Bor-Chun Chen; Pallabi Ghosh; Vlad I. Morariu; Larry S. Davis*

1500 Detection and Localization of Image Forgeries Using Resampling Features and Deep Learning, *Jason Bunk; Jawadul H. Bappy; Tajuddin Manhar Mohammed; Lakshmanan Nataraj; Arjuna Flenner; B.S. Manjunath; Shivkumar Chandrasekaran; Amit K. Roy-Chowdhury; Lawrence Peterson*

1520 Afternoon Break (Kamehameha II)

1535 **Invited Talk:** The Grand Challenges of Media Forensics, *David Doermann (DARPA)*

S4: Camera Identification (1615-1715)

1615 FORMS-Locks: A Dataset for the Evaluation of Similarity Measures for Forensic Toolmark Images, *Manuel Keglevic; Robert Sablatnig*

1635 A C3D-Based Convolutional Neural Network for Frame Dropping Detection in a Single Video Shot, *Chengjiang Long; Eric Smith; Arslan Basharat; Anthony Hoogs*

1655 Spotting Audio-Visual Inconsistencies (SAVI) in Manipulated Video, *Robert Bolles; J. Brian Burns; Martin Graciarena; Andreas Kathol; Aaron Lawson; Mitchell McLaren; Thomas Mensink*

1715 **Closing Remarks**

Notes:

Tensor Methods in Computer Vision

Organizers: Piotr Koniusz
Anoop Cherian
Fatih Porikli

Schedule: Full Day

0900 **Welcome:** *Fatih Porikli (Australian National Univ.)*

0905 **Invited Talk:** Role of Tensors in Deep Learning, *Animashree Anandkumar (Univ. of California Irvine, Amazon AI)*

0935 **Invited Talk:** Tensor Networks for Deep Learning, *Andrzej Cichocki (Brain Science Inst. RIKEN, Skoltech)*

1005 Morning Break (Kamehameha II)

1030 **Invited Talk:** Analysis and Design of Convolutional Networks via Hierarchical Tensor Decompositions, *Nadav Cohen (Hebrew Univ. of Jerusalem)*

1100 **Invited Talk:** Globally Optimal Structured Low-Rank Matrix and Tensor Factorization, *René Vidal (Johns Hopkins Univ.)*

1130 **Invited Talk:** Deep Learning and Tensors for the Approximation of Multivariate Functions: Recent Results and Open Problems, *Ivan Oseledets (Skoltech)*

1200 Lunch (Kamehameha II)

1330 Exploration of Social and Web Image Search Results Using Tensor Decomposition, *Liuqing Yang, Evangelos E. Papalexakis*

1340 Graph-Regularized Generalized Low-Rank Models, *Mihir Paradkar, Madeleine Udell*

1350 Exploring the Granularity of Sparsity in Convolutional Neural Networks, *Huizi Mao, Song Han, Jeff Pool, Wenshuo Li, Xingyu Liu, Yu Wang, William J. Dally*

1400 Human Action Recognition Using Tensor Dynamical System Modeling, *Chan-Su Lee*

1410 Tensor Contraction Layers for Parsimonious Deep Nets, *Jean Kossaifi, Aran Khanna, Zachary Lipton, Tommaso Furlanello, Anima Anandkumar*

1420 **Invited Talk:** Learning Methods and Optimization on Matrix Manifolds and Matrix Lie Groups, *Richard Hartley (Australian National Univ., Data61/CSIRO)*

1450 **Invited Talk:** You've got Data, We've Got Tensors: Linear and Multilinear Tensor Models for Computer Vision, Graphics and Machine Learning, *M. Alex O. Vasilescu (UCLA)*

1520 Afternoon Break (Kamehameha II)

1550 **Invited Talk:** Numerical Optimization Algorithm for Tensor-based Recognition, *Otto Debals, Lieven De Lathauwer (KU Leuven)*

1620 **Invited Talk:** A New Tensor Algebra - Theory and Applications, *Lior Horesh (IBM T. J. Watson Research Center, Columbia Univ.)*

1650 **Invited Talk:** Redeeming the Clinical Promise of Diffusion MRI in Support of the Neurosurgical Workflow, *Luc Florack (Eindhoven Univ. of Technology)*

1720 **Closing Remarks**

Continuous and Open-Set Learning

Organizers: Erik Rodner
Alexander Freytag
Christoph H. Lampert
Terrance E. Boult
Joachim Denzler

Schedule: Full Day

0830 **Welcome & Introduction**

0840 **Invited Talk:** Deep Reinforcement Learning in Sequential Environments, *Raia Hadsell (Google DeepMind)*"

0920 **Invited Talk**: Adaptive Representation Learning for Perception, Action, and Explanation, *Trevor Darrell (Univ. of California, Berkeley)*

1000 Morning Break (Kamehameha II) **& Poster Session**

1030 **Research Teaser Talks – Session 1:** Multiple Speakers

1110 **Invited Talk:** TBA, *Vittorio Ferrari (Google Zurich, Univ. of Edinburgh)*

1200 Lunch (Kamehameha II)

1400 **Research Teaser Talks – Session 2:** Multiple Speakers

1440 **Invited Talk**: TBA

1530 Afternoon Break (Kamehameha II) **& Poster Session**

1615 **Panel Discussion:** Speakers & Organizers

1700 **Closing Remarks**

YouTube-8M Large-Scale Video Understanding Challenge

Organizers: Apostol (Paul) Natsev
Rahul Sukthankar
Joonseok Lee
George Toderici

Schedule: Full Day

S1: Invited Talks I (0900-1030)

0900 **Opening Remarks**
0905 **Overview:** YouTube-8M Dataset & Challenge, *Paul Natsev (Google Research)*
0930 **Invited Talk:** Video Understanding: What We Understood and What We Still Need to Learn, *Alex Hauptmann (CMU)*
1000 **Invited Talk:** Structured Models for Human Action Recognition, *Cordelia Schmid (INRIA)*

1030 Morning Break (Kamehameha II)

S2: Oral Session: Classification Challenge (1045-1200)

1045 **Presentation**
1100 **Presentation**
1115 **Presentation**
1130 **Presentation**
1145 **Presentation**

1200 Lunch (Kamehameha II)

S3: Invited Talks II (1300-1430)

1300 **Invited Talk:** Learning from Synthetic Humans, *Ivan Laptev (INRIA Research)*
1330 **Invited Talk:** Tube Convolutional Neural Network (T-CNN) for Action Detection in Videos, *Mubarak Shah (Univ. of Central Florida)*
1400 **Approaches & Analysis:** YouTube-8M Classification Challenge, *Challenge Organizers*
1430 **Poster Session**

1530 Afternoon Break (Kamehameha II)

S4: Oral Session: General Research (1545-1645)

1545 **Presentation**
1600 **Presentation**
1615 **Presentation**
1630 **Presentation**
1645 **Closing and Award Ceremony**

Visual Question Answering Challenge

Organizers: Aishwarya Agrawal
Yash Goyal
Tejas Khot
Peng Zhang
Jiasen Lu
Larry Zitnick
Dhruv Batra
Devi Parikh

Schedule: Full Day

0900 **Welcome**
0910 **Invited Talk:** *Svetlana Lazebnik (UIUC)*
0935 **Invited Talk:** *Marcus Rohrbach (Facebook AI Research)*
1000 **Invited Talk:** *Kate Saenko (Boston Univ.)*

1025 Morning Break (Kamehameha II)

1045 **Invited Talk:** *Derek Hoiem (UIUC)*
1110 **Overview of Dataset, Challenge, Winner Announcements, Analysis of Results**
1140 **Challenge Honorable Mention Talk**
1145 **Challenge Runner-Up Talk**
1200 **Challenge Winner Talk**

1215 Lunch (Kamehameha II)

1345 **Invited Talk:** *Anton Van Den Hengel (Univ. of Adelaide)*
1410 **Invited Talk:** *Sanja Fidler (Univ. of Toronto)*

1435 Afternoon Break (Kamehameha II) **& Poster Session**

1600 **Invited Talk:** *Jason Weston (Facebook AI Research)*
1625 **Invited Talk:** *Abhishek Das (Georgia Tech)*
1650 **Invited Talk:** *Hugo Larochelle (Google Brain)*
1715 **Panel:** Future Directions
1755 **Closing Remarks**

Faces "In-The-Wild" Workshop-Challenge

Organizers: Stefanos Zafeiriou
Mihalis Nicolaou
Irene Kotsia
Fabian Benitez-Quiroz
Guoying Zhao
Maja Pantic

Schedule: Full Day

0800 **Welcome**

0815 **Invited Talk:** The Emotionet Challenge, *Aleix Martinez (Ohio State Univ.)*

0 900 Estimation of Affective Level in the Wild With Multiple Memory Networks, *Jianshu Li, Yunpeng Chen, Shengtao Xiao, Jian Zhao, Sujoy Roy, Jiashi Feng, Shuicheng Yan, Terence Sim*

0 920 Facial Affect Estimation in the Wild Using Deep Residual and Convolutional Networks, *Behzad Hasani, Mohammad H. Mahoor*

0 940 FATAUVA-Net : An Integrated Deep Learning Framework for Facial Attribute Recognition, Action Unit Detection, and Valence-Arousal Estimation, *Wei-Yi Chang, Shih-Huan Hsu, Jen-Hsien Chien*

1000 Morning Break (Kamehameha II)

1030 Recognition of Affect in the Wild Using Deep Neural Networks, *Dimitrios Kollias, Mihalis A. Nicolaou, Irene Kotsia, Guoying Zhao, Stefanos Zafeiriou*

1050 Aff-Wild: Valence and Arousal 'In-The-Wild' Challenge, *Stefanos Zafeiriou, Dimitrios Kollias, Mihalis A. Nicolaou, Athanasios Papaioannou, Guoying Zhao, Irene Kotsia*

1110 Deep Analysis of Facial Behavioral Dynamics, *Lazaros Zafeiriou, Stefanos Zafeiriou, Maja Pantic*

1130 **Poster Session 1** (Oral papers can also have posters)

- AgeDB: The First Manually Collected, In-The-Wild Age Database, *Stylianos Moschoglou, Athanasios Papaioannou, Christos Sagonas, Jiankang Deng, Irene Kotsia, Stefanos Zafeiriou*
- Marginal Loss for Deep Face Recognition, *Jiankang Deng, Yuxiang Zhou, Stefanos Zafeiriou*
- Deep Face Deblurring, *Grigorios G. Chrysos, Stefanos Zafeiriou*

1230 Lunch (Kamehameha II)

1430 Stacked Hourglass Network for Robust Facial Landmark Localisation, *Jing Yang, Qingshan Liu, Kaihua Zhang*

1450 Deep Alignment Network: A Convolutional Neural Network for Robust Face Alignment, *Marek Kowalski, Jacek Naruniec, Tomasz Trzcinski*

1510 Robust FEC-CNN: A High Accuracy Facial Landmark Detection System, *Zhenliang He, Jie Zhang, Meina Kan, Shiguang Shan, Xilin Chen*

1530 Afternoon Break (Kamehameha II) **& Poster Session 2** (Oral papers can also have posters)

- Convolutional Experts Constrained Local Model for Facial Landmark Detection, *Amir Zadeh, Tadas Baltrušaitis, Louis-Philippe Morency*
- 3D-Assisted Coarse-To-Fine Extreme-Pose Facial Landmark Detection, *Shengtao Xiao, Jianshu Li, Yunpeng Chen, Zhecan Wang, Jiashi Feng, Shuicheng Yan, Ashraf Kassim*
- Unconstrained Face Alignment Without Face Detection, *Xiaohu Shao, Junliang Xing, Jiangjing Lv, Chunlin Xiao, Pengcheng Liu, Youji Feng, Cheng Cheng*
- Multi-Scale Fully Convolutional Network for Face Detection in the Wild, *Yancheng Bai, Bernard Ghanem*
- Delving Deep Into Coarse-To-Fine Framework for Facial Landmark Localization, *Xi Chen, Erjin Zhou, Yuchen Mo, Jiancheng Liu, Zhimin Cao*

1630 Leveraging Intra and Inter-Dataset Variations for Robust Face Alignment, *Wenyan Wu, Shuo Yang*

1650 Face Detection, Bounding Box Aggregation and Pose Estimation for Robust Facial Landmark Localisation in the Wild, *Zhen-Hua Feng, Josef Kittler, Muhammad Awais, Patrik Huber, Xiao-Jun Wu*

1710 The Menpo Facial Landmark Localisation Challenge: A Step Towards the Solution, *Stefanos Zafeiriou, George Trigeorgis, Grigorios Chrysos, Jiankang Deng, Jie Shen*

1730 **Winners & Cosing Remarks**

Scene Understanding and LSUN Challenge

Organizers: Bolei Zhou
Aditya Khosla
Jianxiong Xiao
James Hays
Fisher Yu
Peter Kontschieder
Shuran Song
Ming Jiang
Yinda Zhang
Catherine Qi Zhao
Thomas Funkhouser
Jianxiong Xiao

Schedule: Full Day

S1: Scene Understanding Workshop (SUNw) (0825-1200)

0825 **Welcome**

0830 **Invited Talk:** *Matthias Nießner (TU Munch)*

0900 **Invited Talk:** *Larry Zitnick (Facebook AI Research)*

0930 **Spotlights**

- Understanding Convolution for Semantic Segmentation, *Panqu Wang, Pengfei Chen, Ye Yuan, Ding Liu, Zehua Huang, Xiaodi Hou, Garrison Cottrell*
- C-VQA: A Compositional Split of the Visual Question Answering (VQA) v1.0 Dataset, *Aishwarya Agrawal*
- High-Level Cues for Predicting Motivations, *Arun Mallya, Svetlana Lazebnik*
- SalGAN: Visual Saliency Prediction With Adversarial Networks, *Junting Pan, Xavier Giro-i-Nieto*
- Quantifying the Interpretability of Deep Visual Representations, *David Bau, Bolei Zhou, Antonio Torralba*

0945 Morning Break (Kamehameha II) **& Poster Session**

- CNN Based Repeated Cropping for Photo Re-Composition, *Eunbin Hong, Junho Jeon, Seungyong Lee*
- Object-Level Context Modeling for Scene Classification With Context-CNN, *Syed Javed, Anil Nelakanti*
- Detecting Nonexistent Pedestrians, *Jui Ting Chien, Chia-Jung Chou, Hwann-Tzong Chen*
- Self-Supervised Depth Learning Improves Semantic Segmentation, *Huaizu Jiang, Erik Learned-Miller, Gustav Larsson, Michael Maire, Greg Shakhnarovich*
- Depth Estimation From Monocular Image With Sparse Known Labels, *Yaoyu Li, Yfan Zhang*
- Good Practice on Deep Scene Classification: From Local Supervision to Knowledge Guided Disambiguation, *Yu Qiao, Limin Wang, Sheng Guo, Zhe Wang, Weilin Huang, Yali Wang*
- Uber-Text: A Large-Scale Dataset for Optical Character Recognition From Street-Level Imagery, *Ying Zhang, Ben Kadlec, Peter Zhang, Lionel Gueguen, Keith Seifert*
- Object Recognition and Reconstruction With Partial Appearance, *Lifeng Liu*
- Uniform Representations of Scene Objects for Autonomous Decision Making, *Xiaotian Yin, Lifeng Liu, Yingxuan Zhu, Jun Zhang, Jian Li*
- Weakly Supervised PatchNets: Learning Aggregated Patch Descriptors for Scene Recognition, *Zhe Wang, Limin Wang, Yali Wang, Bowen Zhang, Yu Qiao, Charless Fowlkes*
- Lessons Learned From a Lightweight Autonomous Vehicle Implementation, *Kuan Wang, Dawei Sun, Xiaojian Ma, Zhaoyuan Gu, Hao Zhao*
- What Is Wrong With Them?: Qualitative Analysis on ImageNet Failure Cases, *Han Lee, Heechul Jung, Junmo Kim*
- Object State Recognition for Automatic AR-Based Maintenance Guidance, *Pavel Dvorak, Elisabetta Delponte, Radovan Josth*

1100 **Invited Talk:** *Raquel Urtasun (Univ. of Toronto)*

1130 **Invited Talk:** *Ali Farhadi (Univ. of Washington)*

1200 Lunch (Kamehameha II)

S2: Large SUN Challenge (LSUN) (1345-1730)

1345 **Welcome & Overview**

1400 **Challenge Winner Talk**

1415 **Challenge Winner Talk**

1430 **Scene Segmentation Task:** *Peter Kontschieder*

1445 **Challenge Winner Talk**

1500 **Challenge Winner Talk**

1515 **Challenge Winner Talk**

1530 Afternoon Break (Kamehameha II)

1600 **Invited Talk**
1630 **Invited Talk**
1700 **Challenge Winner Talk**
1715 **Concluding Remarks**

BMTT-PETS Workshop on Tracking and Surveillance

Organizers: Laura Leal-Taixe
Luis Patino
Anton Milan
Tom Cane
Ian Reid
Daniel Cremers
James L. Crowley
Stefan Roth
Konrad Schindler
James Ferryman

Schedule: Full Day

0845 **Welcome and Introduction to Datasets & Dhallenges**
0915 **Invited Talk:** CNN-Based Visual Tracking, *Bohyung Han (Pohang Univ. of Science and Technology)*

S1: Detection & Tracking (0945-1025)

0945 CoMaL Tracking: Tracking Points at the Object Boundaries, *Santhosh K. Ramakrishnan, Swarna Kamlam Ravindran, Anurag Mittal*
1005 Enhancing Detection Model for Multiple Hypothesis Tracking, *Jiahui Chen, Hao Sheng, Yang Zhang, Zhang Xiong*

1025 Morning Break (Kamehameha II)

1045 **Invited Talk:** Human Activity Recognition From Anonymized Videos, *Michael Ryoo (EgoVid, Indiana Univ.)*

S2: Detection, Tracking, & Surveillance I (1115-1155)

1115 Okutama-Action: An Aerial View Video Dataset for Concurrent Human Action Detection, *Mohammadamin Barekatain, Miquel Martí, Hsueh-Fu Shih, Samuel Murray, Kotaro Nakayama, Yutaka Matsuo, Helmut Prendinger*
1135 Abnormal Event Detection on BMTT-PETS 2017 Surveillance Challenge, *Kothapalli Vignesh, Gaurav Yadav, Amit Sethi*

1155 Lunch (Kamehameha II) **& Poster Session**

S3: Detection, Tracking, & Surveillance II (1415-1535)

1415 **Invited Talk:** Performance Measures and the DukeMTMC Benchmark for Multi-Target Multi-Camera Tracking, *Ergys Ristani (Duke Univ.)*
1445 **Invited Talk:** Detection and Tracking Systems – What Counts in the Real World: A Maritime Surveillance Case Study, *Tom Cane (BMT)*
1515 Loitering Behaviour Detection of Boats at Sea, *Luis Patino, James Ferryman*

1535 Afternoon Break (Kamehameha II)

1600 **Invited Talk:** Tracking: Where Has It Been and Where Is It Going? *Robert Collins (Pennsylvania State Univ.)*

S4: Challenge Results & Discussion (1630-1730)

1630 **Evaluation Results & Awards**
1700 **Discussion**
1730 **Closing Remarks**

Visual Understanding by Learning from Web Data

Organizers: Jesse Berent
Abhinav Gupta
Rahul Sukthankar
Luc Van Gool
Wen Li
Limin Wang
Wei Li
Eirikur Agustsson

Schedule: Full Day

Schedue TBA

Joint Bridges to 3D Vision and Non-Rigid Structure From Motion Challenge

Organizers: David Fouhey
Joseph Lim
Qixing Huang
Alessio Del Bue
Henrik Aanæs
Sebastian Nesgaard Jensen
Yaser Sheikh

Schedule: Full Day

S1: NRSfM Challenge (0900-1300)

0900 **Opening Remarks**

0905 **Invited Talk:** Spectral Non-Rigid Partial Correspondence, *Michael M. Bronstein (Università della Svizzera Italiana)*

0950 **Challenge Presentation**

1010 Morning Break (Kamehameha II)

1045 **Invited Talk:** Monocular Reconstruction of Poorly-Textured Deformable Surfaces and Articulated Motion, *Mathieu Salzmann (EPFL)*

1130 **Challenge Results**

1200 **Invited Talk:** Self-Supervised Visual Descriptor Learning for Non-Rigid Dense Correspondence, *Richard Newcombe (Surreal Vision)*

1245 **Discussion & Closing Remarks**

1300 Lunch (Kamehameha II)

S2: Bridges to 3D Vision (1345-1800)

(See bridgesto3d.github.io for final schedule)

1345 **Welcome and Overview**

1400 **Invited Talk:** *Vladlen Koltun (Intel Visual Computing Lab)*

1430 **Invited Talk:** *Derek Hoiem (Univ. of Illinois at Urbana-Champaign)*

1500 **Invited Talk:** *Dinesh Manocha (Univ. of North Carolina)*

1530 Afternoon Break (Kamehameha II) **& Invited Poster Session**

1630 **Invited Talk:** *Raquel Urtasun (Univ. of Toronto)*

1700 **Invited Talk:** *Jitendra Malik (Univ. of California, Berkeley)*

1730 **Panel Discussion**

Deep-Vision: Deep Learning in Computer Vision – Temporal Deep Learning

Organizers: Jose M. Alvarez
Nathan Silberman
Dhruv Batra
Yann LeCun
Kamal Nasrollahi
Sergio Escalera
Ajmal Mian
Gholamreza Anbarjafari
Thomas B. Moeslund

Schedule: Full Day

S1: DeepVision (0845-1235)

0845 **Opening Remarks**

0850 **Invited Talk:** Unsupervised Cross-Domain Mapping, *Lior Wolf (Facebook AI Research)*

0925 **Invited Talk:** Optimization as a Model for Few-Shot Learning, *Hugo Larochelle (Google Brain)*

1000 Morning Break (Kamehameha II)

1035 **Invited Talk:** How Studying Brains Can Help Us Get a Deeper Vision, *Gabriel Kreiman (Harvard Univ.)*

1110 **Invited Talk:** TBA, *Sanja Fidler (Univ. of Toronto)*

1145 **Invited Talk:** Deep RL for Navigation in Complex Environments, *Raia Hadsell (Google DeepMind)*

1235 Lunch (Kamehameha II)

S2: Temporal Deep Learning (1435-1700)

1430 **Session Opening**

1435 **Invited Talk:** DeepVideo: Deep Models for Activity Detection and Description, *Kate Saenko (Boston Univ.)*

1505 **Poster Spotlights**

1535 Afternoon Break (Kamehameha II) **& Poster Session**

- Concurrence-Aware Long Short-Term Sub-Memories for Person-Person Action Recognition, *Xiangbo Shu,*

Jinhui Tang, Guo-Jun Qi, Yan Song, Zechao Li, Liyan Zhang

- Crowd-11: A Dataset for Fine Grained Crowd Behaviour Analysis, *Camille Dupont, Luis Tobías, Bertrand Luvison*
- Temporal Domain Neural Encoder for Video Representation Learning, *Hao Hu, Zhaowen Wang, Joon-Young Lee, Zhe Lin, Guo-Jun Qi*
- Recurrent Memory Addressing for Describing Videos, *Arnav Kumar Jain, Abhinav Agarwalla, Kumar Krishna Agrawal, Pabitra Mitra*
- Temporally Steered Gaussian Attention for Video Understanding, *Shagan Sah, Thang Nguyen, Miguel Dominguez, Felipe Petroski Such, Raymond Ptucha*
- SANet: Structure-Aware Network for Visual Tracking, *Heng Fan, Haibin Ling*
- Fixation Prediction in Videos Using Unsupervised Hierarchical Features, *Julius Wang, Hamed R. Tavakoli, Jorma Laaksonen*
- Learning Latent Temporal Connectionism of Deep Residual Visual Abstractions for Identifying Surgical Tools in Laparoscopy Procedures, *Kaustuv Mishra, Rachana Sathish, Debdoot Sheet*
- Kernalised Multi-Resolution Convnet for Visual Tracking, *Di Wu, Wenbin Zou, Xia Li, Yong Zhao*

1600 **Invited Talk:** Tubelet-Based Video Object Detection, *Xiaogang Wang (Chinese Univ. of Hong Kong)*

1635 **Invited Talk:** TBA, *Maja Pantic (Imperial College London)*

1700 **Closing Remarks**

Visual Understanding Across Modalities

Organizers: Ali Farhadi
Abhinav Gupta

Schedule: Full Day

Schedue TBA

Negative Results in Computer Vision

Organizers: Alexei A. Efros
Aude Oliva
Amir R. Zamir
Ozan Sener
Rahul Sukthankar
Antonio Torralba
David Forsyth
William T. Freeman
Jitendra Malik

Schedule: Full Day

0900 **Opening Remarks:** *Ozan Sener, Amir Zamir, Jitendra Malik*

0930 **Invited Talk:** *Jitendra Malik (Univ. of California at Berkeley)*

1015 Morning Break (Kamehameha II)

1045 **Invited Talk:** TBA, *David Forsyth (Univ. of Illinois at Urbana-Champaign)*

1130 **Spotlights**

- No Need to Worry About Adversarial Examples in Object Detection in Autonomous Vehicles, *Jiajun Lu; Hussein Sibai; Evan Fabry; David Forsyth*
- Imposing Hard Constraints on Deep Networks: Promises and Limitations, *Pablo Márquez-Neila; Mathieu Salzmann; Pascal Fua*
- Being Negative but Constructively: Lessons Learned From Creating Better Visual Question Answering Datasets, *Wei-Lun Chao; Hexiang Hu; Fei Sha*
- Negative Results in Computer Vision: A Perspective, *Ali Borji*

1200 Lunch (Kamehameha II)

1330 **Invited Talk:** *Alexei Efros (Univ. of California at Berkeley)*

1415 **Invited Talk:** *Larry Zitnick (Facebook AI Research)*

1500 **Invited Talk:** *Antonio Torralba (MIT)*

1545 Afternoon Break (Kamehameha II)

1615 **Panel:** *Jitendra Malik, David Forsyth, Larry Zitnick, Alexei Efors*

Deep Affective Learning and Context Modeling

Organizers: Ognjen (Oggi) Rudovic
Tadas Baltrušaitis
Daniel McDuff
Rosalind W. Picard

Schedule: Half Day — Morning

0800 **Welcome**

0810 **Keynote Talk:** *Tanzeem Choudhury (Cornell Univ.)*

0900 Action-Affect-Gender Classification Using Multi-Task Representation Learning, *Timothy J. Shields; Mohamed R. Amer; Max Ehrlich; Amir Tamrakar*

0920 DyadGAN: Generating Facial Expressions in Dyadic Interactions, *Yuchi Huang; Saad M. Khan*

0940 Exploring Contextual Engagement for Trauma Recovery , *Svati Dhamija; Terrance E. Boult*

1000 Morning Break (Kamehameha II) **& Poster Session**

- Facial Expression Recognition Using Enhanced Deep 3D Convolutional Neural Networks, *Behzad Hasani; Mohammad H. Mahoor*
- DeepSpace: Mood-Based Image Texture Generation for Virtual Reality From Music, *Misha Sra; Prashanth Vijayaraghavan; Ognjen (Oggi) Rudovic; Pattie Maes; Deb Roy*
- It's Written All Over Your Face: Full-Face Appearance-Based Gaze Estimation, *Xucong Zhang; Yusuke Sugano; Mario Fritz; Andreas Bulling*

1045 **Keynote Talk:** *Ruslan Salakhutdinov (Carnegie Mellon Univ.)*

1135 EMOTIC: Emotions in Context Dataset, *Ronak Kosti; Jose M. Alvarez; Adria Recasens; Agata Lapedriza*

1155 Personalized Automatic Estimation of Self-Reported Pain Intensity From Facial Expressions, *Daniel Lopez Martinez; Ognjen (Oggi) Rudovic; Rosalind Picard*

1215 Speech-Driven 3D Facial Animation With Implicit Emotional Awareness: A Deep Learning Approach, *Hai X. Pham; Samuel Cheung; Vladimir Pavlovic*

1235 Lunch (Kamehameha II)

The DAVIS Challenge on Video Object Segmentation

Organizers: Jordi Pont-Tuset
Federico Perazzi
Sergi Caelles Prat
Alexander Sorkine-Hornung
Pablo Arbeláez
Luc Van Gool

Schedule: Half Day — Morning

0845 **Welcome**

0900 **Invited Talk:** *Chenliang Xu (Univ. of Rochester)*

0920 **Poster Session:** Top Ten Participants

0940 **Invited Talk:** *Katerina Fragkiadaki (Carnegie Mellon Univ.)*

1000 Morning Break (Kamehameha II)

1030 **Oral Presentation:** 4th Place Method

1040 **Oral Presentation:** 3rd Place Method

1050 **Invited Talk:** *Fuxin Li (Oregon State Univ.)*

1110 **Oral Presentation:** 2nd Place Method

1125 **Oral Presentation:** 1st Place Method

1140 **Invited Talk:** *Jitendra Malik (Univ. of California at Berkeley)*

1200 **Awards & Closing Remarks**

1215 **Panel Discussion**

1245 Lunch (Kamehameha II)

Notes:

Women in Computer Vision

Organizers: Samaneh Azadi
Panna Felsen
Timnit Gebru
Roxane Licandro
Negar Rostamzadeh
Gül Varol

Schedule: Half Day — Afternoon

1330 **Introduction**

1340 **Keynote Talk:** Learning to Segment Moving Objects, *Cordelia Schmid (INRIA)*

1400 **Invited Talk:** Gaze Embeddings for Zero-Shot Image Classification, *Nour Karessli (Max Planck Inst. for Informatics)*

1415 **Invited Talk:** Towards Better Instance-Level Recognition, *Georgia Gkioxari (Facebook AI Research)*

1430 **Keynote Talk:** Interferences in Match Kernels , *Naila Murray (Xerox Research Centre Europe)*

1450 Afternoon Break (Kamehameha II) **& Poster Session**

1615 **Keynote Talk:** Computer Vision for the Blind, *Chieko Asakawa (IBM Research)*

1635 **Invited Talk:** Dynamic Deep Neural Networks: Optimizing Accuracy-Efficiency Trade-Offs by Selective Execution, *Lanlan Liu (Univ. of Michigan)*

1650 **Invited Talk**: Semi and Weakly Supervised Semantic Segmentation Using Generative Adversarial Network, *Nasim Souly (Univ. of Central Florida)*

1705 **Panel:** Increasing Diversity in Computer Vision
Panelists:

- Chieko Asakawa *(IBM Research)*
- Andrea Frome *(Clarifai)*
- Raia Hadsell *(Google DeepMind)*
- Naila Murray *(Xerox Research Centre Europe)*
- Cordelia Schmid *(INRIA)*
- Helge Seetzen *(TandemLaunch)*

1735 **Closing Remarks**

ActivityNet: Large Scale Activity Recognition

Organizers: Fabian Caba Heilbron
Cees G. M. Snoek
Juan Carlos Niebles
Bernard Ghanem

Schedule: Half Day — Afternoon

1400 **Opening Remarks**

1415 **Keynote Talk:** TBA, *Jitendra Malik (Univ. of California, Berkeley)*

1445 **Keynote Talk:** Learning Video Representations From the Kinetics Dataset, *Joao Carreira (Google DeepMind)*

1515 **Trimmed Action Recognition:** Dataset & Results

1525 **Trimmed Action Recognition:** Winner Talk

1545 Afternoon Break (Kamehameha II)

1615 **Untrimmed Video Classification:** Dataset & Results

1625 **Untrimmed Video Classification:** Winner Talk

1645 **Temporal Action Proposals:** Dataset & Results

1655 **Temporal Action Proposals:** Winner Talk

1715 **Temporal Action Localization:** Dataset & Results

1725 **Temporal Action Localization:** Winner Talk

1745 **Dense-Captioning Events in Videos:** Dataset & Results

1755 **Dense-Captioning Events in Videos:** Winner Talk

1815 **Closing Remarks**

Notes:


Silver Donors
UBER ATG
Kitware
DATATANG
VionVision
TOYOTA RESEARCH INSTITUTE
Disney Research
Pinterest
AIMATTER
XILINX ALL PROGRAMMABLE
HRI
SIEMENS Healthineers
Bronze Donors
MathWorks
matterport
Mapillary
Orbital Insight
ELSEVIER
BODY LABS
Second Spectrum
Lunit
OMRON Sensing tomorrow
Springer
FLIR
SPORTLOGiQ
VCC VISUAL COMPUTING CENTER
DEEP VISION
CloudSight
Horizon Robotics
唯品会 vip.com
LINKFACE
ALTUM VIEW
Panasonic
MVTec Software GmbH
AXON
SEGWAY ROBOTICS
头条 TOUTIAO
3dMD
cruise
Zillow
AiCure
ZOOX
Skydio
gumgum
Synaptics
NIO
ObjectVideo LABS
SIEMENS Ingenuity for life
NetPosa
Yandex
playground
CrowdFlower
Startup Donors
MORPX
PERCEPTIVE AUTOMATA
wrnch
iniLabs
AutoX
OCTI
umbo CV
TURING
FEATURE X
ANANTAK
spotscale
deepcogn ai
SURFING TECH
AUGMENTED PIXELS
VanGogh Imaging
EgoVid
ISEE
Shopagon
speechocean 海天瑞声
水滴科技 WATRIX TECHNOLOGY
MUKH TECHNOLOGIES
PIXM
MARKABLE
SAIKOU OPTICS
eyeris
EYENUK
Non-Profit Donors
sama SOURCE
NEXT AI

Platinum Donors

Gold Donors