
IEEE Computer Society Conference on Computer Vision and Pattern Recognition

Pocket Guide (Main Conference)

ustilligitit.

# July 21 July 26, 2017 Honolulu, Hawaii



#### Aloha!

Welcome to Honolulu, Hawaii and the 3oth 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of which 2620 were fully reviewed (the others were administratively rejected for technical or ethical reasons or were withdrawn before review).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 nonconflicted 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

## **QR Codes for the IEEE CPS Mobile App**

## Apple iTunes App Store (iOS)





## **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 Yanxi Liu		Shuicheng Yan Mei Han
	Ying Wu Camillo Taylor	Doctoral Consortium Chairs:	Adriana Kovashka 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: Demos Chair:	Walter Scheirer Le Lu	Program Coordination Lead:	Christopher Funk

## **CVPR 2017 Area Chairs**

Shai Avidan Dhruv Batra Margrit Betke Gabriel Brostow Octavia Camps Peter Carr Xilin Chen Bob Collins Jason Corso David Crandall Kristin Dana Kostas Daniilidis **Rogerio Feris** Cornelia Fermuller Vittorio Ferrari Sanja Fidler David Forsyth Charless Fowlkes Dieter Fox William Freeman Pascal Fua Yasutaka Furukawa Andreas Geiger Ross Girshick Bohvung Han Martial Hebert Gang Hua David Jacobs Jiaya Jia Neel Joshi Sing Bing Kang Kris Kitani Pushmeet Kohli Nikos Komodakis Pawan Kumar Kyros Kutulakos Erik Learned-Miller Honglak Lee Fuxin Li Ce Liu Xiaoming Liu Zicheng Liu David Lowe Le Lu

Simon Lucey liebo Luo Dimitris Metaxas Greg Mori **PJ**Narayanan Ko Nishino Nikos Paragios Devi Parikh Pietro Perona Hamed Pirsiavash Robert Pless Marc Pollefeys Amit Roy-Chowdhury Imari Sato Yoichi Sato Daniel Scharstein Stan Sclaroff Nicu Sebe Greg Shakhnarovich Eli Shechtman Jianbo Shi Jian Sun

Oi Tian Sinisa Todorovic Carlo Tomasi Lorenzo Torresani Alexander Toshev Matthew Turk René Vidal Jingdong Wang Xiaogang Wang Daphna Weinshall Lior Wolf Jianxin Wu Ruigang Yang Zhaozheng Yin Jingyi Yu Junsong Yuan Cha Zhang S. Kevin Zhou Song-Chun Zhu

## **CVPR 2017 Outstanding Reviewers**

We are pleased to recognize the following researchers as "CVPR 2017 Outstanding Reviewers". These reviewers were identified by one or more of the CVPR Area Chairs for their hard work in providing high quality and detailed reviews for their assigned papers.

Radhakrishna Achanta Aishwarya Agrawal Xavier Alameda-Pineda Jose M. Alvarez Yannis Avrithis Xiang Bai Ohad Ben-Shahar Katie Bouman Stefan Breuers Michael Brown Thomas Brox Dylan Campbell Ayan Chakrabarti Anoop Cherian Sunghyun Cho Yung-Yu Chuang Michael Cogswell Daniel Cremers Marco Cristani Abhishek Das Luca del Pero Sergio Escalera Victor Escorcia Francisco Estrada Boris Flach Wolfgang Foerstner Victor Fragoso Orazio Gallo Ioannis Gkioulekas Georgia Gkioxari Michael Goesele Boging Gong Yahong Han

Christian Häne Michal Havlena Kaiming He Jia-Bin Huang Go Irie Nathan Jacobs Varun Jampani C.V. Jawahar Laszlo leni Kevin Karsch Laurent Kneip Piotr Koniusz Adriana Kovashka Abhijit Kundu Jean-Francois Lalonde Diane Larlus Stefan Lee Yinli Chia-Kai Liang Yen-Yulin Jonathan Long Chen Change Loy Mohammed Mahoor Michael Maire Yasushi Makihara Mateusz Malinowski Renaud Marlet Stefan Mathe Mason McGill Thomas Mensink Phillippos Mordohai Roozbeh Mottaghi Nikhil Naik

Martin R. Oswald Tomas Paidla Vishal Patel Ioannis Patras **Bojan Pepik** Rene Ranftl Michal Reinstein Jerome Revaud Elisa Ricci Gernot Riegler Ergys Ristani Emanuele Rodola Gregory Rogez Marcus Rohrbach Matteo Ronchi Samuel Rota Bulò Stefan Roth Fereshteh Sadeghi Torsten Sattler Johannes Schoenberger Alex Schwing Amir Shahroudy Chunhua Shen Boxin Shi Arnold Smeulders Yale Song Alexander Sorkine-Hornung Hang Su Shuochen Su Yusuke Sugano David Suter Chris Sweeney Yuichi Taguchi

Gabriel Taubin Joseph Tighe Federico Tombari Akihiko Torii Georgios Tzimiropoulos Osman Ulusov Michel Valstar Jan van Gemert Grant van Horn Ramakrishna Vedantam Christoph Vogel Christian Vogler Catherine Wah Limin Wang Wei Wang Xiu-Shen Wei Kyle Wilson David Wipf Tianfu Wu Lingxi Xie Kota Yamaguchi linwei Ye Rvo Yonetani Lu Yuan Stefanos Zafeiriou Jianming Zhang Lei Zhang Wei-Shi Zheng Jun-yan Zhu Yixin Zhu Silva Zuffi

# Saturday, July 22

### 0700-1700 Registration (Main Lobby)

#### 0730-0830 Breakfast (Kamehameha II)

0830-0845 Opening Remarks & Paper Awards (Kamehameha III)

#### 0900-1030 Session 1-1A: Machine Learning 1 (Kamehameha III)

#### Papers in this session are also in Poster Session P1-1.

Chairs: Sanja Fidler (Univ. of Toronto) Kris Kitani (Carnegie Mellon Univ.)

#### 0900 Spotlights (S1-1A)

Format (4 min. for presentation; no questions)

- [ogoo] Exclusivity-Consistency Regularized Multi-View Subspace Clustering, Xiaobo Wang, Xiaojie Guo, Zhen Lei, Changqing Zhang, Stan Z. Li
- [0904] Borrowing Treasures From the Wealthy: Deep Transfer Learning Through Selective Joint Fine-Tuning, Weifeng Ge, Yizhou Yu
- [ogo8] The More You Know: Using Knowledge Graphs for Image Classification, Kenneth Marino, Ruslan Salakhutdinov, Abhinav Gupta
- [0912] Dynamic Edge-Conditioned Filters in Convolutional Neural Networks on Graphs, Martin Simonovsky, Nikos Komodakis
- [0916] Convolutional Neural Network Architecture for Geometric Matching, Ignacio Rocco, Relja Arandjelović, Josef Sivic
- 6. **[0920]** Deep Affordance-Grounded Sensorimotor Object Recognition, *Spyridon Thermos, Georgios Th. Papadopoulos, Petros Daras, Gerasimos Potamianos*
- [0924] Discovering Causal Signals in Images, David Lopez-Paz, Robert Nishihara, Soumith Chintala, Bernhard Schölkopf, Léon Bottou

8. **[0928]** On Compressing Deep Models by Low Rank and Sparse Decomposition, *Xiyu Yu, Tongliang Liu, Xinchao Wang, Dacheng Tao* 

#### 0933 Orals (O1-1A)

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

- [0933] PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation, Charles R. Qi, Hao Su, Kaichun Mo, Leonidas J. Guibas
- [0947] Universal Adversarial Perturbations, Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Omar Fawzi, Pascal Frossard
- 11. [1001] Unsupervised Pixel-Level Domain Adaptation With Generative Adversarial Networks, Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, Dilip Krishnan
- 12. [1015] Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network, Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, Wenzhe Shi

#### 0900–1030 Session 1-1B: 3D Vision 1 (Kalākaua Ballroom A-B)

#### Papers in this session are also in Poster Session P1-1.

Chairs: Andreas Geiger (MPI Tübingen) Ruigang Yang (Univ. of Kenucky)

#### 0900 Spotlights (S1-1B)

Format (4 min. for presentation; no questions)

- 13. **[0900]** Global Hypothesis Generation for 6D Object Pose Estimation, Frank Michel, Alexander Kirillov, Eric Brachmann, Alexander Krull, Stefan Gumhold, Bogdan Savchynskyy, Carsten Rother
- 14. **[0904]** A Practical Method for Fully Automatic Intrinsic Camera Calibration Using Directionally Encoded Light, *Mahdi Abbaspour Tehrani, Thabo Beeler, Anselm Grundhöfer*
- 15. [0908] CATS: A Color and Thermal Stereo Benchmark, Wayne Treible, Philip Saponaro, Scott Sorensen, Abhishek Kolagunda, Michael O'Neal, Brian Phelan, Kelly Sherbondy, Chandra Kambhamettu

## Program

- 16. **[0912]** Elastic Shape-From-Template With Spatially Sparse Deforming Forces, *Abed Malti, Cédric Herzet*
- 17. **[0916]** Distinguishing the Indistinguishable: Exploring Structural Ambiguities via Geodesic Context, *Qingan Yan*, Long Yang, Ling Zhang, Chunxia Xiao
- 18. [0920] Multi-Scale Continuous CRFs as Sequential Deep Networks for Monocular Depth Estimation, Dan Xu, Elisa Ricci, Wanli Ouyang, Xiaogang Wang, Nicu Sebe
- 19. **[0924]** Dynamic Time-Of-Flight, Michael Schober, Amit Adam, Omer Yair, Shai Mazor, Sebastian Nowozin
- 20. **[0928]** Training Object Class Detectors With Click Supervision, *Dim P. Papadopoulos, Jasper R. R. Uijlings, Frank Keller, Vittorio Ferrari*

## 0933 Orals (O1-1B)

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

- 21. [0933] Semantic Scene Completion From a Single Depth Image, Shuran Song, Fisher Yu, Andy Zeng, Angel X. Chang, Manolis Savva, Thomas Funkhouser
- 22. **[0947]** 3DMatch: Learning Local Geometric Descriptors From RGB-D Reconstructions, Andy Zeng, Shuran Song, Matthias Nießner, Matthew Fisher, Jianxiong Xiao, Thomas Funkhouser
- 23. **[1001]** Multi-View Supervision for Single-View Reconstruction via Differentiable Ray Consistency, *Shubham Tulsiani, Tinghui Zhou, Alexei A. Efros, Jitendra Malik*
- 24. [1015] On-The-Fly Adaptation of Regression Forests for Online Camera Relocalisation, *Tommaso Cavallari, Stuart Golodetz, Nicholas A. Lord, Julien Valentin, Luigi Di Stefano, Philip H. S. Torr*

## 0900–1030 Session 1-1C: Low- & Mid-Level Vision (Kalākaua Ballroom C)

## Papers in this session are also in Poster Session P1-1.

Chairs: Octavia Camps (Northeastern Univ.) David Jacobs (Univ. of Maryland)

## 0900 Spotlights (S1-1C)

## Format (4 min. for presentation; no questions)

25. [0900] Designing Effective Inter-Pixel Information Flow for Natural Image Matting, Yağiz Aksoy, Tunç Ozan Aydin, Marc Pollefeys

- 26. **[0904]** Deep Video Deblurring for Hand-Held Cameras, Shuochen Su, Mauricio Delbracio, Jue Wang, Guillermo Sapiro, Wolfgang Heidrich, Oliver Wang
- 27. **[0908]** Instance-Level Salient Object Segmentation, *Guanbin Li, Yuan Xie, Liang Lin, Yizhou Yu*
- 28. **[0912]** Deep Multi-Scale Convolutional Neural Network for Dynamic Scene Deblurring, *Seungjun Nah, Tae Hyun Kim, Kyoung Mu Lee*
- 29. [0916] Diversified Texture Synthesis With Feed-Forward Networks, Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, Ming-Hsuan Yang
- 30. [0920] Radiometric Calibration for Internet Photo Collections, Zhipeng Mo, Boxin Shi, Sai-Kit Yeung, Yasuyuki Matsushita
- 31. **[0924]** Deeply Aggregated Alternating Minimization for Image Restoration, *Youngjung Kim, Hyungjoo Jung, Dongbo Min, Kwanghoon Sohn*
- 32. **[0928]** End-To-End Instance Segmentation With Recurrent Attention, *Mengye Ren, Richard S. Zemel*

### 0933 Orals (O1-1C)

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

- 33. **[0933]** SRN: Side-output Residual Network for Object Symmetry Detection in the Wild, *Wei Ke, Jie Chen, Jianbin Jiao, Guoying Zhao, Qixiang Ye*
- 34. **[0947]** Deep Image Matting, *Ning Xu, Brian Price, Scott Cohen, Thomas Huang*
- 35. **[1001]** Wetness and Color From a Single Multispectral Image, *Mihoko Shimano, Hiroki Okawa, Yuta Asano, Ryoma Bise, Ko Nishino, Imari Sato*
- 36. **[1015]** FC<sup>4</sup>: Fully Convolutional Color Constancy With Confidence-Weighted Pooling, *Yuanming Hu, Baoyuan Wang, Stephen Lin*

## 1030-1115 Break (Kamehameha II)

### 1030-1230 Poster Session P1-1 (Kamehameha I)

#### 3D Computer Vision

- Face Normals "In-The-Wild" Using Fully Convolutional Networks, George Trigeorgis, Patrick Snape, Iasonas Kokkinos, Stefanos Zafeiriou
- 38. A Non-Convex Variational Approach to Photometric Stereo Under Inaccurate Lighting, Yvain Quéau, Tao Wu, François Lauze, Jean-Denis Durou, Daniel Cremers
- 39. A Linear Extrinsic Calibration of Kaleidoscopic Imaging System From Single 3D Point, Kosuke Takahashi, Akihiro Miyata, Shohei Nobuhara, Takashi Matsuyama
- Polarimetric Multi-View Stereo, Zhaopeng Cui, Jinwei Gu, Boxin Shi, Ping Tan, Jan Kautz
- 41. An Exact Penalty Method for Locally Convergent Maximum Consensus, *Huu Le, Tat-Jun Chin, David Suter*
- 42. Deep Supervision With Shape Concepts for Occlusion-Aware 3D Object Parsing, Chi Li, M. Zeeshan Zia, Quoc-Huy Tran, Xiang Yu, Gregory D. Hager, Manmohan Chandraker
- 43. Amodal Detection of 3D Objects: Inferring 3D Bounding Boxes From 2D Ones in RGB-Depth Images, *Zhuo Deng*, *Longin Jan Latecki*

#### Analyzing Humans in Images

- 44. Transition Forests: Learning Discriminative Temporal Transitions for Action Recognition and Detection, *Guillermo Garcia-Hernando, Tae-Kyun Kim*
- 45. Scene Flow to Action Map: A New Representation for RGB-D Based Action Recognition With Convolutional Neural Networks, Pichao Wang, Wanqing Li, Zhimin Gao, Yuyao Zhang, Chang Tang, Philip Ogunbona
- 46. Detecting Masked Faces in the Wild With LLE-CNNs, Shiming Ge, Jia Li, Qiting Ye, Zhao Luo
- 47. A Domain Based Approach to Social Relation Recognition, *Qianru Sun, Bernt Schiele, Mario Fritz*
- Spatio-Temporal Naive-Bayes Nearest-Neighbor (ST-NBNN) for Skeleton-Based Action Recognition, Junwu Weng, Chaoqun Weng, Junsong Yuan
- 49. Personalizing Gesture Recognition Using Hierarchical Bayesian Neural Networks, Ajjen Joshi, Soumya Ghosh, Margrit Betke, Stan Sclaroff, Hanspeter Pfister

#### **Applications**

- 50. Real-Time 3D Model Tracking in Color and Depth on a Single CPU Core, *Wadim Kehl, Federico Tombari, Slobodan Ilic, Nassir Navab*
- 51. Multi-Scale FCN With Cascaded Instance Aware Segmentation for Arbitrary Oriented Word Spotting in the Wild, Dafang He, Xiao Yang, Chen Liang, Zihan Zhou, Alexander G. Ororbi II, Daniel Kifer, C. Lee Giles
- 52. Viraliency: Pooling Local Virality, Xavier Alameda-Pineda, Andrea Pilzer, Dan Xu, Nicu Sebe, Elisa Ricci

#### Biomedical Image/Video Analysis

53. A Non-Local Low-Rank Framework for Ultrasound Speckle Reduction, Lei Zhu, Chi-Wing Fu, Michael S. Brown, Pheng-Ann Heng

#### Image Motion & Tracking

- 54. Video Acceleration Magnification, Yichao Zhang, Silvia L. Pintea, Jan C. van Gemert
- 55. Superpixel-Based Tracking-By-Segmentation Using Markov Chains, Donghun Yeo, Jeany Son, Bohyung Han, Joon Hee Han
- 56. BranchOut: Regularization for Online Ensemble Tracking With Convolutional Neural Networks, *Bohyung Han, Jack Sim, Hartwig Adam*
- 57. Learning Motion Patterns in Videos, Pavel Tokmakov, Karteek Alahari, Cordelia Schmid

#### Low- & Mid-Level Vision

- 58. Deep Level Sets for Salient Object Detection, Ping Hu, Bing Shuai, Jun Liu, Gang Wang
- 59. Binary Constraint Preserving Graph Matching, *Bo Jiang*, Jin Tang, Chris Ding, Bin Luo
- 60. From Local to Global: Edge Profiles to Camera Motion in Blurred Images, *Subeesh Vasu, A. N. Rajagopalan*
- 61. What Is the Space of Attenuation Coefficients in Underwater Computer Vision? Derya Akkaynak, Tali Treibitz, Tom Shlesinger, Yossi Loya, Raz Tamir, David Iluz
- 62. Robust Energy Minimization for BRDF-Invariant Shape From Light Fields, Zhengqin Li, Zexiang Xu, Ravi Ramamoorthi, Manmohan Chandraker
- 63. Boundary-Aware Instance Segmentation, Zeeshan Hayder, Xuming He, Mathieu Salzmann

## Program

- 64. Spatially-Varying Blur Detection Based on Multiscale Fused and Sorted Transform Coefficients of Gradient Magnitudes, S. Alireza Golestaneh, Lina J. Karam
- 65. Model-Based Iterative Restoration for Binary Document Image Compression With Dictionary Learning, Yandong Guo, Cheng Lu, Jan P. Allebach, Charles A. Bouman
- 66. FCSS: Fully Convolutional Self-Similarity for Dense Semantic Correspondence, Seungryong Kim, Dongbo Min, Burnsub Ham, Sangryul Jeon, Stephen Lin, Kwanghoon Sohn

#### Machine Learning

- Learning by Association A Versatile Semi-Supervised Training Method for Neural Networks, *Philip Haeusser, Alexander Mordvintsev, Daniel Cremers*
- 68. Dilated Residual Networks, Fisher Yu, Vladlen Koltun, Thomas Funkhouser
- Split-Brain Autoencoders: Unsupervised Learning by Cross-Channel Prediction, *Richard Zhang, Phillip Isola, Alexei A. Efros*
- 70. Nonnegative Matrix Underapproximation for Robust Multiple Model Fitting, Mariano Tepper, Guillermo Sapiro
- 71. Truncated Max-Of-Convex Models, Pankaj Pansari, M. Pawan Kumar
- 72. Additive Component Analysis, Calvin Murdock, Fernando De la Torre
- 73. Subspace Clustering via Variance Regularized Ridge Regression, Chong Peng, Zhao Kang, Qiang Cheng
- 74. The Incremental Multiresolution Matrix Factorization Algorithm, Vamsi K. Ithapu, Risi Kondor, Sterling C. Johnson, Vikas Singh
- 75. Transformation-Grounded Image Generation Network for Novel 3D View Synthesis, *Eunbyung Park, Jimei Yang, Ersin Yumer, Duygu Ceylan, Alexander C. Berg*
- 76. Learning Dynamic Guidance for Depth Image Enhancement, Shuhang Gu, Wangmeng Zuo, Shi Guo, Yunjin Chen, Chongyu Chen, Lei Zhang
- 77. A-Lamp: Adaptive Layout-Aware Multi-Patch Deep Convolutional Neural Network for Photo Aesthetic Assessment, Shuang Ma, Jing Liu, Chang Wen Chen
- 78. Teaching Compositionality to CNNs, Austin Stone, Huayan Wang, Michael Stark, Yi Liu, D. Scott Phoenix, Dileep George

- 79. Using Ranking-CNN for Age Estimation, *Shixing Chen*, *Caojin Zhang, Ming Dong, Jialiang Le, Mike Rao*
- Accurate Single Stage Detector Using Recurrent Rolling Convolution, Jimmy Ren, Xiaohao Chen, Jianbo Liu, Wenxiu Sun, Jiahao Pang, Qiong Yan, Yu-Wing Tai, Li Xu
- A Compact DNN: Approaching GoogLeNet-Level Accuracy of Classification and Domain Adaptation, Chunpeng Wu, Wei Wen, Tariq Afzal, Yongmei Zhang, Yiran Chen, Hai (Helen) Li
- 82. The Impact of Typicality for Informative Representative Selection, Jawadul H. Bappy, Sujoy Paul, Ertem Tuncel, Amit K. Roy-Chowdhury
- 83. Infinite Variational Autoencoder for Semi-Supervised Learning, M. Ehsan Abbasnejad, Anthony Dick, Anton van den Hengel
- 84. SurfNet: Generating 3D Shape Surfaces Using Deep Residual Networks, Ayan Sinha, Asim Unmesh, Qixing Huang, Karthik Ramani
- 85. Intrinsic Grassmann Averages for Online Linear and Robust Subspace Learning, Rudrasis Chakraborty, Søren Hauberg, Baba C. Vemuri
- Variational Bayesian Multiple Instance Learning With Gaussian Processes, Manuel Haußmann, Fred A. Hamprecht, Melih Kandemir
- 87. Temporal Attention-Gated Model for Robust Sequence Classification, *Wenjie Pei, Tadas Baltrušaitis, David M.J. Tax, Louis-Philippe Morency*
- Non-Uniform Subset Selection for Active Learning in Structured Data, Sujoy Paul, Jawadul H. Bappy, Amit K. Roy-Chowdhury
- 89. Colorization as a Proxy Task for Visual Understanding, Gustav Larsson, Michael Maire, Gregory Shakhnarovich
- 90. Shading Annotations in the Wild, Balazs Kovacs, Sean Bell, Noah Snavely, Kavita Bala
- 91. LCNN: Lookup-Based Convolutional Neural Network, Hessam Bagherinezhad, Mohammad Rastegari, Ali Farhadi

#### **Object Recognition & Scene Understanding**

92. Physics Inspired Optimization on Semantic Transfer Features: An Alternative Method for Room Layout Estimation, *Hao Zhao, Ming Lu, Anbang Yao, Yiwen Guo, Yurong Chen, Li Zhang* 

## Program

- Pixelwise Instance Segmentation With a Dynamically Instantiated Network, Anurag Arnab, Philip H. S. Torr
- 94. Object Detection in Videos With Tubelet Proposal Networks, Kai Kang, Hongsheng Li, Tong Xiao, Wanli Ouyang, Junjie Yan, Xihui Liu, Xiaogang Wang
- 95. AMVH: Asymmetric Multi-Valued Hashing, Cheng Da, Shibiao Xu, Kun Ding, Gaofeng Meng, Shiming Xiang, Chunhong Pan
- 96. Spindle Net: Person Re-Identification With Human Body Region Guided Feature Decomposition and Fusion, *Haiyu* Zhao, Maoqing Tian, Shuyang Sun, Jing Shao, Junjie Yan, Shuai Yi, Xiaogang Wang, Xiaoou Tang
- 97. Deep Visual-Semantic Quantization for Efficient Image Retrieval, Yue Cao, Mingsheng Long, Jianmin Wang, Shichen Liu
- 98. Efficient Diffusion on Region Manifolds: Recovering Small Objects With Compact CNN Representations, Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, Teddy Furon, Ondřej Chum
- 99. Feature Pyramid Networks for Object Detection, *Tsung-Yi* Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, Serge Belongie
- 100. Mind the Class Weight Bias: Weighted Maximum Mean Discrepancy for Unsupervised Domain Adaptation, Hongliang Yan, Yukang Ding, Peihua Li, Qilong Wang, Yong Xu, Wangmeng Zuo
- 101. StyleNet: Generating Attractive Visual Captions With Styles, Chuang Gan, Zhe Gan, Xiaodong He, Jianfeng Gao, Li Deng
- 102. Fine-Grained Recognition of Thousands of Object Categories With Single-Example Training, Leonid Karlinsky, Joseph Shtok, Yochay Tzur, Asaf Tzadok
- 103. Improving Interpretability of Deep Neural Networks With Semantic Information, *Yinpeng Dong, Hang Su, Jun Zhu, Bo Zhang*
- 104. Video Captioning With Transferred Semantic Attributes, Yingwei Pan, Ting Yao, Houqiang Li, Tao Mei
- 105. Fast Boosting Based Detection Using Scale Invariant Multimodal Multiresolution Filtered Features, Arthur Daniel Costea, Robert Varga, Sergiu Nedevschi

#### Video Analytics

- 106. Temporal Convolutional Networks for Action Segmentation and Detection, Colin Lea, Michael D. Flynn, René Vidal, Austin Reiter, Gregory D. Hager
- 107. Surveillance Video Parsing With Single Frame Supervision, Si Liu, Changhu Wang, Ruihe Qian, Han Yu, Renda Bao, Yao Sun
- 108.Weakly Supervised Actor-Action Segmentation via Robust Multi-Task Ranking, Yan Yan, Chenliang Xu, Dawen Cai, Jason J. Corso
- 109. Unsupervised Visual-Linguistic Reference Resolution in Instructional Videos, *De-An Huang, Joseph J. Lim, Li Fei-Fei, Juan Carlos Niebles*
- 110. Zero-Shot Action Recognition With Error-Correcting Output Codes, Jie Qin, Li Liu, Ling Shao, Fumin Shen, Bingbing Ni, Jiaxin Chen, Yunhong Wang
- 111. Enhancing Video Summarization via Vision-Language Embedding, Bryan A. Plummer, Matthew Brown, Svetlana Lazebnik
- 112. Synthesizing Dynamic Patterns by Spatial-Temporal Generative ConvNet, *Jianwen Xie, Song-Chun Zhu, Ying Nian Wu*

## 1030-1230 Demos (Kamehameha I)

- Tensorflow Object Detection API, Jonathan Huang, Vivek Rathod, Chen Sun, Sergio Guadarrama, Tyler Zhu, George Papandreou, Menglong Zhu, Alireza Fathi, Derek Chow, Kevin Murphy (Google)
- DenseReg: Fully Convolutional Dense Shape Regression In-The-Wild, Rıza Alp Güler, George Trigeorgis, Epameinondas Antonakos, Patrick Snape, Stefanos Zafeiriou, Iasonas Kokkinos (INRIA/CentraleSupélec, Imperial College, Univ. College London/Facebook Al Research)
- UnrealCV: Connecting Computer Vision to Unreal Engine, Weichao Qiu, Yi Zhang, Fangwei Zhong, Siyuan Qiao, Yizhou Wang, Alan Yuille (Johns Hopkins Univ., Tsinghua Univ., Peking Univ.)
- Minimum Delay Moving Object Detection, Samim Zahoor, Dong Lao, Ganesh Sundaramoorthi (King Abdullah Univ. of Science and Technology (KAUST))

## 1030-1230 Exhibitors (Kamehameha I)

### Booth # in parenthesis—see map on next page (pg. 11).

### <u>Platinum Level</u>

- A9 (305)
- Alibaba (116)
- Amazon (112)
- Apple (117)
- DeepGlint & UISEE (405)
- Didi Chuxing (605)
- Facebook (205)
- Google (106)
- Intel Nervana (635)

## <u>Gold Level</u>

- Adobe Research (737)
- Argo Al (705)
- Bosch (856)
- Cognex (727)
- CVTE (222)
- DJI (707)
- HERE (135)
- HiScene (647)
- IBM (230)
- iRobot (127)
- Malong (226)

#### Silver Level

- AIMATTER (749)
- Datatang (246)
- Disney Research (147)
- Honda Research Inst. (853)
- Kitware (242)
- Pinterest (143)

#### Bronze Level

- 3dMD (857)
- AiCure (759)
- AltumView Systems (157)
- Axon (750)
- Body Labs (238)
- CloudSight (649)
- CrowdFlower (849)
- Cruise Automation (254)
- DeepVision (725)
- Elsevier (159)
- FLIR Systems (250)
- GumGum (653)

- JD-X RDC (146)
- Microsoft (136)
- Momenta (111)
- Movidius (629)
- NAVER (505)
- NVIDIA (105)
- Samsung (617)
- Tencent (126)
- Megvii Face++ (855)
- Meitu (129)
- MERL (224)
- Mighty Ai (125)
- nuTonomy (752)
- Qihoo/360 (850)
- SAP SE (861)
- SenseTime (131)
- Sighthound (228)
- Snapchat (123)
- TuSimple (234)
- Siemens Healthineers (852)
- Toyota Research Inst. (711)
- Uber ATG (236)
- Vion Technologies (741)
- Xilinx (848)
- NIO (747)
- ObjectVideo Labs (757)
- Omron (139)
- Orbital Insight (625)
- Panasonic (713)
- Playaround (713
- Playground (756)Second Spectrum (155)
- Second Spectrum (155)
  Segway Robotics (560)
- Siemens (854)
- Skydio (256)
- SportLogig (252)
- SportLogiq (252)
- Springer (248)

• Horizon Robotics (149)

Program

Synaptics (755)

• Vipshop US (655)

Yandex (709)

• Zillow (151)

ZOOX (715)

Octi (717)

• Pixm (731)

TouTiao/ByteDance (722)

Perceptive Automata (724)

Saikou Optics (754)

Speechocean (851)

Surfing Tech (845)

• TuringVideo (563)

VanGogh Imaging (753)

Watrix Technology (660)

Shopagon (758)

Spotscale (739)

Umbo CV (657)

WRNCH (641)

Samasource (735)

Tesla (729)

- KAUST (258)
- Linkface (748)
- Lunit (240)
- Mapillary (137)
- MathWorks (153)
- MatterPort (141)
- NetPosa (847)

#### <u>Start-Up Level</u>

- Anantak (746)
- Augmented Pixels (751)
- AutoX (561)
- DeepCognAl (661)
- EgoVid (743)
- Eyenuk (846)
- Eyeris (651)
- FeatureX (740)

Markable (461)

Morgan & Claypool (742)

1230-1330 Lunch (Kamehameha II)

• iniLabs (659) • ISEE AI (761)

• Morpx (760)


• MUKH (719)

<u>Non-Profit Level</u>

NextAl (723)

Notes:

10



## Program

#### 1330–1500 Session 1-2A: Object Recognition & Scene Understanding - Computer Vision & Language (Kamehameha III)

Papers in this session are also in Poster Session P1-2.

Chairs: Devi Parikh (Georgia Tech) Nicu Sebe (Univ. of Trento)

#### 1330 Spotlights (S1-2A)

Format (4 min. for presentation; no questions)

- [1330] Context-Aware Captions From Context-Agnostic Supervision, Ramakrishna Vedantam, Samy Bengio, Kevin Murphy, Devi Parikh, Gal Chechik
- [1334] Visual Dialog, Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José M. F. Moura, Devi Parikh, Dhruv Batra
- [1338] Discriminative Bimodal Networks for Visual Localization and Detection With Natural Language Queries, Yuting Zhang, Luyao Yuan, Yijie Guo, Zhiyuan He, I-An Huang, Honglak Lee
- 4. [1342] Automatic Understanding of Image and Video Advertisements, Zaeem Hussain, Mingda Zhang, Xiaozhong Zhang, Keren Ye, Christopher Thomas, Zuha Agha, Nathan Ong, Adriana Kovashka
- [1346] Discover and Learn New Objects From Documentaries, Kai Chen, Hang Song, Chen Change Loy, Dahua Lin
- 6. [1350] Spatial-Semantic Image Search by Visual Feature Synthesis, Long Mai, Hailin Jin, Zhe Lin, Chen Fang, Jonathan Brandt, Feng Liu
- [1354] Fully-Adaptive Feature Sharing in Multi-Task Networks With Applications in Person Attribute Classification, Yongxi Lu, Abhishek Kumar, Shuangfei Zhai, Yu Cheng, Tara Javidi, Rogerio Feris
- [1358] Semantic Compositional Networks for Visual Captioning, Zhe Gan, Chuang Gan, Xiaodong He, Yunchen Pu, Kenneth Tran, Jianfeng Gao, Lawrence Carin, Li Deng

#### 1403 Orals (O1-2A)

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

9. **[1403]** Deep Reinforcement Learning-Based Image Captioning With Embedding Reward, *Zhou Ren, Xiaoyu Wang, Ning Zhang, Xutao Lv, Li-Jia Li* 

- 10. **[1417]** From Red Wine to Red Tomato: Composition With Context, *Ishan Misra, Abhinav Gupta, Martial Hebert*
- 11. [1431] Captioning Images With Diverse Objects, Subhashini Venugopalan, Lisa Anne Hendricks, Marcus Rohrbach, Raymond Mooney, Trevor Darrell, Kate Saenko
- 12. [1445] Self-Critical Sequence Training for Image Captioning, Steven J. Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, Vaibhava Goel

### 1330–1500 Session 1-2B: Analyzing Humans 1 (Kalākaua Ballroom A-B)

#### Papers in this session are also in Poster Session P1-2.

Chairs: Greg Mori (Simon Fraser Univ.) Amit Roy-Chowdhury (UC Riverside)

#### 1330 Spotlights (S1-2B)

Format (4 min. for presentation; no questions)

- 13. [1330] Crossing Nets: Combining GANs and VAEs With a Shared Latent Space for Hand Pose Estimation, Chengde Wan, Thomas Probst, Luc Van Gool, Angela Yao
- 14. **[1334]** Predicting Behaviors of Basketball Players From First Person Videos, *Shan Su, Jung Pyo Hong, Jianbo Shi, Hyun Soo Park*
- 15. **[1338]** LCR-Net: Localization-Classification-Regression for Human Pose, *Grégory Rogez, Philippe Weinzaepfel, Cordelia Schmid*
- 16. **[1342]** Learning Residual Images for Face Attribute Manipulation, *Wei Shen, Rujie Liu*
- 17. **[1346]** Seeing What Is Not There: Learning Context to Determine Where Objects Are Missing, *Jin Sun, David W. Jacobs*
- 18. **[1350]** Deep Learning on Lie Groups for Skeleton-Based Action Recognition, *Zhiwu Huang, Chengde Wan, Thomas Probst, Luc Van Gool*
- 19. **[1354]** Harvesting Multiple Views for Marker-Less 3D Human Pose Annotations, *Georgios Pavlakos, Xiaowei* Zhou, Konstantinos G. Derpanis, Kostas Daniilidis
- 20. [1358] Coarse-To-Fine Volumetric Prediction for Single-Image 3D Human Pose, Georgios Pavlakos, Xiaowei Zhou, Konstantinos G. Derpanis, Kostas Daniilidis

## Program

#### 1403 Orals (O1-2B)

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

- 21. [1403] Weakly Supervised Action Learning With RNN Based Fine-To-Coarse Modeling, Alexander Richard, Hilde Kuehne, Juergen Gall
- 22. **[1417]** Disentangled Representation Learning GAN for Pose-Invariant Face Recognition, *Luan Tran, Xi Yin, Xiaoming Liu*
- 23. [1431] ArtTrack: Articulated Multi-Person Tracking in the Wild, Eldar Insafutdinov, Mykhaylo Andriluka, Leonid Pishchulin, Siyu Tang, Evgeny Levinkov, Bjoern Andres, Bernt Schiele
- 24. [1445] Realtime Multi-Person 2D Pose Estimation Using Part Affinity Fields, *Zhe Cao, Tomas Simon, Shih-En Wei, Yaser Sheikh*

### 1330–1500 Session 1-2C: Image Motion & Tracking; Video Analysis (Kalākaua Ballroom C)

#### Papers in this session are also in Poster Session P1-2.

Chairs: Bob Collins (Pennsylvania State Univ.) René Vidal (Johns Hopkins Univ.)

#### 1330 Spotlights (S1-2C)

Format (4 min. for presentation; no questions)

- 25. [1330] Template Matching With Deformable Diversity Similarity, Itamar Talmi, Roey Mechrez, Lihi Zelnik-Manor
- 26. [1334] Beyond Triplet Loss: A Deep Quadruplet Network for Person Re-Identification, Weihua Chen, Xiaotang Chen, Jianguo Zhang, Kaiqi Huang
- 27. [1338] Agent-Centric Risk Assessment: Accident Anticipation and Risky Region Localization, Kuo-Hao Zeng, Shih-Han Chou, Fu-Hsiang Chan, Juan Carlos Niebles, Min Sun
- 28. [1342] Bidirectional Multirate Reconstruction for Temporal Modeling in Videos, *Linchao Zhu, Zhongwen Xu, Yi Yang*
- 29. [1346] Action-Decision Networks for Visual Tracking With Deep Reinforcement Learning, Sangdoo Yun, Jongwon Choi, Youngjoon Yoo, Kimin Yun, Jin Young Choi
- 30. [1350] TGIF-QA: Toward Spatio-Temporal Reasoning in Visual Question Answering, Yunseok Jang, Yale Song, Youngjae Yu, Youngjin Kim, Gunhee Kim

- 31. **[1354]** Making 360° Video Watchable in 2D: Learning Videography for Click Free Viewing, Yu-Chuan Su, Kristen Grauman
- 32. [1358] Unsupervised Adaptive Re-Identification in Open World Dynamic Camera Networks, Rameswar Panda, Amran Bhuiyan, Vittorio Murino, Amit K. Roy-Chowdhury

#### 1403 Orals (O1-2C)

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

- 33. **[1403]** Context-Aware Correlation Filter Tracking, Matthias Mueller, Neil Smith, Bernard Ghanem
- 34. [1417] Deep 360 Pilot: Learning a Deep Agent for Piloting Through 360° Sports Videos, Hou-Ning Hu, Yen-Chen Lin, Ming-Yu Liu, Hsien-Tzu Cheng, Yung-Ju Chang, Min Sun
- 35. [1431] Slow Flow: Exploiting High-Speed Cameras for Accurate and Diverse Optical Flow Reference Data, *Joel* Janai, Fatma Güney, Jonas Wulff, Michael J. Black, Andreas Geiger
- 36. [1445] CDC: Convolutional-De-Convolutional Networks for Precise Temporal Action Localization in Untrimmed Videos, Zheng Shou, Jonathan Chan, Alireza Zareian, Kazuyuki Miyazawa, Shih-Fu Chang

## 1500-1545 Break (Kamehameha II)

### 1500-1700 Poster Session P1-2 (Kamehameha I)

#### 3D Computer Vision

- 37. Exploiting 2D Floorplan for Building-Scale Panorama RGBD Alignment, *Erik Wijmans, Yasutaka Furukawa*
- 38. A Combinatorial Solution to Non-Rigid 3D Shape-To-Image Matching, Florian Bernard, Frank R. Schmidt, Johan Thunberg, Daniel Cremers
- 39. NID-SLAM: Robust Monocular SLAM Using Normalised Information Distance, Geoffrey Pascoe, Will Maddern, Michael Tanner, Pedro Piniés, Paul Newman
- 40. End-To-End Training of Hybrid CNN-CRF Models for Stereo, Patrick Knöbelreiter, Christian Reinbacher, Alexander Shekhovtsov, Thomas Pock
- 41. Learning Shape Abstractions by Assembling Volumetric Primitives, Shubham Tulsiani, Hao Su, Leonidas J. Guibas, Alexei A. Efros, Jitendra Malik

## Program

- 42. Locality-Sensitive Deconvolution Networks With Gated Fusion for RGB-D Indoor Semantic Segmentation, Yanhua Cheng, Rui Cai, Zhiwei Li, Xin Zhao, Kaiqi Huang
- 43. Acquiring Axially-Symmetric Transparent Objects Using Single-View Transmission Imaging, Jaewon Kim, Ilya Reshetouski, Abhijeet Ghosh
- 44. Regressing Robust and Discriminative 3D Morphable Models With a Very Deep Neural Network, Anh Tuấn Trần, Tal Hassner, Iacopo Masi, Gérard Medioni
- 45. End-To-End 3D Face Reconstruction With Deep Neural Networks, Pengfei Dou, Shishir K. Shah, Ioannis A. Kakadiaris
- 46. DUST: Dual Union of Spatio-Temporal Subspaces for Monocular Multiple Object 3D Reconstruction, Antonio Agudo, Francesc Moreno-Noguer

#### Analyzing Humans in Images

- 47. Finding Tiny Faces, Peiyun Hu, Deva Ramanan
- 48. Dynamic Facial Analysis: From Bayesian Filtering to Recurrent Neural Network, Jinwei Gu, Xiaodong Yang, Shalini De Mello, Jan Kautz
- 49. Deep Temporal Linear Encoding Networks, Ali Diba, Vivek Sharma, Luc Van Gool
- 50. Joint Registration and Representation Learning for Unconstrained Face Identification, Munawar Hayat, Salman H. Khan, Naoufel Werghi, Roland Goecke
- 51. 3D Human Pose Estimation From a Single Image via Distance Matrix Regression, *Francesc Moreno-Noguer*
- 52. One-Shot Metric Learning for Person Re-Identification, Slawomir Bąk, Peter Carr
- 53. Generalized Rank Pooling for Activity Recognition, Anoop Cherian, Basura Fernando, Mehrtash Harandi, Stephen Gould
- 54. Deep Representation Learning for Human Motion Prediction and Classification, Judith Bütepage, Michael J. Black, Danica Kragic, Hedvig Kjellström
- 55. Interspecies Knowledge Transfer for Facial Keypoint Detection, Maheen Rashid, Xiuye Gu, Yong Jae Lee
- 56. Recurrent Convolutional Neural Networks for Continuous Sign Language Recognition by Staged Optimization, Runpeng Cui, Hu Liu, Changshui Zhang

#### **Applications**

57. Modeling Sub-Event Dynamics in First-Person Action Recognition, Hasan F. M. Zaki, Faisal Shafait, Ajmal Mian

#### Computational Photography

- 58. Turning an Urban Scene Video Into a Cinemagraph, Hang Yan, Yebin Liu, Yasutaka Furukawa
- 59. Light Field Reconstruction Using Deep Convolutional Network on EPI, *Gaochang Wu*, *Mandan Zhao*, *Liangyong Wang*, *Qionghai Dai*, *Tianyou Chai*, *Yebin Liu*

#### Image Motion & Tracking

 FlowNet 2.0: Evolution of Optical Flow Estimation With Deep Networks, Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, Thomas Brox

#### Low- & Mid-Level Vision

- Attention-Aware Face Hallucination via Deep Reinforcement Learning, Qingxing Cao, Liang Lin, Yukai Shi, Xiaodan Liang, Guanbin Li
- 62. Simple Does It: Weakly Supervised Instance and Semantic Segmentation, Anna Khoreva, Rodrigo Benenson, Jan Hosang, Matthias Hein, Bernt Schiele
- 63. Anti-Glare: Tightly Constrained Optimization for Eyeglass Reflection Removal, *Tushar Sandhan, Jin Young Choi*
- 64. Deep Joint Rain Detection and Removal From a Single Image, Wenhan Yang, Robby T. Tan, Jiashi Feng, Jiaying Liu, Zongming Guo, Shuicheng Yan
- 65. Radiometric Calibration From Faces in Images, Chen Li, Stephen Lin, Kun Zhou, Katsushi Ikeuchi
- 66. Webly Supervised Semantic Segmentation, Bin Jin, Maria V. Ortiz Segovia, Sabine Süsstrunk
- 67. Removing Rain From Single Images via a Deep Detail Network, Xueyang Fu, Jiabin Huang, Delu Zeng, Yue Huang, Xinghao Ding, John Paisley
- 68. Deep Crisp Boundaries, Yupei Wang, Xin Zhao, Kaiqi Huang
- 69. Coarse-To-Fine Segmentation With Shape-Tailored Continuum Scale Spaces, *Naeemullah Khan, Byung-Woo Hong, Anthony Yezzi, Ganesh Sundaramoorthi*
- Large Kernel Matters Improve Semantic Segmentation by Global Convolutional Network, Chao Peng, Xiangyu Zhang, Gang Yu, Guiming Luo, Jian Sun

## Program

- 71. Single Image Reflection Suppression, Nikolaos Arvanitopoulos, Radhakrishna Achanta, Sabine Süsstrunk
- 72. CASENet: Deep Category-Aware Semantic Edge Detection, Zhiding Yu, Chen Feng, Ming-Yu Liu, Srikumar Ramalingam
- 73. Reflectance Adaptive Filtering Improves Intrinsic Image Estimation, *Thomas Nestmeyer, Peter V. Gehler*

#### Machine Learning

- 74. Conditional Similarity Networks, Andreas Veit, Serge Belongie, Theofanis Karaletsos
- 75. Spatially Adaptive Computation Time for Residual Networks, Michael Figurnov, Maxwell D. Collins, Yukun Zhu, Li Zhang, Jonathan Huang, Dmitry Vetrov, Ruslan Salakhutdinov
- 76. Xception: Deep Learning With Depthwise Separable Convolutions, *François Chollet*
- 77. Feedback Networks, Amir R. Zamir, Te-Lin Wu, Lin Sun, William B. Shen, Bertram E. Shi, Jitendra Malik, Silvio Savarese
- 78. Online Summarization via Submodular and Convex Optimization, Ehsan Elhamifar, M. Clara De Paolis Kaluza
- 79. Deep MANTA: A Coarse-To-Fine Many-Task Network for Joint 2D and 3D Vehicle Analysis From Monocular Image, Florian Chabot, Mohamed Chaouch, Jaonary Rabarisoa, Céline Teulière, Thierry Chateau
- 80. Improving Pairwise Ranking for Multi-Label Image Classification, *Yuncheng Li*, *Yale Song, Jiebo Luo*
- 81. Active Convolution: Learning the Shape of Convolution for Image Classification, *Yunho Jeon, Junmo Kim*
- 82. Linking Image and Text With 2-Way Nets, Aviv Eisenschtat, Lior Wolf
- 83. Stacked Generative Adversarial Networks, Xun Huang, Yixuan Li, Omid Poursaeed, John Hopcroft, Serge Belongie
- 84. Image Splicing Detection via Camera Response Function Analysis, Can Chen, Scott McCloskey, Jingyi Yu
- 85. Building a Regular Decision Boundary With Deep Networks, *Edouard Oyallon*
- 86. More Is Less: A More Complicated Network With Less Inference Complexity, *Xuanyi Dong, Junshi Huang, Yi Yang, Shuicheng Yan*

- 87. Joint Graph Decomposition and Node Labeling: Problem, Algorithms, Applications, Evgeny Levinkov, Jonas Uhrig, Siyu Tang, Mohamed Omran, Eldar Insafutdinov, Alexander Kirillov, Carsten Rother, Thomas Brox, Bernt Schiele, Bjoern Andres
- 88. Scale-Aware Face Detection, Zekun Hao, Yu Liu, Hongwei Qin, Junjie Yan, Xiu Li, Xiaolin Hu
- 89. Deep Unsupervised Similarity Learning Using Partially Ordered Sets, *Miguel A. Bautista, Artsiom Sanakoyeu, Björn Ommer*
- 90. Generative Hierarchical Learning of Sparse FRAME Models, Jianwen Xie, Yifei Xu, Erik Nijkamp, Ying Nian Wu, Song-Chun Zhu

#### **Object Recognition & Scene Understanding**

- 91. Generating Holistic 3D Scene Abstractions for Text-Based Image Retrieval, Ang Li, Jin Sun, Joe Yue-Hei Ng, Ruichi Yu, Vlad I. Morariu, Larry S. Davis
- 92. Perceptual Generative Adversarial Networks for Small Object Detection, Jianan Li, Xiaodan Liang, Yunchao Wei, Tingfa Xu, Jiashi Feng, Shuicheng Yan
- 93. Emotion Recognition in Context, *Ronak Kosti, Jose M. Alvarez, Adria Recasens, Agata Lapedriza*
- 94. Deep Learning of Human Visual Sensitivity in Image Quality Assessment Framework, *Jongyoo Kim, Sanghoon Lee*
- 95. Dense Captioning With Joint Inference and Visual Context, *Linjie Yang, Kevin Tang, Jianchao Yang, Li-Jia Li*
- 96. CLEVR: A Diagnostic Dataset for Compositional Language and Elementary Visual Reasoning, Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C. Lawrence Zitnick, Ross Girshick
- 97. Cross-View Image Matching for Geo-Localization in Urban Environments, *Yicong Tian, Chen Chen, Mubarak Shah*
- Matrix Tri-Factorization With Manifold Regularizations for Zero-Shot Learning, Xing Xu, Fumin Shen, Yang Yang, Dongxiang Zhang, Heng Tao Shen, Jingkuan Song
- 99. Self-Supervised Learning of Visual Features Through Embedding Images Into Text Topic Spaces, *Lluis Gomez*, Yash Patel, Marçal Rusiñol, Dimosthenis Karatzas, C. V. Jawahar

- 100. Learning Spatial Regularization With Image-Level Supervisions for Multi-Label Image Classification, Feng Zhu, Hongsheng Li, Wanli Ouyang, Nenghai Yu, Xiaogang Wang
- 101. Semantically Consistent Regularization for Zero-Shot Recognition, *Pedro Morgado, Nuno Vasconcelos*
- 102. Can Walking and Measuring Along Chord Bunches Better Describe Leaf Shapes? Bin Wang, Yongsheng Gao, Changming Sun, Michael Blumenstein, John La Salle

#### Video Analytics

- 103. Self-Learning Scene-Specific Pedestrian Detectors Using a Progressive Latent Model, Qixiang Ye, Tianliang Zhang, Wei Ke, Qiang Qiu, Jie Chen, Guillermo Sapiro, Baochang Zhang
- 104. Predictive-Corrective Networks for Action Detection, Achal Dave, Olga Russakovsky, Deva Ramanan
- 105. Budget-Aware Deep Semantic Video Segmentation, Behrooz Mahasseni, Sinisa Todorovic, Alan Fern
- 106. Unified Embedding and Metric Learning for Zero-Exemplar Event Detection, Noureldien Hussein, Efstratios Gavves, Arnold W.M. Smeulders
- 107. Spatiotemporal Pyramid Network for Video Action Recognition, Yunbo Wang, Mingsheng Long, Jianmin Wang, Philip S. Yu
- 108.ER3: A Unified Framework for Event Retrieval, Recognition and Recounting, Zhanning Gao, Gang Hua, Dongqing Zhang, Nebojsa Jojic, Le Wang, Jianru Xue, Nanning Zheng
- 109. FusionSeg: Learning to Combine Motion and Appearance for Fully Automatic Segmentation of Generic Objects in Videos, Suyog Dutt Jain, Bo Xiong, Kristen Grauman
- 110. Query-Focused Video Summarization: Dataset, Evaluation, and a Memory Network Based Approach, Aidean Sharghi, Jacob S. Laurel, Boqing Gong
- 111. Flexible Spatio-Temporal Networks for Video Prediction, Chaochao Lu, Michael Hirsch, Bernhard Schölkopf
- 112. Temporal Action Co-Segmentation in 3D Motion Capture Data and Videos, Konstantinos Papoutsakis, Costas Panagiotakis, Antonis A. Argyros

### 1500-1700 Demos (Kamehameha I)

• Same as Saturday morning Demos (see pg. 9)

### 1500-1700 Exhibits (Kamehameha I)

• Same as Saturday morning Exhibits (see pg. 10)


### 1715-1830 Plenary Session (Kamehameha III)

• Keynote Talk: Extracting Social Meaning From Language, Dan Jurafsky (Stanford Univ.)

Abstract: I describe research in our lab on computationally extracting social meaning from language, meaning that takes into account social relationships between people. I describe our study of interactions between police and community members in traffic stops recorded in bodyworn camera footage. We automatically measure the quality of the interaction from language, study the role of race in the interaction, and draw suggestions for going forward in this fraught area. In another we computationally model the language of scientific papers together with the network formed by scientists and their research areas to better understand scientific innovation, how it progresses, and the role of interdisciplinarity. I discuss implications for the history of science and specifically of artificial intelligence. Both studies highlight the importance of social context and social models for interpreting the latent meanings behind the words we use.

### 1840–2040 PAMI Technical Committee Meeting (Kamehameha III)

#### Notes:



## Program

# Sunday, July 23

## 0730-1700 Registration (Main Lobby)

### 0730-0830 Breakfast (Kamehameha II)

### 0830–1000 Session 2-1A: Machine Learning 2 (Kamehameha III)

#### Papers in this session are also in Poster Session P2-1.

Chairs: Lorenzo Torresani (Dartmouth) Alexander Toshev (Google)

#### 0830 Spotlights (S2-1A)

#### Format (4 min. for presentation; no questions)

- [0830] Dual Attention Networks for Multimodal Reasoning and Matching, Hyeonseob Nam, Jung-Woo Ha, Jeonghee Kim
- [0834] DESIRE: Distant Future Prediction in Dynamic Scenes With Interacting Agents, Namhoon Lee, Wongun Choi, Paul Vernaza, Christopher B. Choy, Philip H. S. Torr, Manmohan Chandraker
- 3. [0838] Interpretable Structure-Evolving LSTM, Xiaodan Liang, Liang Lin, Xiaohui Shen, Jiashi Feng, Shuicheng Yan, Eric P. Xing
- [0842] ShapeOdds: Variational Bayesian Learning of Generative Shape Models, Shireen Elhabian, Ross Whitaker
- [0846] Fast Video Classification via Adaptive Cascading of Deep Models, Haichen Shen, Seungyeop Han, Matthai Philipose, Arvind Krishnamurthy
- 6. **[0850]** Deep Metric Learning via Facility Location, *Hyun Oh* Song, Stefanie Jegelka, Vivek Rathod, Kevin Murphy
- 7. **[0854]** Semi-Supervised Deep Learning for Monocular Depth Map Prediction, Yevhen Kuznietsov, Jörg Stückler, Bastian Leibe
- [0858] Weakly Supervised Semantic Segmentation Using Web-Crawled Videos, Seunghoon Hong, Donghun Yeo, Suha Kwak, Honglak Lee, Bohyung Han

#### 0903 Orals (O2-1A)

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

- [0903] Making Deep Neural Networks Robust to Label Noise: A Loss Correction Approach, Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, Lizhen Qu
- 10. [0917] Learning From Simulated and Unsupervised Images Through Adversarial Training, Ashish Shrivastava, Tomas Pfister, Oncel Tuzel, Joshua Susskind, Wenda Wang, Russell Webb
- [0931] Inverse Compositional Spatial Transformer Networks, Chen-Hsuan Lin, Simon Lucey
- 12. [0945] Densely Connected Convolutional Networks, Gao Huang, Zhuang Liu, Laurens van der Maaten, Kilian Q. Weinberger

## 0830-1000 Session 2-1B: Computational Photography

(Kalākaua Ballroom A-B)

#### Papers in this session are also in Poster Session P2-1.

Chairs: Imari Sato (National Inst. of Informatics) Eli Shechtman (Adobe Research)

#### o830 Spotlights (S2-1B)

#### Format (4 min. for presentation; no questions)

- 13. **[0830]** Video Frame Interpolation via Adaptive Convolution, *Simon Niklaus, Long Mai, Feng Liu*
- 14. [0834] FastMask: Segment Multi-Scale Object Candidates in One Shot, *Hexiang Hu, Shiyi Lan, Yuning Jiang, Zhimin Cao, Fei Sha*
- 15. **[0838]** Reconstructing Transient Images From Single-Photon Sensors, *Matthew O'Toole, Felix Heide, David B. Lindell, Kai Zang, Steven Diamond, Gordon Wetzstein*
- 16. **[0842]** Deep Sketch Hashing: Fast Free-Hand Sketch-Based Image Retrieval, *Li Liu, Fumin Shen, Yuming Shen, Xianglong Liu, Ling Shao*
- 17. **[0846]** DeshadowNet: A Multi-Context Embedding Deep Network for Shadow Removal, *Liangqiong Qu, Jiandong Tian, Shengfeng He, Yandong Tang, Rynson W. H. Lau*
- 18. [0850] Illuminant-Camera Communication to Observe Moving Objects Under Strong External Light by Spread Spectrum Modulation, Ryusuke Sagawa, Yutaka Satoh

## Program

- 19. [0854] Photorealistic Facial Texture Inference Using Deep Neural Networks, Shunsuke Saito, Lingyu Wei, Liwen Hu, Koki Nagano, Hao Li
- 20. [0858] The Geometry of First-Returning Photons for Non-Line-Of-Sight Imaging, Chia-Yin Tsai, Kiriakos N. Kutulakos, Srinivasa G. Narasimhan, Aswin C. Sankaranarayanan

#### 0903 Orals (O2-1B)

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

- 21. [0903] Unrolling the Shutter: CNN to Correct Motion Distortions, *Vijay Rengarajan, Yogesh Balaji, A. N. Rajagopalan*
- 22. **[0917]** Light Field Blind Motion Deblurring, *Pratul P. Srinivasan, Ren Ng, Ravi Ramamoorthi*
- 23. [0931] Computational Imaging on the Electric Grid, Mark Sheinin, Yoav Y. Schechner, Kiriakos N. Kutulakos
- 24. **[0945]** Deep Outdoor Illumination Estimation, Yannick Hold-Geoffroy, Kalyan Sunkavalli, Sunil Hadap, Emiliano Gambaretto, Jean-François Lalonde

#### 0830–1000 Session 2-1C: 3D Vision 2 (Kalākaua Ballroom C)

#### Papers in this session are also in Poster Session P2-1.

Chairs: Yasutaka Furukawa (Washington Univ. St. Louis) Jingyi Yu (Univ. of Delaware)

#### o830 Spotlights (S2-1C)

Format (4 min. for presentation; no questions)

- 25. [0830] Efficient Solvers for Minimal Problems by Syzygy-Based Reduction, Viktor Larsson, Kalle Åström, Magnus Oskarsson
- 26. **[0834]** HSfM: Hybrid Structure-from-Motion, Hainan Cui, Xiang Gao, Shuhan Shen, Zhanyi Hu
- 27. [0838] Efficient Global Point Cloud Alignment Using Bayesian Nonparametric Mixtures, Julian Straub, Trevor Campbell, Jonathan P. How, John W. Fisher III
- 28. [0842] A New Rank Constraint on Multi-View Fundamental Matrices, and Its Application to Camera Location Recovery, Soumyadip Sengupta, Tal Amir, Meirav Galun, Tom Goldstein, David W. Jacobs, Amit Singer, Ronen Basri
- 29. [0846] IM2CAD, Hamid Izadinia, Qi Shan, Steven M. Seitz

- 30. [0850] ScanNet: Richly-Annotated 3D Reconstructions of Indoor Scenes, Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, Matthias Nießner
- 31. **[0854]** Noise Robust Depth From Focus Using a Ring Difference Filter, Jaeheung Surh, Hae-Gon Jeon, Yunwon Park, Sunghoon Im, Hyowon Ha, In So Kweon
- 32. [0858] Group-Wise Point-Set Registration Based on Rényi's Second Order Entropy, Luis G. Sanchez Giraldo, Erion Hasanbelliu, Murali Rao, Jose C. Principe

#### 0903 Orals (O2-1C)

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

- 33. **[ogo3]** A Point Set Generation Network for 3D Object Reconstruction From a Single Image, *Haoqiang Fan, Hao Su, Leonidas J. Guibas*
- 34. [0917] 3D Point Cloud Registration for Localization Using a Deep Neural Network Auto-Encoder, Gil Elbaz, Tamar Avraham, Anath Fischer
- 35. **[0931]** Flight Dynamics-Based Recovery of a UAV Trajectory Using Ground Cameras, Artem Rozantsev, Sudipta N. Sinha, Debadeepta Dey, Pascal Fua
- 36. [0945] DSAC Differentiable RANSAC for Camera Localization, Eric Brachmann, Alexander Krull, Sebastian Nowozin, Jamie Shotton, Frank Michel, Stefan Gumhold, Carsten Rother

## 1000-1045 Break (Kamehameha II)

#### 1000-1200 Poster Session P2-1 (Kamehameha I)

#### 3D Computer Vision

- Scalable Surface Reconstruction From Point Clouds With Extreme Scale and Density Diversity, Christian Mostegel, Rudolf Prettenthaler, Friedrich Fraundorfer, Horst Bischof
- 38. Synthesizing 3D Shapes via Modeling Multi-View Depth Maps and Silhouettes With Deep Generative Networks, Amir Arsalan Soltani, Haibin Huang, Jiajun Wu, Tejas D. Kulkarni, Joshua B. Tenenbaum
- 39. General Models for Rational Cameras and the Case of Two-Slit Projections, Matthew Trager, Bernd Sturmfels, John Canny, Martial Hebert, Jean Ponce

## Program

- 40. Accurate Depth and Normal Maps From Occlusion-Aware Focal Stack Symmetry, Michael Strecke, Anna Alperovich, Bastian Goldluecke
- 41. A Multi-View Stereo Benchmark With High-Resolution Images and Multi-Camera Videos, Thomas Schöps, Johannes L. Schönberger, Silvano Galliani, Torsten Sattler, Konrad Schindler, Marc Pollefeys, Andreas Geiger
- 42. Non-Contact Full Field Vibration Measurement Based on Phase-Shifting, *Hiroyuki Kayaba*, *Yuji Kokumai*
- 43. A Minimal Solution for Two-View Focal-Length Estimation Using Two Affine Correspondences, Daniel Barath, Tekla Toth, Levente Hajder
- 44. PoseAgent: Budget-Constrained 6D Object Pose Estimation via Reinforcement Learning, Alexander Krull, Eric Brachmann, Sebastian Nowozin, Frank Michel, Jamie Shotton, Carsten Rother
- 45. An Efficient Background Term for 3D Reconstruction and Tracking With Smooth Surface Models, Mariano Jaimez, Thomas J. Cashman, Andrew Fitzgibbon, Javier Gonzalez-Jimenez, Daniel Cremers

#### Analyzing Humans in Images

- 46. Reliable Crowdsourcing and Deep Locality-Preserving Learning for Expression Recognition in the Wild, Shan Li, Weihong Deng, JunPing Du
- 47. Procedural Generation of Videos to Train Deep Action Recognition Networks, *César Roberto de Souza, Adrien Gaidon, Yohann Cabon, Antonio Manuel López*
- 48. BigHand2.2M Benchmark: Hand Pose Dataset and State of the Art Analysis, Shanxin Yuan, Qi Ye, Björn Stenger, Siddhant Jain, Tae-Kyun Kim
- 49. DenseReg: Fully Convolutional Dense Shape Regression In-The-Wild, Rıza Alp Güler, George Trigeorgis, Epameinondas Antonakos, Patrick Snape, Stefanos Zafeiriou, Iasonas Kokkinos
- Adaptive Class Preserving Representation for Image Classification, Jian-Xun Mi, Qiankun Fu, Weisheng Li

#### **Applications**

51. Generalized Semantic Preserving Hashing for N-Label Cross-Modal Retrieval, *Devraj Mandal, Kunal N. Chaudhury, Soma Biswas* 

- 52. EAST: An Efficient and Accurate Scene Text Detector, Xinyu Zhou, Cong Yao, He Wen, Yuzhi Wang, Shuchang Zhou, Weiran He, Jiajun Liang
- 53. VidLoc: A Deep Spatio-Temporal Model for 6-DoF Video-Clip Relocalization, *Ronald Clark, Sen Wang, Andrew Markham, Niki Trigoni, Hongkai Wen*

#### Biomedical Image/Video Analysis

54. Improving RANSAC-Based Segmentation Through CNN Encapsulation, *Dustin Morley, Hassan Foroosh* 

#### Computational Photography

- 55. Position Tracking for Virtual Reality Using Commodity WiFi, Manikanta Kotaru, Sachin Katti
- 56. Designing Illuminant Spectral Power Distributions for Surface Classification, *Henryk Blasinski, Joyce Farrell, Brian Wandell*
- 57. One-Shot Hyperspectral Imaging Using Faced Reflectors, Tsuyoshi Takatani, Takahito Aoto, Yasuhiro Mukaigawa

#### Image Motion & Tracking

- 58. Direct Photometric Alignment by Mesh Deformation, Kaimo Lin, Nianjuan Jiang, Shuaicheng Liu, Loong-Fah Cheong, Minh Do, Jiangbo Lu
- 59. CNN-Based Patch Matching for Optical Flow With Thresholded Hinge Embedding Loss, Christian Bailer, Kiran Varanasi, Didier Stricker
- 60. Optical Flow Estimation Using a Spatial Pyramid Network, Anurag Ranjan, Michael J. Black
- 61. Deep Network Flow for Multi-Object Tracking, Samuel Schulter, Paul Vernaza, Wongun Choi, Manmohan Chandraker

#### Low- & Mid-Level Vision

- 62. Material Classification Using Frequency- and Depth-Dependent Time-Of-Flight Distortion, *Kenichiro Tanaka*, Yasuhiro Mukaigawa, Takuya Funatomi, Hiroyuki Kubo, Yasuyuki Matsushita, Yasushi Yagi
- 63. Benchmarking Denoising Algorithms With Real Photographs, *Tobias Plötz, Stefan Roth*
- 64. A Unified Approach of Multi-Scale Deep and Hand-Crafted Features for Defocus Estimation, *Jinsun Park*, Yu-*Wing Tai, Donghyeon Cho, In So Kweon*

## Program

- StyleBank: An Explicit Representation for Neural Image Style Transfer, Dongdong Chen, Lu Yuan, Jing Liao, Nenghai Yu, Gang Hua
- 66. Specular Highlight Removal in Facial Images, Chen Li, Stephen Lin, Kun Zhou, Kαtsushi Ikeuchi
- 67. Image Super-Resolution via Deep Recursive Residual Network, *Ying Tai, Jian Yang, Xiaoming Liu*
- 68. Deep Image Harmonization, Yi-Hsuan Tsai, Xiaohui Shen, Zhe Lin, Kalyan Sunkavalli, Xin Lu, Ming-Hsuan Yang
- 69. Learning Deep CNN Denoiser Prior for Image Restoration, Kai Zhang, Wangmeng Zuo, Shuhang Gu, Lei Zhang
- 70. A Novel Tensor-Based Video Rain Streaks Removal Approach via Utilizing Discriminatively Intrinsic Priors, Tai-Xiang Jiang, Ting-Zhu Huang, Xi-Le Zhao, Liang-Jian Deng, Yao Wang
- 71. GMS: Grid-based Motion Statistics for Fast, Ultra-Robust Feature Correspondence, JiaWang Bian, Wen-Yan Lin, Yasuyuki Matsushita, Sai-Kit Yeung, Tan-Dat Nguyen, Ming-Ming Cheng
- 72. Video Desnowing and Deraining Based on Matrix Decomposition, *Weihong Ren, Jiandong Tian, Zhi Han, Antoni Chan, Yandong Tang*
- 73. Real-Time Video Super-Resolution With Spatio-Temporal Networks and Motion Compensation, *Jose Caballero, Christian Ledig, Andrew Aitken, Alejandro Acosta, Johannes Totz, Zehan Wang, Wenzhe Shi*
- 74. Deep Watershed Transform for Instance Segmentation, Min Bai, Raquel Urtasun
- 75. AnchorNet: A Weakly Supervised Network to Learn Geometry-Sensitive Features for Semantic Matching, David Novotny, Diane Larlus, Andrea Vedaldi
- 76. Learning Diverse Image Colorization, Aditya Deshpande, Jiajun Lu, Mao-Chuang Yeh, Min Jin Chong, David Forsyth
- 77. Awesome Typography: Statistics-Based Text Effects Transfer, Shuai Yang, Jiaying Liu, Zhouhui Lian, Zongming Guo

#### Machine Learning

- Unsupervised Video Summarization With Adversarial LSTM Networks, Behrooz Mahasseni, Michael Lam, Sinisa Todorovic
- 79. Deep TEN: Texture Encoding Network, Hang Zhang, Jia Xue, Kristin Dana

- 80. Order-Preserving Wasserstein Distance for Sequence Matching, *Bing Su, Gang Hua*
- Attend in Groups: A Weakly-Supervised Deep Learning Framework for Learning From Web Data, Bohan Zhuang, Lingqiao Liu, Yao Li, Chunhua Shen, Ian Reid
- Hierarchical Multimodal Metric Learning for Multimodal Classification, Heng Zhang, Vishal M. Patel, Rama Chellappa
- Efficient Linear Programming for Dense CRFs, Thalaiyasingam Ajanthan, Alban Desmaison, Rudy Bunel, Mathieu Salzmann, Philip H. S. Torr, M. Pawan Kumar
- 84. Variational Autoencoded Regression: High Dimensional Regression of Visual Data on Complex Manifold, YoungJoon Yoo, Sangdoo Yun, Hyung Jin Chang, Yiannis Demiris, Jin Young Choi
- Learning Random-Walk Label Propagation for Weakly-Supervised Semantic Segmentation, Paul Vernaza, Manmohan Chandraker
- 86. Adversarial Discriminative Domain Adaptation, Eric Tzeng, Judy Hoffman, Kate Saenko, Trevor Darrell
- 87. Low-Rank-Sparse Subspace Representation for Robust Regression, Yongqiang Zhang, Daming Shi, Junbin Gao, Dansong Cheng

#### **Object Recognition & Scene Understanding**

- Generating the Future With Adversarial Transformers, Carl Vondrick, Antonio Torralba
- 89. Semantic Amodal Segmentation, Yan Zhu, Yuandong Tian, Dimitris Metaxas, Piotr Dollár
- 90. Learning a Deep Embedding Model for Zero-Shot Learning, *Li Zhang, Tao Xiang, Shaogang Gong*
- 91. BIND: Binary Integrated Net Descriptors for Texture-Less Object Recognition, Jacob Chan, Jimmy Addison Lee, Qian Kemao
- 92. Growing a Brain: Fine-Tuning by Increasing Model Capacity, Yu-Xiong Wang, Deva Ramanan, Martial Hebert
- 93. A-Fast-RCNN: Hard Positive Generation via Adversary for Object Detection, Xiaolong Wang, Abhinav Shrivastava, Abhinav Gupta
- 94. Multiple Instance Detection Network With Online Instance Classifier Refinement, *Peng Tang, Xinggang Wang, Xiang Bai, Wenyu Liu*

- 95. Kernel Pooling for Convolutional Neural Networks, Yin Cui, Feng Zhou, Jiang Wang, Xiao Liu, Yuanqing Lin, Serge Belongie
- 96. Learning Cross-Modal Embeddings for Cooking Recipes and Food Images, Amaia Salvador, Nicholas Hynes, Yusuf Aytar, Javier Marin, Ferda Ofli, Ingmar Weber, Antonio Torralba
- 97. Zero-Shot Learning the Good, the Bad and the Ugly, Yongqin Xian, Bernt Schiele, Zeynep Akata
- 98. DeepNav: Learning to Navigate Large Cities, Samarth Brahmbhatt, James Hays
- 99. Scene Graph Generation by Iterative Message Passing, Danfei Xu, Yuke Zhu, Christopher B. Choy, Li Fei-Fei
- 100. Visual Translation Embedding Network for Visual Relation Detection, Hanwang Zhang, Zawlin Kyaw, Shih-Fu Chang, Tat-Seng Chua
- 101. Unsupervised Part Learning for Visual Recognition, Ronan Sicre, Yannis Avrithis, Ewa Kijak, Frédéric Jurie
- 102. Comprehension-Guided Referring Expressions, *Ruotian Luo, Gregory Shakhnarovich*
- 103. Top-Down Visual Saliency Guided by Captions, Vasili Ramanishka, Abir Das, Jianming Zhang, Kate Saenko

#### Theory

104. Grassmannian Manifold Optimization Assisted Sparse Spectral Clustering, *Qiong Wang, Junbin Gao, Hong Li* 

#### Video Analytics

- 105. Video Propagation Networks, Varun Jampani, Raghudeep Gadde, Peter V. Gehler
- 106.ActionVLAD: Learning Spatio-Temporal Aggregation for Action Classification, *Rohit Girdhar, Deva Ramanan, Abhinav Gupta, Josef Sivic, Bryan Russell*
- 107. SCC: Semantic Context Cascade for Efficient Action Detection, Fabian Caba Heilbron, Wayner Barrios, Victor Escorcia, Bernard Ghanem
- 108. Hierarchical Boundary-Aware Neural Encoder for Video Captioning, *Lorenzo Baraldi, Costantino Grana, Rita Cucchiara*
- 109. HOPE: Hierarchical Object Prototype Encoding for Efficient Object Instance Search in Videos, *Tan Yu, Yuwei Wu, Junsong Yuan*

- 110. Spatio-Temporal Vector of Locally Max Pooled Features for Action Recognition in Videos, *Ionut Cosmin Duta*, *Bogdan Ionescu, Kiyoharu Aizawa, Nicu Sebe*
- 111. Temporal Action Localization by Structured Maximal Sums, Zehuan Yuan, Jonathan C. Stroud, Tong Lu, Jia Deng
- 112. Predicting Salient Face in Multiple-Face Videos, Yufan Liu, Songyang Zhang, Mai Xu, Xuming He

### 1000-1200 Demos (Kamehameha I)

- FlowNet 2.0: Evolution of Optical Flow Estimation With Deep Networks, Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, Thomas Brox (Univ. of Freiburg)
- A Low Power, Fully Event-Based Gesture Recognition System, Arnon Amir, Brian Taba, David Berg, Timothy Melano, Jeffrey McKinstry, Carmelo Di Nolfo, Tapan Nayak, Alexander Andreopoulos, Guillaume Garreau, Marcela Mendoza, Jeff Kusnitz, Michael Debole, Steve Esser, Tobi Delbruck, Myron Flickner, Dharmendra Modha (IBM Almaden Research Center)
- Fast Moving Objects Detection, Recognition, Description, De-Blurring, Denys Rozumnyi, Aleš Hrabalík, Jan Kotera, Filip Šroubek, Jiří Matas (Czech Technical Univ. in Prague, Czech Academy of Sciences)
- OpenPose: A Real-Time Multi-Person Keypoint Detection Library, Gines Hidalgo, Zhe Cao, Tomas Simon, Shih-En Wei, Hanbyul Joo, Yaser Sheikh (Carnegie Mellon Univ.)

## 1000-1200 Exhibits (Kamehameha I)

• Same as Saturday morning Exhibits (see pg. 10)

## 1200-1300 Lunch (Kamehameha II)

## Program

#### 1300-1430 Session 2-2A: Object Recognition & Scene Understanding 1 (Kamehameha III)

#### Papers in this session are also in Poster Session P2-2.

Chairs: Vittorio Ferrari (Univ. of Edinburgh) Cha Zhang (Microsoft Research)

#### 1300 Spotlights (S2-2A)

#### Format (4 min. for presentation; no questions)

- [1300] Graph-Structured Representations for Visual Question Answering, Damien Teney, Lingqiao Liu, Anton van den Hengel
- [1304] Knowing When to Look: Adaptive Attention via a Visual Sentinel for Image Captioning, *Jiasen Lu, Caiming Xiong, Devi Parikh, Richard Socher*
- [1308] Learned Contextual Feature Reweighting for Image Geo-Localization, Hyo Jin Kim, Enrique Dunn, Jan-Michael Frahm
- 4. [1312] End-To-End Concept Word Detection for Video Captioning, Retrieval, and Question Answering, Youngjae Yu, Hyungjin Ko, Jongwook Choi, Gunhee Kim
- 5. [1316] Deep Cross-Modal Hashing, Qing-Yuan Jiang, Wu-Jun Li
- [1320] Unambiguous Text Localization and Retrieval for Cluttered Scenes, Xuejian Rong, Chucai Yi, Yingli Tian
- 7. [1324] Bayesian Supervised Hashing, Zihao Hu, Junxuan Chen, Hongtao Lu, Tongzhen Zhang
- [1328] Speed/Accuracy Trade-Offs for Modern Convolutional Object Detectors, Jonathan Huang, Vivek Rathod, Chen Sun, Menglong Zhu, Anoop Korattikara, Alireza Fathi, Ian Fischer, Zbigniew Wojna, Yang Song, Sergio Guadarrama, Kevin Murphy

#### 1333 Orals (O2-2A)

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

- [1333] Detecting Visual Relationships With Deep Relational Networks, Bo Dai, Yuqi Zhang, Dahua Lin
- [1347] Full-Resolution Residual Networks for Semantic Segmentation in Street Scenes, Tobias Pohlen, Alexander Hermans, Markus Mathias, Bastian Leibe
- 11. [1401] Network Dissection: Quantifying Interpretability of Deep Visual Representations, David Bau, Bolei Zhou, Aditya Khosla, Aude Oliva, Antonio Torralba

12. [1415] AGA: Attribute-Guided Augmentation, Mandar Dixit, Roland Kwitt, Marc Niethammer, Nuno Vasconcelos

### 1300–1430 Session 2-2B: Analyzing Humans 2 (Kalākaua Ballroom A-B)

#### Papers in this session are also in Poster Session P2-2.

Chairs: Jingdong Wang (Microsoft Research Asia) Daphna Weinshall (Hebrew Univ. of Jerusalem)

#### 1300 Spotlights (S2-2B)

#### Format (4 min. for presentation; no questions)

- 13. [1300] A Hierarchical Approach for Generating Descriptive Image Paragraphs, Jonathan Krause, Justin Johnson, Ranjay Krishna, Li Fei-Fei
- 14. **[1304]** Person Re-Identification in the Wild, *Liang Zheng,* Hengheng Zhang, Shaoyan Sun, Manmohan Chandraker, Yi Yang, Qi Tian
- 15. **[1308]** Scalable Person Re-Identification on Supervised Smoothed Manifold, *Song Bai, Xiang Bai, Qi Tian*
- 16. [1312] Binge Watching: Scaling Affordance Learning From Sitcoms, Xiaolong Wang, Rohit Girdhar, Abhinav Gupta
- [1316] Joint Detection and Identification Feature Learning for Person Search, Tong Xiao, Shuang Li, Bochao Wang, Liang Lin, Xiaogang Wang
- 18. [1320] Synthesizing Normalized Faces From Facial Identity Features, Forrester Cole, David Belanger, Dilip Krishnan, Aaron Sarna, Inbar Mosseri, William T. Freeman
- [1324] Consistent-Aware Deep Learning for Person Re-Identification in a Camera Network, Ji Lin, Liangliang Ren, Jiwen Lu, Jianjiang Feng, Jie Zhou
- 20. [1328] Level Playing Field for Million Scale Face Recognition, Aaron Nech, Ira Kemelmacher-Shlizerman

#### 1333 Orals (O2-2B)

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

- 21. [1333] Re-Sign: Re-Aligned End-To-End Sequence Modelling With Deep Recurrent CNN-HMMs, Oscar Koller, Sepehr Zargaran, Hermann Ney
- 22. **[1347]** Social Scene Understanding: End-To-End Multi-Person Action Localization and Collective Activity Recognition, *Timur Bagautdinov, Alexandre Alahi, François Fleuret, Pascal Fua, Silvio Savarese*

## Program

- 23. [1401] Detangling People: Individuating Multiple Close People and Their Body Parts via Region Assembly, *Hao Jiang, Kristen Grauman*
- 24. [1415] Lip Reading Sentences in the Wild, Joon Son Chung, Andrew Senior, Oriol Vinyals, Andrew Zisserman

#### 1300–1430 Session 2-2C: Applications (Kalākaua Ballroom C)

#### Papers in this session are also in Poster Session P2-2.

Chairs: Zicheng Liu (Microsoft Research) S. Kevin Zhou (Siemens Research)

#### 1300 Spotlights (S2-2C)

#### Format (4 min. for presentation; no questions)

- 25. **[1300]** Deep Matching Prior Network: Toward Tighter Multi-Oriented Text Detection, *Yuliang Liu, Lianwen Jin*
- 26. [1304] ChestX-ray8: Hospital-Scale Chest X-Ray Database and Benchmarks on Weakly-Supervised Classification and Localization of Common Thorax Diseases, Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Mohammadhadi Bagheri, Ronald M. Summers
- 27. [1308] Attentional Push: A Deep Convolutional Network for Augmenting Image Salience With Shared Attention Modeling in Social Scenes, Siavash Gorji, James J. Clark
- 28. [1312] Detecting Oriented Text in Natural Images by Linking Segments, Baoguang Shi, Xiang Bai, Serge Belongie
- 29. [1316] Learning Video Object Segmentation From Static Images, Federico Perazzi, Anna Khoreva, Rodrigo Benenson, Bernt Schiele, Alexander Sorkine-Hornung
- 30. **[1320]** Seeing Invisible Poses: Estimating 3D Body Pose From Egocentric Video, *Hao Jiang, Kristen Grauman*
- 31. [1324] Plug & Play Generative Networks: Conditional Iterative Generation of Images in Latent Space, Anh Nguyen, Jeff Clune, Yoshua Bengio, Alexey Dosovitskiy, Jason Yosinski
- 32. **[1328]** A Joint Speaker-Listener-Reinforcer Model for Referring Expressions, *Licheng Yu, Hao Tan, Mohit Bansal, Tamara L. Berg*

#### 1333 Orals (O2-2C)

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

- 33. [1333] End-To-End Learning of Driving Models From Large-Scale Video Datasets, *Huazhe Xu*, *Yang Gao*, *Fisher Yu*, *Trevor Darrell*
- 34. [1347] Deep Future Gaze: Gaze Anticipation on Egocentric Videos Using Adversarial Networks, Mengmi Zhang, Keng Teck Ma, Joo Hwee Lim, Qi Zhao, Jiashi Feng
- 35. **[1401]** MDNet: A Semantically and Visually Interpretable Medical Image Diagnosis Network, *Zizhao Zhang, Yuanpu Xie, Fuyong Xing, Mason McGough, Lin Yang*

## 1430-1515 Break (Kamehameha II)

## 1430-1630 Poster Session P2-2 (Kamehameha I)

#### 3D Computer Vision

- 36. Surface Motion Capture Transfer With Gaussian Process Regression, Adnane Boukhayma, Jean-Sébastien Franco, Edmond Boyer
- 37. Visual-Inertial-Semantic Scene Representation for 3D Object Detection, *Jingming Dong, Xiaohan Fei, Stefano Soatto*
- 38. Template-Based Monocular 3D Recovery of Elastic Shapes Using Lagrangian Multipliers, Nazim Haouchine, Stephane Cotin
- 39. Learning Category-Specific 3D Shape Models From Weakly Labeled 2D Images, Dingwen Zhang, Junwei Han, Yang Yang, Dong Huang
- 40. Simultaneous Geometric and Radiometric Calibration of a Projector-Camera Pair, Marjan Shahpaski, Luis Ricardo Sapaico, Gaspard Chevassus, Sabine Süsstrunk
- A Clever Elimination Strategy for Efficient Minimal Solvers, Zuzana Kukelova, Joe Kileel, Bernd Sturmfels, Tomas Pajdla
- 42. Learning Barycentric Representations of 3D Shapes for Sketch-Based 3D Shape Retrieval, Jin Xie, Guoxian Dai, Fan Zhu, Yi Fang
- 43. Geodesic Distance Descriptors, Gil Shamai, Ron Kimmel

#### Analyzing Humans in Images

- 44. Modeling Temporal Dynamics and Spatial Configurations of Actions Using Two-Stream Recurrent Neural Networks, Hongsong Wang, Liang Wang
- 45. Forecasting Human Dynamics From Static Images, Yu-Wei Chao, Jimei Yang, Brian Price, Scott Cohen, Jia Deng
- 46. Re-Ranking Person Re-Identification With *k*-Reciprocal Encoding, *Zhun Zhong, Liang Zheng, Donglin Cao, Shaozi Li*
- 47. Deep Sequential Context Networks for Action Prediction, Yu Kong, Zhiqiang Tao, Yun Fu
- 48. Global Context-Aware Attention LSTM Networks for 3D Action Recognition, *Jun Liu, Gang Wang, Ping Hu, Ling-Yu Duan, Alex C. Kot*
- 49. Dynamic Attention-Controlled Cascaded Shape Regression Exploiting Training Data Augmentation and Fuzzy-Set Sample Weighting, Zhen-Hua Feng, Josef Kittler, William Christmas, Patrik Huber, Xiao-Jun Wu
- 50. A Deep Regression Architecture With Two-Stage Re-Initialization for High Performance Facial Landmark Detection, Jiangjing Lv, Xiaohu Shao, Junliang Xing, Cheng Cheng, Xi Zhou
- 51. Multiple People Tracking by Lifted Multicut and Person Re-Identification, *Siyu Tang, Mykhaylo Andriluka, Bjoern Andres, Bernt Schiele*
- 52. Towards Accurate Multi-Person Pose Estimation in the Wild, George Papandreou, Tyler Zhu, Nori Kanazawa, Alexander Toshev, Jonathan Tompson, Chris Bregler, Kevin Murphy

#### **Applications**

- 53. Towards a Quality Metric for Dense Light Fields, Vamsi Kiran Adhikarla, Marek Vinkler, Denis Sumin, Rafał K. Mantiuk, Karol Myszkowski, Hans-Peter Seidel, Piotr Didyk
- 54. Controlling Perceptual Factors in Neural Style Transfer, Leon A. Gatys, Alexander S. Ecker, Matthias Bethge, Aaron Hertzmann, Eli Shechtman

#### Biomedical Image/Video Analysis

- 55. Joint Sequence Learning and Cross-Modality Convolution for 3D Biomedical Segmentation, Kuan-Lun Tseng, Yen-Liang Lin, Winston Hsu, Chung-Yang Huang
- 56. LSTM Self-Supervision for Detailed Behavior Analysis, Biagio Brattoli, Uta Büchler, Anna-Sophia Wahl, Martin E. Schwab, Björn Ommer

#### Computational Photography

57. A Wide-Field-Of-View Monocentric Light Field Camera, Donald G. Dansereau, Glenn Schuster, Joseph Ford, Gordon Wetzstein

#### Image Motion & Tracking

- 58. S2F: Slow-To-Fast Interpolator Flow, Yanchao Yang, Stefano Soatto
- 59. CLKN: Cascaded Lucas-Kanade Networks for Image Alignment, Che-Han Chang, Chun-Nan Chou, Edward Y. Chang
- 60. Multi-Object Tracking With Quadruplet Convolutional Neural Networks, *Jeany Son, Mooyeol Baek, Minsu Cho, Bohyung Han*

#### Low- & Mid-Level Vision

- Learning to Detect Salient Objects With Image-Level Supervision, Lijun Wang, Huchuan Lu, Yifan Wang, Mengyang Feng, Dong Wang, Baocai Yin, Xiang Ruan
- 62. From Motion Blur to Motion Flow: A Deep Learning Solution for Removing Heterogeneous Motion Blur, Dong Gong, Jie Yang, Lingqiao Liu, Yanning Zhang, Ian Reid, Chunhua Shen, Anton van den Hengel, Qinfeng Shi
- 63. Co-Occurrence Filter, Roy J. Jevnisek, Shai Avidan
- 64. Fractal Dimension Invariant Filtering and Its CNN-Based Implementation, *Hongteng Xu*, *Junchi Yan*, *Nils Persson*, *Weiyao Lin*, *Hongyuan Zha*
- 65. Noise-Blind Image Deblurring, *Meiguang Jin, Stefan Roth, Paolo Favaro*
- 66. Simultaneous Visual Data Completion and Denoising Based on Tensor Rank and Total Variation Minimization and Its Primal-Dual Splitting Algorithm, *Tatsuya Yokota*, *Hidekata Hontani*
- 67. HPatches: A Benchmark and Evaluation of Handcrafted and Learned Local Descriptors, Vassileios Balntas, Karel Lenc, Andrea Vedaldi, Krystian Mikolajczyk
- Hyperspectral Image Super-Resolution via Non-Local Sparse Tensor Factorization, *Renwei Dian, Leyuan Fang,* Shutao Li
- 69. Reflection Removal Using Low-Rank Matrix Completion, Byeong-Ju Han, Jae-Young Sim
- 70. Object Co-Skeletonization With Co-Segmentation, Koteswar Rao Jerripothula, Jianfei Cai, Jiangbo Lu, Junsong Yuan

## Program

#### Machine Learning

- Mining Object Parts From CNNs via Active Question-Answering, Quanshi Zhang, Ruiming Cao, Ying Nian Wu, Song-Chun Zhu
- 72. PolyNet: A Pursuit of Structural Diversity in Very Deep Networks, Xingcheng Zhang, Zhizhong Li, Chen Change Loy, Dahua Lin
- 73. The VQA-Machine: Learning How to Use Existing Vision Algorithms to Answer New Questions, *Peng Wang, Qi Wu, Chunhua Shen, Anton van den Hengel*
- 74. Joint Discriminative Bayesian Dictionary and Classifier Learning, Naveed Akhtar, Ajmal Mian, Fatih Porikli
- 75. Quad-Networks: Unsupervised Learning to Rank for Interest Point Detection, Nikolay Savinov, Akihito Seki, Ľubor Ladický, Torsten Sattler, Marc Pollefeys
- 76. Outlier-Robust Tensor PCA, Pan Zhou, Jiashi Feng
- 77. Learning Adaptive Receptive Fields for Deep Image Parsing Network, *Zhen Wei, Yao Sun, Jinqiao Wang, Hanjiang Lai, Si Liu*
- 78. Learning an Invariant Hilbert Space for Domain Adaptation, Samitha Herath, Mehrtash Harandi, Fatih Porikli
- 79. Fixed-Point Factorized Networks, *Peisong Wang, Jian* Cheng
- Discriminative Optimization: Theory and Applications to Point Cloud Registration, Jayakorn Vongkulbhisal, Fernando De la Torre, João P. Costeira
- 81. Online Asymmetric Similarity Learning for Cross-Modal Retrieval, Yiling Wu, Shuhui Wang, Qingming Huang
- Improving Training of Deep Neural Networks via Singular Value Bounding, Kui Jia, Dacheng Tao, Shenghua Gao, Xiangmin Xu
- S3Pool: Pooling With Stochastic Spatial Sampling, Shuangfei Zhai, Hui Wu, Abhishek Kumar, Yu Cheng, Yongxi Lu, Zhongfei Zhang, Rogerio Feris
- 84. Sports Field Localization via Deep Structured Models, Namdar Homayounfar, Sanja Fidler, Raquel Urtasun
- Noisy Softmax: Improving the Generalization Ability of DCNN via Postponing the Early Softmax Saturation, Binghui Chen, Weihong Deng, Junping Du

- Switching Convolutional Neural Network for Crowd Counting, Deepak Babu Sam, Shiv Surya, R. Venkatesh Babu
- Network Sketching: Exploiting Binary Structure in Deep CNNs, Yiwen Guo, Anbang Yao, Hao Zhao, Yurong Chen
- 88. Multi-Task Clustering of Human Actions by Sharing Information, *Xiaoqiang Yan, Shizhe Hu, Yangdong Ye*
- 89. Soft-Margin Mixture of Regressions, Dong Huang, Longfei Han, Fernando De la Torre
- 90. Multigrid Neural Architectures, Tsung-Wei Ke, Michael Maire, Stella X. Yu
- 91. High-Resolution Image Inpainting Using Multi-Scale Neural Patch Synthesis, *Chao Yang, Xin Lu, Zhe Lin, Eli* Shechtman, Oliver Wang, Hao Li
- 92. Deep Quantization: Encoding Convolutional Activations With Deep Generative Model, *Zhaofan Qiu*, *Ting Yao*, *Tao Mei*
- 93. DOPE: Distributed Optimization for Pairwise Energies, Jose Dolz, Ismail Ben Ayed, Christian Desrosiers
- 94. Improved Texture Networks: Maximizing Quality and Diversity in Feed-Forward Stylization and Texture Synthesis, Dmitry Ulyanov, Andrea Vedaldi, Victor Lempitsky

#### **Object Recognition & Scene Understanding**

- 95. Polyhedral Conic Classifiers for Visual Object Detection and Classification, *Hakan Cevikalp, Bill Triggs*
- 96. Incremental Kernel Null Space Discriminant Analysis for Novelty Detection, Juncheng Liu, Zhouhui Lian, Yi Wang, Jianguo Xiao
- 97. Predicting Ground-Level Scene Layout From Aerial Imagery, *Menghua Zhai, Zachary Bessinger, Scott Workman, Nathan Jacobs*
- 98. Deep Feature Flow for Video Recognition, *Xizhou Zhu, Yuwen Xiong, Jifeng Dai, Lu Yuan, Yichen Wei*
- 99. Object-Aware Dense Semantic Correspondence, Fan Yang, Xin Li, Hong Cheng, Jianping Li, Leiting Chen
- 100. Semantic Regularisation for Recurrent Image Annotation, Feng Liu, Tao Xiang, Timothy M. Hospedales, Wankou Yang, Changyin Sun
- 101. Video2Shop: Exact Matching Clothes in Videos to Online Shopping Images, Zhi-Qi Cheng, Xiao Wu, Yang Liu, Xian-Sheng Hua

- 102.Fast-At: Fast Automatic Thumbnail Generation Using Deep Neural Networks, Seyed A. Esmaeili, Bharat Singh, Larry S. Davis
- 103. Multi-Level Attention Networks for Visual Question Answering, *Dongfei Yu*, *Jianlong Fu*, *Tao Mei*, *Yong Rui*
- 104. Generating Descriptions With Grounded and Co-Referenced People, Anna Rohrbach, Marcus Rohrbach, Siyu Tang, Seong Joon Oh, Bernt Schiele
- 105. Straight to Shapes: Real-Time Detection of Encoded Shapes, Saumya Jetley, Michael Sapienza, Stuart Golodetz, Philip H. S. Torr
- 106. Simultaneous Feature Aggregating and Hashing for Large-Scale Image Search, *Thanh-Toan Do, Dang-Khoa Le Tan, Trung T. Pham, Ngai-Man Cheung*
- 107. Improving Facial Attribute Prediction Using Semantic Segmentation, Mahdi M. Kalayeh, Boqing Gong, Mubarak Shah

#### Video Analytics

- 108. Learning Cross-Modal Deep Representations for Robust Pedestrian Detection, Dan Xu, Wanli Ouyang, Elisa Ricci, Xiaogang Wang, Nicu Sebe
- 109. Spatio-Temporal Self-Organizing Map Deep Network for Dynamic Object Detection From Videos, Yang Du, Chunfeng Yuan, Bing Li, Weiming Hu, Stephen Maybank
- 110. CERN: Confidence-Energy Recurrent Network for Group Activity Recognition, *Tianmin Shu, Sinisa Todorovic, Song-Chun Zhu*
- 111. Understanding Traffic Density From Large-Scale Web Camera Data, Shanghang Zhang, Guanhang Wu, João P. Costeira, José M. F. Moura
- 112. Collaborative Summarization of Topic-Related Videos, Rameswar Panda, Amit K. Roy-Chowdhury

#### 1430-1630 Demos (Kamehameha I)

Same as Sunday morning Demos (see pg. 21)

#### 1430-1630 Exhibits (Kamehameha I)

• Same as Saturday morning Exhibits (see pg. 10)

## 1645-1800 Plenary Session (Kamehameha III)

- Keynote Talk: Commercializing Computer Vision: Success Stories and Lessons Learned, Harry Shum (Microsoft Research)
- Abstract: It is an exciting time for all of us computer vision researchers and practitioners. We have seen an unprecedented growth in the conversion of years of progress into marketable technologies. Microsoft has long been committed to developing new computer vision technologies, making them available to developers, and incorporating them into many products. In this talk, I will first briefly review 25 years of computer vision research at Microsoft Research (MSR), highlighting MSR's contributions to the vision community and emphasizing the importance of longterm commitment to funding successful industrial research labs. I will also describe some of our latest research work in computational photography, image understanding, and vision and language before detailing our commercialization successes. In particular, I will share our experiences in developing three products: Microsoft Pix, HoloLens, and Cognitive Services, which leverage computer vision systems and technologies in different ways.

Pix is an Al-powered camera app that makes taking great pictures easy and fun: "point, shoot, perfect!" It has incorporated technologies from more than a dozen CVPR, ICCV, and SIGGRAPH papers from MSR. HoloLens is the first commercially available mixed reality system in the market. Cognitive Services allow you to build useful Albased apps using just a few lines of code, across different devices and platforms. I will show IRIS, which is an interactive visual learning service for developers to create image recognition applications. I will also show the latest cool demos using HoloLens, including real-time environment understanding. There are challenges in accelerating the cycle from research to product, and I will discuss the lessons learned in productizing Pix, HoloLens, and Cognitive Services.

### **1900–2200 Luau Reception** (The Museum Quad, Hale Koa Hotel beachfront)

**Directions:** On the beachfront between the Hale Koa Hotel (2055 Kalia Rd.) and the Fort DeRussy US Army Museum (2161 Kalia Rd.).

# Monday, July 24

### 0730-1200 Registration (Main Lobby)

### 0730-0830 Breakfast (Kamehameha II)

0830-1000 Session 3-1A: Machine Learning 3 (Kamehameha III)

#### Papers in this session are also in Poster Session P3-1.

Chairs: Simon Lucey (Carnegie Mellon Univ.) Gabriel Brostow (Univ. College London)

#### o830 Spotlights (S3-1A)

#### Format (4 min. for presentation; no questions)

- 1. [0830] Local Binary Convolutional Neural Networks, Felix Juefei-Xu, Vishnu Naresh Boddeti, Marios Savvides
- [0834] Deep Self-Taught Learning for Weakly Supervised Object Localization, Zequn Jie, Yunchao Wei, Xiaojie Jin, Jiashi Feng, Wei Liu
- 3. [0838] Multi-Modal Mean-Fields via Cardinality-Based Clamping, Pierre Baqué, François Fleuret, Pascal Fua
- 4. **[0842]** Probabilistic Temporal Subspace Clustering, Behnam Gholami, Vladimir Pavlovic
- [0846] Provable Self-Representation Based Outlier Detection in a Union of Subspaces, Chong You, Daniel P. Robinson, René Vidal
- 6. **[0850]** Latent Multi-View Subspace Clustering, Changqing Zhang, Qinghua Hu, Huazhu Fu, Pengfei Zhu, Xiaochun Cao
- 7. [0854] Learning to Extract Semantic Structure From Documents Using Multimodal Fully Convolutional Neural Networks, Xiao Yang, Ersin Yumer, Paul Asente, Mike Kraley, Daniel Kifer, C. Lee Giles
- 8. **[0858]** Age Progression/Regression by Conditional Adversarial Autoencoder, *Zhifei Zhang, Yang Song, Hairong Qi*

#### 0903 Orals (O3-1A)

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

9. **[0903]** Compact Matrix Factorization With Dependent Subspaces, *Viktor Larsson, Carl Olsson* 

- 10. **[0917]** FFTLasso: Large-Scale LASSO in the Fourier Domain, *Adel Bibi, Hani Itani, Bernard Ghanem*
- 11. [0931] On the Global Geometry of Sphere-Constrained Sparse Blind Deconvolution, Yuqian Zhang, Yenson Lau, Han-wen Kuo, Sky Cheung, Abhay Pasupathy, John Wright
- 12. **[0945]** Global Optimality in Neural Network Training, Benjamin D. Haeffele, René Vidal

#### 0830-1000 Session 3-1B: Object Recognition & Scene Understanding 2 (Kalākaua Ballroom)

#### Papers in this session are also in Poster Session P3-1.

Chairs: Gang Hua (Microsoft Research Asia) Hamed Pirsiavash (Univ. of Maryland Baltimore County)

#### o830 Spotlights (S3-1B)

Format (4 min. for presentation; no questions)

- 13. [0830] What Is and What Is Not a Salient Object? Learning Salient Object Detector by Ensembling Linear Exemplar Regressors, Changqun Xia, Jia Li, Xiaowu Chen, Anlin Zheng, Yu Zhang
- 14. **[0834]** Deep Variation-Structured Reinforcement Learning for Visual Relationship and Attribute Detection, *Xiaodan Liang, Lisa Lee, Eric P. Xing*
- 15. [0838] Modeling Relationships in Referential Expressions With Compositional Modular Networks, Ronghang Hu, Marcus Rohrbach, Jacob Andreas, Trevor Darrell, Kate Saenko
- 16. **[0842]** Counting Everyday Objects in Everyday Scenes, Prithvijit Chattopadhyay, Ramakrishna Vedantam, Ramprasaath R. Selvaraju, Dhruv Batra, Devi Parikh
- 17. **[0846]** Fully Convolutional Instance-Aware Semantic Segmentation, Yi Li, Haozhi Qi, Jifeng Dai, Xiangyang Ji, Yichen Wei
- 18. **[0850]** Semantic Autoencoder for Zero-Shot Learning, *Elyor Kodirov, Tao Xiang, Shaogang Gong*
- 19. [0854] CityPersons: A Diverse Dataset for Pedestrian Detection, Shanshan Zhang, Rodrigo Benenson, Bernt Schiele
- 20. [0858] GuessWhat?! Visual Object Discovery Through Multi-Modal Dialogue, Harm de Vries, Florian Strub, Sarath Chandar, Olivier Pietquin, Hugo Larochelle, Aaron Courville

## Program

### 0903 Orals (O3-1B)

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

- 21. [0903] Look Closer to See Better: Recurrent Attention Convolutional Neural Network for Fine-Grained Image Recognition, *Jianlong Fu*, *Heliang Zheng*, *Tao Mei*
- 22. [0917] Annotating Object Instances With a Polygon-RNN, Lluís Castrejón, Kaustav Kundu, Raquel Urtasun, Sanja Fidler
- 23. [0931] Connecting Look and Feel: Associating the Visual and Tactile Properties of Physical Materials, *Wenzhen Yuan, Shaoxiong Wang, Siyuan Dong, Edward Adelson*
- 24. [0945] Deep Learning Human Mind for Automated Visual Classification, Concetto Spampinato, Simone Palazzo, Isaak Kavasidis, Daniela Giordano, Nasim Souly, Mubarak Shah

## 1000-1045 Break (Kamehameha II)

## 1000-1200 Poster Session P3-1 (Kamehameha I)

### 3D Computer Vision

- 25. Self-Calibration-Based Approach to Critical Motion Sequences of Rolling-Shutter Structure From Motion, *Eisuke Ito, Takayuki Okatani*
- 26. Semi-Calibrated Near Field Photometric Stereo, Fotios Logothetis, Roberto Mecca, Roberto Cipolla
- 27. Semantic Multi-View Stereo: Jointly Estimating Objects and Voxels, *Ali Osman Ulusoy, Michael J. Black, Andreas Geiger*
- 28. Learning to Predict Stereo Reliability Enforcing Local Consistency of Confidence Maps, Matteo Poggi, Stefano Mattoccia
- 29. The Misty Three Point Algorithm for Relative Pose, *Tobias Palmér, Kalle Åström, Jan-Michael Frahm*
- 30. The Surfacing of Multiview 3D Drawings via Lofting and Occlusion Reasoning, Anil Usumezbas, Ricardo Fabbri, Benjamin B. Kimia
- 31. A New Representation of Skeleton Sequences for 3D Action Recognition, Qiuhong Ke, Mohammed Bennamoun, Senjian An, Ferdous Sohel, Farid Boussaid
- 32. A General Framework for Curve and Surface Comparison and Registration With Oriented Varifolds, *Irène Kaltenmark, Benjamin Charlier, Nicolas Charon*

- 33. Learning to Align Semantic Segmentation and 2.5D Maps for Geolocalization, Anil Armagan, Martin Hirzer, Peter M. Roth, Vincent Lepetit
- 34. A Generative Model for Depth-Based Robust 3D Facial Pose Tracking, Lu Sheng, Jianfei Cai, Tat-Jen Cham, Vladimir Pavlovic, King Ngi Ngan
- 35. Fast 3D Reconstruction of Faces With Glasses, Fabio Maninchedda, Martin R. Oswald, Marc Pollefeys
- 36. An Efficient Algebraic Solution to the Perspective-Three-Point Problem, *Tong Ke, Stergios I. Roumeliotis*

#### Analyzing Humans in Images

- 37. Learning From Synthetic Humans, Gül Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J. Black, Ivan Laptev, Cordelia Schmid
- 38. Forecasting Interactive Dynamics of Pedestrians With Fictitious Play, Wei-Chiu Ma, De-An Huang, Namhoon Lee, Kris M. Kitani
- 39. Hand Keypoint Detection in Single Images Using Multiview Bootstrapping, Tomas Simon, Hanbyul Joo, Iain Matthews, Yaser Sheikh
- 40. PoseTrack: Joint Multi-Person Pose Estimation and Tracking, Umar Iqbal, Anton Milan, Juergen Gall
- 41. Expecting the Unexpected: Training Detectors for Unusual Pedestrians With Adversarial Imposters, Shiyu Huang, Deva Ramanan
- 42. On Human Motion Prediction Using Recurrent Neural Networks, Julieta Martinez, Michael J. Black, Javier Romero
- 43. Learning and Refining of Privileged Information-Based RNNs for Action Recognition From Depth Sequences, *Zhiyuan Shi, Tae-Kyun Kim*
- 44. Quality Aware Network for Set to Set Recognition, Yu Liu, Junjie Yan, Wanli Ouyang
- 45. Unite the People: Closing the Loop Between 3D and 2D Human Representations, *Christoph Lassner, Javier Romero, Martin Kiefel, Federica Bogo, Michael J. Black, Peter V. Gehler*
- 46. Deep Multitask Architecture for Integrated 2D and 3D Human Sensing, Alin-Ionut Popa, Mihai Zanfir, Cristian Sminchisescu
- 47. Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset, *João Carreira, Andrew Zisserman*

## Program

## **Applications**

- 48. Identifying First-Person Camera Wearers in Third-Person Videos, Chenyou Fan, Jangwon Lee, Mingze Xu, Krishna Kumar Singh, Yong Jae Lee, David J. Crandall, Michael S. Ryoo
- 49. Learning to Rank Retargeted Images, Yang Chen, Yong-Jin Liu, Yu-Kun Lai

### Biomedical Image/Video Analysis

- 50. Parsing Images of Overlapping Organisms With Deep Singling-Out Networks, Victor Yurchenko, Victor Lempitsky
- 51. Fine-Tuning Convolutional Neural Networks for Biomedical Image Analysis: Actively and Incrementally, Zongwei Zhou, Jae Shin, Lei Zhang, Suryakanth Gurudu, Michael Gotway, Jianming Liang

### Computational Photography

- 52. Depth From Defocus in the Wild, Huixuan Tang, Scott Cohen, Brian Price, Stephen Schiller, Kiriakos N. Kutulakos
- 53. Matting and Depth Recovery of Thin Structures Using a Focal Stack, *Chao Liu, Srinivasa G. Narasimhan, Artur W. Dubrawski*

## Image Motion & Tracking

- 54. Robust Interpolation of Correspondences for Large Displacement Optical Flow, *Yinlin Hu, Yunsong Li, Rui Song*
- 55. Large Margin Object Tracking With Circulant Feature Maps, *Mengmeng Wang*, *Yong Liu*, *Zeyi Huang*
- 56. Minimum Delay Moving Object Detection, *Dong Lao*, *Ganesh Sundaramoorthi*
- 57. Multi-Task Correlation Particle Filter for Robust Object Tracking, *Tianzhu Zhang, Changsheng Xu, Ming-Hsuan* Yang
- 58. Attentional Correlation Filter Network for Adaptive Visual Tracking, Jongwon Choi, Hyung Jin Chang, Sangdoo Yun, Tobias Fischer, Yiannis Demiris, Jin Young Choi
- 59. The World of Fast Moving Objects, Denys Rozumnyi, Jan Kotera, Filip Šroubek, Lukáš Novotný, Jiří Matas
- 60. Discriminative Correlation Filter With Channel and Spatial Reliability, Alan Lukežič, Tomáš Vojíř, Luka Čehovin Zajc, Jiří Matas, Matej Kristan

## Low- & Mid-Level Vision

- 61. Learning Deep Binary Descriptor With Multi-Quantization, Yueqi Duan, Jiwen Lu, Ziwei Wang, Jianjiang Feng, Jie Zhou
- 62. One-To-Many Network for Visually Pleasing Compression Artifacts Reduction, *Jun Guo, Hongyang Chao*
- 63. Gated Feedback Refinement Network for Dense Image Labeling, *Md Amirul Islam, Mrigank Rochan, Neil D. B. Bruce, Yang Wang*
- 64. BRISKS: Binary Features for Spherical Images on a Geodesic Grid, *Hao Guan*, *William A. P. Smith*
- 65. Superpixels and Polygons Using Simple Non-Iterative Clustering, Radhakrishna Achanta, Sabine Süsstrunk
- 66. Hardware-Efficient Guided Image Filtering for Multi-Label Problem, Longquan Dai, Mengke Yuan, Zechao Li, Xiaopeng Zhang, Jinhui Tang
- 67. Alternating Direction Graph Matching, D. Khuê Lê-Huu, Nikos Paragios
- Learning Discriminative and Transformation Covariant Local Feature Detectors, Xu Zhang, Felix X. Yu, Svebor Karaman, Shih-Fu Chang

### Machine Learning

- 69. Correlational Gaussian Processes for Cross-Domain Visual Recognition, *Chengjiang Long, Gang Hua*
- 70. DeLiGAN : Generative Adversarial Networks for Diverse and Limited Data, Swaminathan Gurumurthy, Ravi Kiran Sarvadevabhatla, R. Venkatesh Babu
- A Dual Ascent Framework for Lagrangean Decomposition of Combinatorial Problems, Paul Swoboda, Jan Kuske, Bogdan Savchynskyy
- 72. Oriented Response Networks, Yanzhao Zhou, Qixiang Ye, Qiang Qiu, Jianbin Jiao
- 73. Missing Modalities Imputation via Cascaded Residual Autoencoder, Luan Tran, Xiaoming Liu, Jiayu Zhou, Rong Jin
- 74. Efficient Optimization for Hierarchically-structured Interacting Segments (HINTS), Hossam Isack, Olga Veksler, Ipek Oguz, Milan Sonka, Yuri Boykov
- 75. A Message Passing Algorithm for the Minimum Cost Multicut Problem, *Paul Swoboda, Bjoern Andres*

## Program

- 76. End-To-End Representation Learning for Correlation Filter Based Tracking, Jack Valmadre, Luca Bertinetto, João Henriques, Andrea Vedaldi, Philip H. S. Torr
- 77. Filter Flow Made Practical: Massively Parallel and Lock-Free, Sathya N. Ravi, Yunyang Xiong, Lopamudra Mukherjee, Vikas Singh
- 78. Online Graph Completion: Multivariate Signal Recovery in Computer Vision, Won Hwa Kim, Mona Jalal, Seongjae Hwang, Sterling C. Johnson, Vikas Singh
- 79. Point to Set Similarity Based Deep Feature Learning for Person Re-Identification, Sanping Zhou, Jinjun Wang, Jiayun Wang, Yihong Gong, Nanning Zheng
- Exploiting Saliency for Object Segmentation From Image Level Labels, Seong Joon Oh, Rodrigo Benenson, Anna Khoreva, Zeynep Akata, Mario Fritz, Bernt Schiele
- Consensus Maximization With Linear Matrix Inequality Constraints, Pablo Speciale, Danda Pani Paudel, Martin R. Oswald, Till Kroeger, Luc Van Gool, Marc Pollefeys
- 82. Physically-Based Rendering for Indoor Scene Understanding Using Convolutional Neural Networks, Yinda Zhang, Shuran Song, Ersin Yumer, Manolis Savva, Joon-Young Lee, Hailin Jin, Thomas Funkhouser
- 83. Deep Multimodal Representation Learning From Temporal Data, Xitong Yang, Palghat Ramesh, Radha Chitta, Sriganesh Madhvanath, Edgar A. Bernal, Jiebo Luo
- 84. All You Need Is Beyond a Good Init: Exploring Better Solution for Training Extremely Deep Convolutional Neural Networks With Orthonormality and Modulation, Di Xie, Jiang Xiong, Shiliang Pu
- Hard Mixtures of Experts for Large Scale Weakly Supervised Vision, Sam Gross, Marc'Aurelio Ranzato, Arthur Szlam
- 86. A Reinforcement Learning Approach to the View Planning Problem, Mustafa Devrim Kaba, Mustafa Gokhan Uzunbas, Ser Nam Lim
- 87. Zero-Shot Classification With Discriminative Semantic Representation Learning, *Meng Ye, Yuhong Guo*

#### **Object Recognition & Scene Understanding**

 Automatic Discovery, Association Estimation and Learning of Semantic Attributes for a Thousand Categories, Ziad Al-Halah, Rainer Stiefelhagen

- Scene Parsing Through ADE2oK Dataset, Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, Antonio Torralba
- 90. Weakly Supervised Cascaded Convolutional Networks, Ali Diba, Vivek Sharma, Ali Pazandeh, Hamed Pirsiavash, Luc Van Gool
- 91. Discretely Coding Semantic Rank Orders for Supervised Image Hashing, *Li Liu, Ling Shao, Fumin Shen, Mengyang* Yu
- 92. Joint Geometrical and Statistical Alignment for Visual Domain Adaptation, *Jing Zhang, Wanqing Li, Philip Ogunbona*
- 93. Weakly Supervised Dense Video Captioning, Zhiqiang Shen, Jianguo Li, Zhou Su, Minjun Li, Yurong Chen, Yu-Gang Jiang, Xiangyang Xue
- 94. RefineNet: Multi-Path Refinement Networks for High-Resolution Semantic Segmentation, *Guosheng Lin, Anton Milan, Chunhua Shen, Ian Reid*
- 95. Semantic Segmentation via Structured Patch Prediction, Context CRF and Guidance CRF, *Falong Shen, Rui Gan, Shuicheng Yan, Gang Zeng*
- 96. Person Search With Natural Language Description, Shuang Li, Tong Xiao, Hongsheng Li, Bolei Zhou, Dayu Yue, Xiaogang Wang
- 97. Weakly Supervised Affordance Detection, Johann Sawatzky, Abhilash Srikantha, Juergen Gall
- 98. Zero-Shot Recognition Using Dual Visual-Semantic Mapping Paths, Yanan Li, Donghui Wang, Huanhang Hu, Yuetan Lin, Yueting Zhuang
- 99. Neural Aggregation Network for Video Face Recognition, Jiaolong Yang, Peiran Ren, Dongqing Zhang, Dong Chen, Fang Wen, Hongdong Li, Gang Hua
- 100. Relationship Proposal Networks, Ji Zhang, Mohamed Elhoseiny, Scott Cohen, Walter Chang, Ahmed Elgammal
- 101. Learning Object Interactions and Descriptions for Semantic Image Segmentation, *Guangrun Wang, Ping Luo, Liang Lin, Xiaogang Wang*
- 102.RON: Reverse Connection With Objectness Prior Networks for Object Detection, *Tao Kong, Fuchun Sun, Anbang Yao, Huaping Liu, Ming Lu, Yurong Chen*

## Program

- 103. Weakly-Supervised Visual Grounding of Phrases With Linguistic Structures, Fanyi Xiao, Leonid Sigal, Yong Jae Lee
- 104. Incorporating Copying Mechanism in Image Captioning for Learning Novel Objects, *Ting Yao, Yingwei Pan, Yehao Li, Tao Mei*
- 105. Beyond Instance-Level Image Retrieval: Leveraging Captions to Learn a Global Visual Representation for Semantic Retrieval, *Albert Gordo, Diane Larlus*
- 106.MuCaLe-Net: Multi Categorical-Level Networks to Generate More Discriminating Features, Youssef Tamaazousti, Hervé Le Borgne, Céline Hudelot
- 107. Zero Shot Learning via Multi-Scale Manifold Regularization, *Shay Deutsch, Soheil Kolouri, Kyungnam Kim, Yuri Owechko, Stefano Soatto*

#### Theory

- 108. Deeply Supervised Salient Object Detection With Short Connections, Qibin Hou, Ming-Ming Cheng, Xiaowei Hu, Ali Borji, Zhuowen Tu, Philip H. S. Torr
- 109. A Matrix Splitting Method for Composite Function Minimization, Ganzhao Yuan, Wei-Shi Zheng, Bernard Ghanem

#### Video Analytics

- 110. One-Shot Video Object Segmentation, Sergi Caelles, Kevis-Kokitsi Maninis, Jordi Pont-Tuset, Laura Leal-Taixé, Daniel Cremers, Luc Van Gool
- 111. Fast Person Re-Identification via Cross-Camera Semantic Binary Transformation, Jiaxin Chen, Yunhong Wang, Jie Qin, Li Liu, Ling Shao
- 112. SPFTN: A Self-Paced Fine-Tuning Network for Segmenting Objects in Weakly Labelled Videos, *Dingwen Zhang, Le Yang, Deyu Meng, Dong Xu, Junwei Han*

## 1000-1200 Demos (Kamehameha I)

 VNect: Real-Time 3D Human Pose Estimation With a Single RGB Camera, Dushyant Mehta, Srinath Sridhar, Oleksandr Sotnychenko, Helge Rhodin, Franziska Mueller, Weipeng Xu, Dan Casas, Christian Theobalt (Max-Planck-Institut für Informatik, École Polytechnique Fédérale de Lausanne, Universidad Rey Juan Carlos)


- ShapeSearch: A Generic Search Engine for 3D Models, Images and Sketches, Flora Ponjou Tasse (Univ. of Cambridge/Selerio)
- Online Social Media Image Processing Using AIDR 2.0: Artificial Intelligence for Digital Response, *Firoj Alam, Muhammad Imran, Ferda Ofli (Qatar Computing Research Inst., Hamad Bin Khalifa Univ.)*

### 1000-1200 Exhibits (Kamehameha I)

• Same as Saturday morning Exhibits (see pg. 10)

### 1200 Lunch (on your own)

#### Notes:



## 1200-1400 Doctoral Consortium (Kamehameha

II) (by invitation only)

Supported by: National Science Foundation

Mataa

- Shervin Ardeshir Univ. of Central Florida ٠
- Xiao Chu CUHK
- Zhengming Ding Northeastern Univ. •
- Christoph Feichtenhofer Graz Univ. of Technology
- Je Hyeong Hong Univ. of Cambridge
- Ahmet Iscen INRIA/Renne
- Dinesh Jayaraman Univ. of Texas at Austin ٠
- Dinghuang Ji UNC Chapel Hill •
- Felix Juefei-Xu Carnegie Mellon Univ. ٠
- Mahdi Kalayeh Univ. of Central Florida ٠
- Kuldeep Kulkarni Arizona State Univ. ٠
- Ang Li Univ. of Maryland at College Park
- Dong Li Tsinghua Univ. ٠
- Yang Long Univ. of Sheffield ٠
- Ondrej Miksik Univ. of Oxford •
- Saeid Motiian West Virginia Univ.
- Tae-Hyun Oh KAIST ٠
- Eshed Ohn-Bar Univ. of California San Diego ٠
- Jinshan Pan Dalian Univ. of Technology
- Wenjie Pei Delft Univ. of Technology •
- Wenqi Ren Tianjin Univ. ٠
- Christos Sagonas Imperial College London
- Sergey Tulyakov Univ. of Trento ٠
- Subhashini Venugopalan Univ. of Texas at Austin ٠
- Yu-Xiong Wang Carnegie Mellon Univ.
- Scott Workman Univ. of Kentucky
- Zhongwen Xu Univ. of Technology Sydney ٠
- Qingan Yan Wuhan Univ.
- Fisher Yu Princeton Univ.
- Hang Zhang Rutgers Univ.

 		 	 		 _
 					_
 					_
 	 	 		 	_
 					 _
 					_

## Program

# Tuesday, July 25

## 0730-1700 Registration (Main Lobby)

## 0730-0830 Breakfast (Kamehameha II)

### 0830-1000 Session 4-1A: Machine Learning 4 (Kamehameha III)

### Papers in this session are also in Poster Session P4-1.

Chairs: Greg Shakhnarovich (TTI Chicago) Lior Wolf (Tel Aviv Univ.)

### 0830 Spotlights (S4-1A)

### Format (4 min. for presentation; no questions)

- 1. [0830] Hidden Layers in Perceptual Learning, Gad Cohen, Daphna Weinshall
- 2. [0834] Few-Shot Object Recognition From Machine-Labeled Web Images, *Zhongwen Xu, Linchao Zhu, Yi Yang*
- [0838] Hallucinating Very Low-Resolution Unaligned and Noisy Face Images by Transformative Discriminative Autoencoders, Xin Yu, Fatih Porikli
- 4. [0842] Are You Smarter Than a Sixth Grader? Textbook Question Answering for Multimodal Machine Comprehension, Aniruddha Kembhavi, Minjoon Seo, Dustin Schwenk, Jonghyun Choi, Ali Farhadi, Hannaneh Hajishirzi
- 5. **[0846]** Deep Hashing Network for Unsupervised Domain Adaptation, *Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, Sethuraman Panchanathan*
- 6. **[0850]** Generalized Deep Image to Image Regression, Venkataraman Santhanam, Vlad I. Morariu, Larry S. Davis
- 7. **[0854]** Deep Learning With Low Precision by Half-Wave Gaussian Quantization, *Zhaowei Cai, Xiaodong He, Jian Sun, Nuno Vasconcelos*
- 8. **[0858]** Creativity: Generating Diverse Questions Using Variational Autoencoders, *Unnat Jain, Ziyu Zhang, Alexander G. Schwing*

#### 0903 Orals (O4-1A)

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

- [0903] Geometric Deep Learning on Graphs and Manifolds Using Mixture Model CNNs, Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodolà, Jan Svoboda, Michael M. Bronstein
- [0917] Full Resolution Image Compression With Recurrent Neural Networks, George Toderici, Damien Vincent, Nick Johnston, Sung Jin Hwang, David Minnen, Joel Shor, Michele Covell
- 11. [0931] Neural Face Editing With Intrinsic Image Disentangling, Zhixin Shu, Ersin Yumer, Sunil Hadap, Kalyan Sunkavalli, Eli Shechtman, Dimitris Samaras
- 12. [0945] Ubernet: Training a Universal Convolutional Neural Network for Low-, Mid-, and High-Level Vision Using Diverse Datasets and Limited Memory, *Iasonas Kokkinos*

### 0830-1000 Session 4-1B: Analyzing Humans with 3D Vision (Kalākaua Ballroom)

### Papers in this session are also in Poster Session P4-1.

Chairs: Xilin Chen (Chinese Academy of Sciences) Matthew Turk (UC Santa Barbara)

### 0830 Spotlights (S4-1B)

### Format (4 min. for presentation; no questions)

- 13. [0830] 3D Face Morphable Models "In-The-Wild", James Booth, Epameinondas Antonakos, Stylianos Ploumpis, George Trigeorgis, Yannis Panagakis, Stefanos Zafeiriou
- 14. [0834] KillingFusion: Non-Rigid 3D Reconstruction Without Correspondences, Miroslava Slavcheva, Maximilian Baust, Daniel Cremers, Slobodan Ilic
- 15. **[0838]** Detailed, Accurate, Human Shape Estimation From Clothed 3D Scan Sequences, *Chao Zhang, Sergi Pujades, Michael J. Black, Gerard Pons-Moll*
- 16. **[0842]** POSEidon: Face-From-Depth for Driver Pose Estimation, *Guido Borghi, Marco Venturelli, Roberto Vezzani, Rita Cucchiara*
- 17. **[0846]** Human Shape From Silhouettes Using Generative HKS Descriptors and Cross-Modal Neural Networks, *Endri* Dibra, Himanshu Jain, Cengiz Öztireli, Remo Ziegler, Markus Gross

## Program

- Io850] Parametric T-Spline Face Morphable Model for Detailed Fitting in Shape Subspace, Weilong Peng, Zhiyong Feng, Chao Xu, Yong Su
- [0854] 3D Menagerie: Modeling the 3D Shape and Pose of Animals, Silvia Zuffi, Angjoo Kanazawa, David W. Jacobs, Michael J. Black
- 20. [0858] iCaRL: Incremental Classifier and Representation Learning, Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, Christoph H. Lampert

#### 0903 Orals (O4-1B)

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

- 21. [0903] Recurrent 3D Pose Sequence Machines, Mude Lin, Liang Lin, Xiaodan Liang, Keze Wang, Hui Cheng
- 22. [0917] Learning Detailed Face Reconstruction From a Single Image, *Elad Richardson, Matan Sela, Roy Or-El, Ron Kimmel*
- 23. [0931] Thin-Slicing Network: A Deep Structured Model for Pose Estimation in Videos, *Jie Song, Limin Wang, Luc Van Gool, Otmar Hilliges*
- 24. [0945] Dynamic FAUST: Registering Human Bodies in Motion, Federica Bogo, Javier Romero, Gerard Pons-Moll, Michael J. Black

## 1000-1045 Break (Kamehameha II)

## 1000-1200 Poster Session P4-1 (Kamehameha I)

#### 3D Computer Vision

- 25. Semantically Coherent Co-Segmentation and Reconstruction of Dynamic Scenes, *Armin Mustafa, Adrian Hilton*
- 26. On the Two-View Geometry of Unsynchronized Cameras, Cenek Albl, Zuzana Kukelova, Andrew Fitzgibbon, Jan Heller, Matej Smid, Tomas Pajdla
- 27. Using Locally Corresponding CAD Models for Dense 3D Reconstructions From a Single Image, *Chen Kong, Chen-Hsuan Lin, Simon Lucey*
- 28. Convex Global 3D Registration With Lagrangian Duality, Jesus Briales, Javier Gonzalez-Jimenez

- 29. DeMoN: Depth and Motion Network for Learning Monocular Stereo, Benjamin Ummenhofer, Huizhong Zhou, Jonas Uhrig, Nikolaus Mayer, Eddy Ilg, Alexey Dosovitskiy, Thomas Brox
- 3D Bounding Box Estimation Using Deep Learning and Geometry, Arsalan Mousavian, Dragomir Anguelov, John Flynn, Jana Košecká
- 31. A Dataset for Benchmarking Image-Based Localization, Xun Sun, Yuanfan Xie, Pei Luo, Liang Wang

#### Analyzing Humans in Images

- 32. Asynchronous Temporal Fields for Action Recognition, Gunnar A. Sigurdsson, Santosh Divvala, Ali Farhadi, Abhinav Gupta
- 33. Sequential Person Recognition in Photo Albums With a Recurrent Network, Yao Li, Guosheng Lin, Bohan Zhuang, Lingqiao Liu, Chunhua Shen, Anton van den Hengel
- 34. Multi-Context Attention for Human Pose Estimation, Xiao Chu, Wei Yang, Wanli Ouyang, Cheng Ma, Alan L. Yuille, Xiaogang Wang
- 35. 3D Convolutional Neural Networks for Efficient and Robust Hand Pose Estimation From Single Depth Images, Liuhao Ge, Hui Liang, Junsong Yuan, Daniel Thalmann
- 36. Lifting From the Deep: Convolutional 3D Pose Estimation From a Single Image, Denis Tome, Chris Russell, Lourdes Agapito
- 37. AdaScan: Adaptive Scan Pooling in Deep Convolutional Neural Networks for Human Action Recognition in Videos, Amlan Kar, Nishant Rai, Karan Sikka, Gaurav Sharma
- 38. Deep Structured Learning for Facial Action Unit Intensity Estimation, Robert Walecki, Ognjen (Oggi) Rudovic, Vladimir Pavlovic, Bjöern Schuller, Maja Pantic
- 39. Simultaneous Facial Landmark Detection, Pose and Deformation Estimation Under Facial Occlusion, Yue Wu, Chao Gou, Qiang Ji
- 40. Self-Supervised Video Representation Learning With Odd-One-Out Networks, *Basura Fernando, Hakan Bilen, Efstratios Gavves, Stephen Gould*
- 41. Robust Joint and Individual Variance Explained, Christos Sagonas, Yannis Panagakis, Alina Leidinger, Stefanos Zafeiriou

## Program

- 42. Discriminative Covariance Oriented Representation Learning for Face Recognition With Image Sets, *Wen Wang, Ruiping Wang, Shiguang Shan, Xilin Chen*
- 43. 3D Human Pose Estimation = 2D Pose Estimation + Matching, Ching-Hang Chen, Deva Ramanan

#### **Applications**

44. Joint Gap Detection and Inpainting of Line Drawings, Kazuma Sasaki, Satoshi lizuka, Edgar Simo-Serra, Hiroshi Ishikawa

#### Biomedical Image/Video Analysis

- 45. Riemannian Nonlinear Mixed Effects Models: Analyzing Longitudinal Deformations in Neuroimaging, Hyunwoo J. Kim, Nagesh Adluru, Heemanshu Suri, Baba C. Vemuri, Sterling C. Johnson, Vikas Singh
- 46. Simultaneous Super-Resolution and Cross-Modality Synthesis of 3D Medical Images Using Weakly-Supervised Joint Convolutional Sparse Coding, Yawen Huang, Ling Shao, Alejandro F. Frangi

#### Computational Photography

47. Multiple-Scattering Microphysics Tomography, Aviad Levis, Yoav Y. Schechner, Anthony B. Davis

#### Image Motion & Tracking

- Accurate Optical Flow via Direct Cost Volume Processing, Jia Xu, René Ranftl, Vladlen Koltun
- 49. Event-Based Visual Inertial Odometry, Alex Zihao Zhu, Nikolay Atanasov, Kostas Daniilidis
- 50. Robust Visual Tracking Using Oblique Random Forests, Le Zhang, Jagannadan Varadarajan, Ponnuthurai Nagaratnam Suganthan, Narendra Ahuja, Pierre Moulin

#### Low- & Mid-Level Vision

- 51. Deep Laplacian Pyramid Networks for Fast and Accurate Super-Resolution, Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, Ming-Hsuan Yang
- 52. Learning Non-Lambertian Object Intrinsics Across ShapeNet Categories, Jian Shi, Yue Dong, Hao Su, Stella X. Yu
- 53. MCMLSD: A Dynamic Programming Approach to Line Segment Detection, *Emilio J. Almazàn, Ron Tal, Yiming Qian, James H. Elder*

- 54. Contour-Constrained Superpixels for Image and Video Processing, Se-Ho Lee, Won-Dong Jang, Chang-Su Kim
- 55. Richer Convolutional Features for Edge Detection, Yun Liu, Ming-Ming Cheng, Xiaowei Hu, Kai Wang, Xiang Bai
- 56. Non-Local Color Image Denoising With Convolutional Neural Networks, *Stamatios Lefkimmiatis*
- 57. Generative Face Completion, Yijun Li, Sifei Liu, Jimei Yang, Ming-Hsuan Yang
- 58. Hyper-Laplacian Regularized Unidirectional Low-Rank Tensor Recovery for Multispectral Image Denoising, *Yi Chang, Luxin Yan, Sheng Zhong*
- 59. Unsupervised Semantic Scene Labeling for Streaming Data, *Maggie Wigness, John G. Rogers III*
- Why You Should Forget Luminance Conversion and Do Something Better, Rang M. H. Nguyen, Michael S. Brown
- 61. Deep Semantic Feature Matching, *Nikolai Ufer, Björn* Ommer

#### Machine Learning

- 62. Revisiting the Variable Projection Method for Separable Nonlinear Least Squares Problems, *Je Hyeong Hong, Christopher Zach, Andrew Fitzgibbon*
- Efficient Multiple Instance Metric Learning Using Weakly Supervised Data, Marc T. Law, Yaoliang Yu, Raquel Urtasun, Richard S. Zemel, Eric P. Xing
- 64. WILDCAT: Weakly Supervised Learning of Deep ConvNets for Image Classification, Pointwise Localization and Segmentation, *Thibaut Durand*, *Taylor Mordan*, *Nicolas Thome*, *Matthieu Cord*
- Image-To-Image Translation With Conditional Adversarial Networks, Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, Alexei A. Efros
- 66. Deep Roots: Improving CNN Efficiency With Hierarchical Filter Groups, Yani Ioannou, Duncan Robertson, Roberto Cipolla, Antonio Criminisi
- 67. Aggregated Residual Transformations for Deep Neural Networks, Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, Kaiming He
- MIML-FCN+: Multi-Instance Multi-Label Learning via Fully Convolutional Networks With Privileged Information, Hao Yang, Joey Tianyi Zhou, Jianfei Cai, Yew Soon Ong
- 69. Low-Rank Embedded Ensemble Semantic Dictionary for Zero-Shot Learning, *Zhengming Ding, Ming Shao, Yun Fu*

## Program

- 70. Factorized Variational Autoencoders for Modeling Audience Reactions to Movies, *Zhiwei Deng, Rajitha* Navarathna, Peter Carr, Stephan Mandt, Yisong Yue, Iain Matthews, Greg Mori
- 71. Learning Features by Watching Objects Move, Deepak Pathak, Ross Girshick, Piotr Dollár, Trevor Darrell, Bharath Hariharan
- 72. What Can Help Pedestrian Detection? Jiayuan Mao, Tete Xiao, Yuning Jiang, Zhimin Cao
- 73. DeepPermNet: Visual Permutation Learning, Rodrigo Santa Cruz, Basura Fernando, Anoop Cherian, Stephen Gould
- 74. Learning the Multilinear Structure of Visual Data, Mengjiao Wang, Yannis Panagakis, Patrick Snape, Stefanos Zafeiriou
- 75. Adaptive and Move Making Auxiliary Cuts for Binary Pairwise Energies, *Lena Gorelick, Yuri Boykov, Olga Veksler*
- 76. Designing Energy-Efficient Convolutional Neural Networks Using Energy-Aware Pruning, *Tien-Ju Yang, Yu-Hsin Chen, Vivienne Sze*
- 77. Joint Multi-Person Pose Estimation and Semantic Part Segmentation, *Fangting Xia, Peng Wang, Xianjie Chen, Alan L. Yuille*
- 78. Deep Feature Interpolation for Image Content Changes, Paul Upchurch, Jacob Gardner, Geoff Pleiss, Robert Pless, Noah Snavely, Kavita Bala, Kilian Weinberger
- 79. FASON: First and Second Order Information Fusion Network for Texture Recognition, *Xiyang Dai, Joe Yue-Hei* Ng, Larry S. Davis
- Lean Crowdsourcing: Combining Humans and Machines in an Online System, Steve Branson, Grant Van Horn, Pietro Perona

#### **Object Recognition & Scene Understanding**

- Supervising Neural Attention Models for Video Captioning by Human Gaze Data, Youngjae Yu, Jongwook Choi, Yeonhwa Kim, Kyung Yoo, Sang-Hun Lee, Gunhee Kim
- 82. L2-Net: Deep Learning of Discriminative Patch Descriptor in Euclidean Space, Yurun Tian, Bin Fan, Fuchao Wu
- Convolutional Random Walk Networks for Semantic Image Segmentation, Gedas Bertasius, Lorenzo Torresani, Stella X. Yu, Jianbo Shi

- 84. Knowledge Acquisition for Visual Question Answering via Iterative Querying, Yuke Zhu, Joseph J. Lim, Li Fei-Fei
- 85. Memory-Augmented Attribute Manipulation Networks for Interactive Fashion Search, *Bo Zhao, Jiashi Feng, Xiao Wu, Shuicheng Yan*
- 86. From Zero-Shot Learning to Conventional Supervised Classification: Unseen Visual Data Synthesis, Yang Long, Li Liu, Ling Shao, Fumin Shen, Guiguang Ding, Jungong Han
- Are Large-Scale 3D Models Really Necessary for Accurate Visual Localization? Torsten Sattler, Akihiko Torii, Josef Sivic, Marc Pollefeys, Hajime Taira, Masatoshi Okutomi, Tomas Pajdla
- 88. Asymmetric Feature Maps With Application to Sketch Based Retrieval, *Giorgos Tolias, Ondřej Chum*
- 89. Diverse Image Annotation, Baoyuan Wu, Fan Jia, Wei Liu, Bernard Ghanem
- 30. AMC: Attention guided Multi-modal Correlation Learning for Image Search, Kan Chen, Trung Bui, Chen Fang, Zhaowen Wang, Ram Nevatia
- 91. Multi-Attention Network for One Shot Learning, *Peng* Wang, Lingqiao Liu, Chunhua Shen, Zi Huang, Anton van den Hengel, Heng Tao Shen
- 92. Fried Binary Embedding for High-Dimensional Visual Features, *Weixiang Hong, Junsong Yuan, Sreyasee Das Bhattacharjee*
- 93. Pyramid Scene Parsing Network, Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, Jiaya Jia
- 94. Learning Deep Match Kernels for Image-Set Classification, *Haoliang Sun, Xiantong Zhen, Yuanjie Zheng, Gongping Yang, Yilong Yin, Shuo Li*
- 95. Task-Driven Dynamic Fusion: Reducing Ambiguity in Video Description, *Xishan Zhang, Ke Gao, Yongdong Zhang, Dongming Zhang, Jintao Li, Qi Tian*
- 96. Learning Multifunctional Binary Codes for Both Category and Attribute Oriented Retrieval Tasks, *Haomiao Liu*, *Ruiping Wang, Shiguang Shan, Xilin Chen*
- 97. Indoor Scene Parsing With Instance Segmentation, Semantic Labeling and Support Relationship Inference, Wei Zhuo, Mathieu Salzmann, Xuming He, Miaomiao Liu

- 98. Episodic CAMN: Contextual Attention-Based Memory Networks With Iterative Feedback for Scene Labeling, *Abrar H. Abdulnabi, Bing Shuai, Stefan Winkler, Gang Wang*
- 99. Link the Head to the "Beak": Zero Shot Learning From Noisy Text Description at Part Precision, Mohamed Elhoseiny, Yizhe Zhu, Han Zhang, Ahmed Elgammal
- 100. SCA-CNN: Spatial and Channel-Wise Attention in Convolutional Networks for Image Captioning, Long Chen, Hanwang Zhang, Jun Xiao, Liqiang Nie, Jian Shao, Wei Liu, Tat-Seng Chua
- 101. Deep Pyramidal Residual Networks, Dongyoon Han, Jiwhan Kim, Junmo Kim
- 102. Product Split Trees, Artem Babenko, Victor Lempitsky
- 103. Making the v in VQA Matter: Elevating the Role of Image Understanding in Visual Question Answering, Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, Devi Parikh
- 104. Commonly Uncommon: Semantic Sparsity in Situation Recognition, Mark Yatskar, Vicente Ordonez, Luke Zettlemoyer, Ali Farhadi
- 105. Cross-Modality Binary Code Learning via Fusion Similarity Hashing, Hong Liu, Rongrong Ji, Yongjian Wu, Feiyue Huang, Baochang Zhang

#### Theory

- 106. Saliency Revisited: Analysis of Mouse Movements Versus Fixations, Hamed R. Tavakoli, Fawad Ahmed, Ali Borji, Jorma Laaksonen
- 107. InterpoNet, a Brain Inspired Neural Network for Optical Flow Dense Interpolation, *Shay Zweig, Lior Wolf*

#### Video Analytics

- 108.SST: Single-Stream Temporal Action Proposals, Shyamal Buch, Victor Escorcia, Chuanqi Shen, Bernard Ghanem, Juan Carlos Niebles
- 109. Video Segmentation via Multiple Granularity Analysis, *Rui* Yang, Bingbing Ni, Chao Ma, Yi Xu, Xiaokang Yang
- 110. Spatio-Temporal Alignment of Non-Overlapping Sequences From Independently Panning Cameras, Seyed Morteza Safdarnejad, Xiaoming Liu
- 111. UntrimmedNets for Weakly Supervised Action Recognition and Detection, Limin Wang, Yuanjun Xiong, Dahua Lin, Luc Van Gool

## 1000-1200 Demos (Kamehameha I)


- Open-Source Simulator in UE4: Photorealistic Fully Annotated Datasets, Real-Time Object Tracking With VR, Deep Learning Interface, Matthias Mueller, Neil Smith, Bernard Ghanem (King Abdullah Univ. of Science and Technology)
- CloudCV, Deshraj Yadav, Viraj Prabhu, Prithvijit Chattopadhyay, Abhishek Das, Devi Parikh, Dhruv Batra (Virginia Tech, Georgia Tech)
- Procedural Human Action Videos, César De Souza, Yohann Cabon, Adrien Gaidon, Antonio M. Lopez (Xerox Research Centre Europe)
- Hierarchical 3D Fully Convolutional Networks for Multi-Organ and Vessel Segmentation Used in Surgical Navigation, Holger R. Roth, Yusuke Tetsumura, Hirohisa Oda, Yuichiro Hayashi, Masahiro Oda, Natsuki Shimizu, Michitaka Fujiwara, Kazunari Misawa, Kensaku Mori (Nagoya Univ., Nagoya Univ. Graduate School of Medicine, Aichi Cancer Center)

## 1000-1200 Exhibits (Kamehameha I)

• Same as Saturday morning Exhibits (see pg. 10)

## 1200-1300 Lunch (Kamehameha II)

Notes:



#### 1300-1430 Session 4-2A: Object Recognition & Scene Understanding 3 (Kamehameha III)

#### Papers in this session are also in Poster Session P4-2.

Chairs: Kostas Daniilidis (Univ. of Pennsylvania) Junsong Yuan (Nanyang Technological Univ)

#### 1300 Spotlights (S4-2A)

#### Format (4 min. for presentation; no questions)

- [1300] Gaze Embeddings for Zero-Shot Image Classification, Nour Karessli, Zeynep Akata, Bernt Schiele, Andreas Bulling
- [1304] What's in a Question: Using Visual Questions as a Form of Supervision, Siddha Ganju, Olga Russakovsky, Abhinav Gupta
- [1308] Attend to You: Personalized Image Captioning With Context Sequence Memory Networks, Cesc Chunseong Park, Byeongchang Kim, Gunhee Kim
- 4. [1312] Adversarially Tuned Scene Generation, VSR Veeravasarapu, Constantin Rothkopf, Ramesh Visvanathan
- 5. **[1316]** Residual Attention Network for Image Classification, *Fei Wang, Mengqing Jiang, Chen Qian, Shuo Yang, Cheng Li, Honggang Zhang, Xiaogang Wang, Xiaoou Tang*
- 6. [1320] Not All Pixels Are Equal: Difficulty-Aware Semantic Segmentation via Deep Layer Cascade, Xiaoxiao Li, Ziwei Liu, Ping Luo, Chen Change Loy, Xiaoou Tang
- 7. [1324] Learning Non-Maximum Suppression, Jan Hosang, Rodrigo Benenson, Bernt Schiele
- [1328] The Amazing Mysteries of the Gutter: Drawing Inferences Between Panels in Comic Book Narratives, Mohit Iyyer, Varun Manjunatha, Anupam Guha, Yogarshi Vyas, Jordan Boyd-Graber, Hal Daumé III, Larry S. Davis

#### 1333 Orals (O4-2A)

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

- 9. [1333] Object Region Mining With Adversarial Erasing: A Simple Classification to Semantic Segmentation Approach, Yunchao Wei, Jiashi Feng, Xiaodan Liang, Ming-Ming Cheng, Yao Zhao, Shuicheng Yan
- 10. **[1347]** Fine-Grained Recognition as HSnet Search for Informative Image Parts, *Michael Lam, Behrooz Mahasseni, Sinisa Todorovic*

- 11. [1401] G<sup>2</sup>DeNet: Global Gaussian Distribution Embedding Network and Its Application to Visual Recognition, *Qilong* Wang, Peihua Li, Lei Zhang
- 12. **[1415]** YOLO9000: Better, Faster, Stronger, Joseph Redmon, Ali Farhadi

### 1300–1430 Session 4-2B: Machine Learning for 3D Vision (Kalākaua Ballroom)

#### Papers in this session are also in Poster Session P4-2.

Chairs: Dieter Fox (Univ. of Washington) Marc Pollefeys (ETH Zurich)

#### 1300 Spotlights (S4-2B)

Format (4 min. for presentation; no questions)

- 13. **[1300]** Multi-View 3D Object Detection Network for Autonomous Driving, *Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, Tian Xia*
- 14. **[1304]** UltraStereo: Efficient Learning-Based Matching for Active Stereo Systems, Sean Ryan Fanello, Julien Valentin, Christoph Rhemann, Adarsh Kowdle, Vladimir Tankovich, Philip Davidson, Shahram Izadi
- 15. **[1308]** Shape Completion Using 3D-Encoder-Predictor CNNs and Shape Synthesis, *Angela Dai, Charles Ruizhongtai Qi, Matthias Nießner*
- 16. **[1312]** Geometric Loss Functions for Camera Pose Regression With Deep Learning, *Alex Kendall, Roberto Cipolla*
- 17. **[1316]** CNN-SLAM: Real-Time Dense Monocular SLAM With Learned Depth Prediction, *Keisuke Tateno, Federico Tombari, Iro Laina, Nassir Navab*
- [1320] Learning From Noisy Large-Scale Datasets With Minimal Supervision, Andreas Veit, Neil Alldrin, Gal Chechik, Ivan Krasin, Abhinav Gupta, Serge Belongie
- 19. **[1324]** SyncSpecCNN: Synchronized Spectral CNN for 3D Shape Segmentation, *Li Yi, Hao Su, Xingwen Guo, Leonidas J. Guibas*
- 20. [1328] Non-Local Deep Features for Salient Object Detection, Zhiming Luo, Akshaya Mishra, Andrew Achkar, Justin Eichel, Shaozi Li, Pierre-Marc Jodoin

## Program

## Program

### 1333 Orals (O4-2B)

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

- 21. [1333] Unsupervised Monocular Depth Estimation With Left-Right Consistency, Clément Godard, Oisin Mac Aodha, Gabriel J. Brostow
- 22. [1347] Unsupervised Learning of Depth and Ego-Motion From Video, *Tinghui Zhou, Matthew Brown, Noah Snavely, David G. Lowe*
- 23. [1401] OctNet: Learning Deep 3D Representations at High Resolutions, *Gernot Riegler, Ali Osman Ulusoy, Andreas Geiger*
- 24. **[1415]** 3D Shape Segmentation With Projective Convolutional Networks, *Evangelos Kalogerakis, Melinos Averkiou, Subhransu Maji, Siddhartha Chaudhuri*

## 1430-1515 Break (Kamehameha II)

## 1430-1630 Poster Session P4-2 (Kamehameha I)

#### 3D Computer Vision

- SGM-Nets: Semi-Global Matching With Neural Networks, Akihito Seki, Marc Pollefeys
- 26. Stereo-Based 3D Reconstruction of Dynamic Fluid Surfaces by Global Optimization, *Yiming Qian, Minglun Gong, Yee-Hong Yang*
- 27. Fine-To-Coarse Global Registration of RGB-D Scans, Maciej Halber, Thomas Funkhouser
- 28. Analyzing Computer Vision Data The Good, the Bad and the Ugly, Oliver Zendel, Katrin Honauer, Markus Murschitz, Martin Humenberger, Gustavo Fernández Domínguez
- 29. Product Manifold Filter: Non-Rigid Shape Correspondence via Kernel Density Estimation in the Product Space, Matthias Vestner, Roee Litman, Emanuele Rodolà, Alex Bronstein, Daniel Cremers
- 30. Unsupervised Vanishing Point Detection and Camera Calibration From a Single Manhattan Image With Radial Distortion, Michel Antunes, João P. Barreto, Djamila Aouada, Björn Ottersten
- 31. Toroidal Constraints for Two-Point Localization Under High Outlier Ratios, Federico Camposeco, Torsten Sattler, Andrea Cohen, Andreas Geiger, Marc Pollefeys

- 32. 4D Light Field Superpixel and Segmentation, *Hao Zhu, Qi Zhang, Qing Wang*
- 33. Exploiting Symmetry and/or Manhattan Properties for 3D Object Structure Estimation From Single and Multiple Images, Yuan Gao, Alan L. Yuille

### Analyzing Humans in Images

- 34. Binary Coding for Partial Action Analysis With Limited Observation Ratios, Jie Qin, Li Liu, Ling Shao, Bingbing Ni, Chen Chen, Fumin Shen, Yunhong Wang
- 35. SphereFace: Deep Hypersphere Embedding for Face Recognition, *Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, Le Song*
- 36. IRINA: Iris Recognition (Even) in Inaccurately Segmented Data, Hugo Proença, João C. Neves
- 37. Look Into Person: Self-Supervised Structure-Sensitive Learning and a New Benchmark for Human Parsing, Ke Gong, Xiaodan Liang, Dongyu Zhang, Xiaohui Shen, Liang Lin
- 38. Action Unit Detection With Region Adaptation, Multi-Labeling Learning and Optimal Temporal Fusing, Wei Li, Farnaz Abtahi, Zhigang Zhu
- 39. See the Forest for the Trees: Joint Spatial and Temporal Recurrent Neural Networks for Video-Based Person Re-Identification, Zhen Zhou, Yan Huang, Wei Wang, Liang Wang, Tieniu Tan
- 40. Joint Intensity and Spatial Metric Learning for Robust Gait Recognition, Yasushi Makihara, Atsuyuki Suzuki, Daigo Muramatsu, Xiang Li, Yasushi Yagi
- 41. Pose-Aware Person Recognition, Vijay Kumar, Anoop Namboodiri, Manohar Paluri, C. V. Jawahar
- 42. Not Afraid of the Dark: NIR-VIS Face Recognition via Cross-Spectral Hallucination and Low-Rank Embedding, José Lezama, Qiang Qiu, Guillermo Sapiro

## **Applications**

- 43. Jointly Learning Energy Expenditures and Activities Using Egocentric Multimodal Signals, Katsuyuki Nakamura, Serena Yeung, Alexandre Alahi, Li Fei-Fei
- 44. Binarized Mode Seeking for Scalable Visual Pattern Discovery, Wei Zhang, Xiaochun Cao, Rui Wang, Yuanfang Guo, Zhineng Chen

## Program

45. Scribbler: Controlling Deep Image Synthesis With Sketch and Color, Patsorn Sangkloy, Jingwan Lu, Chen Fang, Fisher Yu, James Hays

#### Biomedical Image/Video Analysis

- 46. Multi-Way Multi-Level Kernel Modeling for Neuroimaging Classification, Lifang He, Chun-Ta Lu, Hao Ding, Shen Wang, Linlin Shen, Philip S. Yu, Ann B. Ragin
- 47. WSISA: Making Survival Prediction From Whole Slide Histopathological Images, Xinliang Zhu, Jiawen Yao, Feiyun Zhu, Junzhou Huang

#### Computational Photography

- 48. On the Effectiveness of Visible Watermarks, Tali Dekel, Michael Rubinstein, Ce Liu, William T. Freeman
- 49. Snapshot Hyperspectral Light Field Imaging, Zhiwei Xiong, Lizhi Wang, Huigun Li, Dong Liu, Feng Wu
- 50. Semantic Image Inpainting With Deep Generative Models, Raymond A. Yeh, Chen Chen, Teck Yian Lim, Alexander G. Schwing, Mark Hasegawa-Johnson, Minh N. Do

#### Image Motion & Tracking

- 51. Fast Multi-Frame Stereo Scene Flow With Motion Segmentation, Tatsunori Taniai, Sudipta N. Sinha, Yoichi Sato
- 52. Improved Stereo Matching With Constant Highway Networks and Reflective Confidence Learning, Amit Shaked, Lior Wolf
- 53. Optical Flow in Mostly Rigid Scenes, Jonas Wulff, Laura Sevilla-Lara, Michael J. Black
- 54. Optical Flow Requires Multiple Strategies (but Only One Network), *Tal Schuster, Lior Wolf, David Gadot*
- 55. ECO: Efficient Convolution Operators for Tracking, Martin Danelljan, Goutam Bhat, Fahad Shahbaz Khan, Michael Felsberg

#### Low- & Mid-Level Vision

- 56. Differential Angular Imaging for Material Recognition, Jia Xue, Hang Zhang, Kristin Dana, Ko Nishino
- 57. Fast Fourier Color Constancy, Jonathan T. Barron, Yun-Ta Tsai
- 58. Comparative Evaluation of Hand-Crafted and Learned Local Features, *Johannes L. Schönberger, Hans Hardmeier, Torsten Sattler, Marc Pollefeys*

- 59. Learning Fully Convolutional Networks for Iterative Non-Blind Deconvolution, *Jiawei Zhang, Jinshan Pan, Wei-Sheng Lai, Rynson W. H. Lau, Ming-Hsuan Yang*
- 60. Image Deblurring via Extreme Channels Prior, Yanyang Yan, Wenqi Ren, Yuanfang Guo, Rui Wang, Xiaochun Cao
- 61. Simultaneous Stereo Video Deblurring and Scene Flow Estimation, *Liyuan Pan*, *Yuchao Dai, Miaomiao Liu, Fatih Porikli*
- 62. Deep Photo Style Transfer, Fujun Luan, Sylvain Paris, Eli Shechtman, Kavita Bala
- 63. Generative Attribute Controller With Conditional Filtered Generative Adversarial Networks, *Takuhiro Kaneko, Kaoru Hiramatsu, Kunio Kashino*
- 64. Fast Haze Removal for Nighttime Image Using Maximum Reflectance Prior, Jing Zhang, Yang Cao, Shuai Fang, Yu Kang, Chang Wen Chen

#### Machine Learning

- 65. Low-Rank Bilinear Pooling for Fine-Grained Classification, Shu Kong, Charless Fowlkes
- 66. Neural Scene De-Rendering, Jiajun Wu, Joshua B. Tenenbaum, Pushmeet Kohli
- 67. Real-Time Neural Style Transfer for Videos, Haozhi Huang, Hao Wang, Wenhan Luo, Lin Ma, Wenhao Jiang, Xiaolong Zhu, Zhifeng Li, Wei Liu
- 68. A Graph Regularized Deep Neural Network for Unsupervised Image Representation Learning, Shijie Yang, Liang Li, Shuhui Wang, Weigang Zhang, Qingming Huang
- 69. A Study of Lagrangean Decompositions and Dual Ascent Solvers for Graph Matching, Paul Swoboda, Carsten Rother, Hassan Abu Alhaija, Dagmar Kainmüller, Bogdan Savchynskyy
- 70. Collaborative Deep Reinforcement Learning for Joint Object Search, Xiangyu Kong, Bo Xin, Yizhou Wang, Gang Hua
- 71. Loss Max-Pooling for Semantic Image Segmentation, Samuel Rota Bulò, Gerhard Neuhold, Peter Kontschieder
- 72. Deep View Morphing, Dinghuang Ji, Junghyun Kwon, Max McFarland, Silvio Savarese
- 73. Unsupervised Learning of Long-Term Motion Dynamics for Videos, Zelun Luo, Boya Peng, De-An Huang, Alexandre Alahi, Li Fei-Fei

- 74. Revisiting Metric Learning for SPD Matrix Based Visual Representation, *Luping Zhou, Lei Wang, Jianjia Zhang, Yinghuan Shi, Yang Gao*
- 75. Expert Gate: Lifelong Learning With a Network of Experts, Rahaf Aljundi, Punarjay Chakravarty, Tinne Tuytelaars
- 76. A Gift From Knowledge Distillation: Fast Optimization, Network Minimization and Transfer Learning, Junho Yim, Donggyu Joo, Jihoon Bae, Junmo Kim
- 77. Domain Adaptation by Mixture of Alignments of Secondor Higher-Order Scatter Tensors, *Piotr Koniusz, Yusuf Tas, Fatih Porikli*
- 78. Deep Mixture of Linear Inverse Regressions Applied to Head-Pose Estimation, Stéphane Lathuilière, Rémi Juge, Pablo Mesejo, Rafael Muñoz-Salinas, Radu Horaud
- 79. STD2P: RGBD Semantic Segmentation Using Spatio-Temporal Data-Driven Pooling, Yang He, Wei-Chen Chiu, Margret Keuper, Mario Fritz
- Harmonic Networks: Deep Translation and Rotation Equivariance, Daniel E. Worrall, Stephan J. Garbin, Daniyar Turmukhambetov, Gabriel J. Brostow
- Multimodal Transfer: A Hierarchical Deep Convolutional Neural Network for Fast Artistic Style Transfer, Xin Wang, Geoffrey Oxholm, Da Zhang, Yuan-Fang Wang
- 82. Detect, Replace, Refine: Deep Structured Prediction for Pixel Wise Labeling, *Spyros Gidaris, Nikos Komodakis*
- 83. Weighted-Entropy-Based Quantization for Deep Neural Networks, Eunhyeok Park, Junwhan Ahn, Sungjoo Yoo
- Residual Expansion Algorithm: Fast and Effective Optimization for Nonconvex Least Squares Problems, Daiki Ikami, Toshihiko Yamasaki, Kiyoharu Aizawa
- 85. Bidirectional Beam Search: Forward-Backward Inference in Neural Sequence Models for Fill-In-The-Blank Image Captioning, Qing Sun, Stefan Lee, Dhruv Batra
- Newton-Type Methods for Inference in Higher-Order Markov Random Fields, Hariprasad Kannan, Nikos Komodakis, Nikos Paragios
- 87. Adaptive Relaxed ADMM: Convergence Theory and Practical Implementation, *Zheng Xu*, *Mário A. T. Figueiredo, Xiaoming Yuan, Christoph Studer, Tom Goldstein*

#### **Object Recognition & Scene Understanding**

- ViP-CNN: Visual Phrase Guided Convolutional Neural Network, Yikang Li, Wanli Ouyang, Xiaogang Wang, Xiao'ou Tang
- Instance-Aware Image and Sentence Matching With Selective Multimodal LSTM, Yan Huang, Wei Wang, Liang Wang
- 90. Kernel Square-Loss Exemplar Machines for Image Retrieval, Rafael S. Rezende, Joaquin Zepeda, Jean Ponce, Francis Bach, Patrick Pérez
- 91. Cognitive Mapping and Planning for Visual Navigation, Saurabh Gupta, James Davidson, Sergey Levine, Rahul Sukthankar, Jitendra Malik
- 92. Combining Bottom-Up, Top-Down, and Smoothness Cues for Weakly Supervised Image Segmentation, Anirban Roy, Sinisa Todorovic
- 93. Seeing Into Darkness: Scotopic Visual Recognition, Bo Chen, Pietro Perona
- 94. Deep Co-Occurrence Feature Learning for Visual Object Recognition, Ya-Fang Shih, Yang-Ming Yeh, Yen-Yu Lin, Ming-Fang Weng, Yi-Chang Lu, Yung-Yu Chuang
- 95. An Empirical Evaluation of Visual Question Answering for Novel Objects, Santhosh K. Ramakrishnan, Ambar Pal, Gaurav Sharma, Anurag Mittal
- 96. InstanceCut: From Edges to Instances With MultiCut, Alexander Kirillov, Evgeny Levinkov, Bjoern Andres, Bogdan Savchynskyy, Carsten Rother
- 97. Fine-Grained Image Classification via Combining Vision and Language, *Xiangteng He, Yuxin Peng*
- 98. Mimicking Very Efficient Network for Object Detection, Quanquan Li, Shengying Jin, Junjie Yan
- 99. Tracking by Natural Language Specification, Zhenyang Li, Ran Tao, Efstratios Gavves, Cees G. M. Snoek, Arnold W.M. Smeulders
- 100. A Dataset and Exploration of Models for Understanding Video Data Through Fill-In-The-Blank Question-Answering, Tegan Maharaj, Nicolas Ballas, Anna Rohrbach, Aaron Courville, Christopher Pal
- 101. Learning Detection With Diverse Proposals, Samaneh Azadi, Jiashi Feng, Trevor Darrell

102. Skeleton Key: Image Captioning by Skeleton-Attribute Decomposition, Yufei Wang, Zhe Lin, Xiaohui Shen, Scott Cohen, Garrison W. Cottrell

#### Theory

103. A Low Power, Fully Event-Based Gesture Recognition System, Arnon Amir, Brian Taba, David Berg, Timothy Melano, Jeffrey McKinstry, Carmelo Di Nolfo, Tapan Nayak, Alexander Andreopoulos, Guillaume Garreau, Marcela Mendoza, Jeff Kusnitz, Michael Debole, Steve Esser, Tobi Delbruck, Myron Flickner, Dharmendra Modha

#### Video Analytics

- 104. Learning Deep Context-Aware Features Over Body and Latent Parts for Person Re-Identification, *Dangwei Li, Xiaotang Chen, Zhang Zhang, Kaiqi Huang*
- 105. Recurrent Modeling of Interaction Context for Collective Activity Recognition, *Minsi Wang, Bingbing Ni, Xiaokang Yang*
- 106. Primary Object Segmentation in Videos Based on Region Augmentation and Reduction, Yeong Jun Koh, Chang-Su Kim
- 107. ROAM: A Rich Object Appearance Model With Application to Rotoscoping, Ondrej Miksik, Juan-Manuel Pérez-Rúa, Philip H. S. Torr, Patrick Pérez
- 108. Temporal Residual Networks for Dynamic Scene Recognition, Christoph Feichtenhofer, Axel Pinz, Richard P. Wildes
- 109. Spatiotemporal Multiplier Networks for Video Action Recognition, Christoph Feichtenhofer, Axel Pinz, Richard P. Wildes
- 110. Learning to Learn From Noisy Web Videos, Serena Yeung, Vignesh Ramanathan, Olga Russakovsky, Liyue Shen, Greg Mori, Li Fei-Fei
- 111. YouTube-BoundingBoxes: A Large High-Precision Human-Annotated Data Set for Object Detection in Video, Esteban Real, Jonathon Shlens, Stefano Mazzocchi, Xin Pan, Vincent Vanhoucke
- 112. Online Video Object Segmentation via Convolutional Trident Network, *Won-Dong Jang, Chang-Su Kim*

#### 1430-1630 Demos (Kamehameha I)

• Same as Tuesday morning Demos (see pg. 37)

## 1430-1630 Exhibits (Kamehameha I)

• Same as Saturday morning Exhibits (see pg. 10)

### 1645-1800 Plenary Session (Kamehameha III)

 Keynote Talk: The Science of Natural Intelligence (NI): Reverse Engineering Primate Visual Perception, James J. DiCarlo (MIT)

Abstract: The fields of neuroscience and cognitive science are hard at work on one of our last great scientific guests - to reverse engineer the human mind. In comparison to other areas of science, these sciences are still in their infancy. Not surprisingly, forward engineering approaches that aim to emulate human intelligence in artificial systems (AI) are also still in their infancy. Yet the intelligence and cognitive flexibility apparent in human behavior are an existence proof that machines can be constructed to emulate and work alongside the human mind. In this talk, I will argue that these challenges of reverse engineering the mind will be solved by tightly combining the efforts of brain and cognitive scientists (hypothesis generation and data acquisition), and forward engineering aiming to emulate the mind (hypothesis instantiation and data prediction). To support that thesis, I will focus on one aspect of perceptual intelligence - object categorization and detection — and I will tell the story of how work in brain science, cognitive science and computer science converged to create deep neural networks that can support such tasks. These networks not only reach human performance for many images, but their internal workings are modeled after—and largely emulate — the internal workings of the primate visual system. Yet, the primate visual system (NI) still outperforms current generation artificial deep neural networks (AI), and I will show some new clues that neuroscience can offer. More broadly, this is just the beginning of the last great human science quest — to understand natural intelligence — and I hope to motivate others to engage that frontier alongside us.

### END OF MAIN CVPR 2017 CONFERENCE



