

Google Research

CVPR 2020

At Google, research in machine perception tackles the hard problems of understanding images, sounds, music, and video. This research allows us to better serve our users by enabling a variety of applications such as content-based search in Google Photos and Image Search, natural handwriting interfaces for Android, optical character recognition for Google Drive documents, and recommendation systems that understand music and YouTube videos. Our approach is driven by developing cutting-edge algorithms that benefit from processing very large, partially-labeled datasets using state-of-the-art hardware, leveraging the ability to mine meaningful information from multimedia. While at CVPR, we hope you'll connect with our researchers who are actively pursuing the next generation of intelligent systems that utilize the latest machine learning techniques. To learn more about the research we're presenting at CVPR 2020, please see the list below (Googlers **bolded**).

Committees

General Chair: **Ramin Zabih**
Program Chair: **Ce Liu**
Workshop Chair: **Tali Dekel**
Website Chair: **Tianfan Xue**
Technical Chair: **Daniel Vlasic**
Area Chairs: **Alexander Toshev, Alexey Dosovitskiy, Boqing Gong, Caroline Pantofaru, Chen Sun, Deqing Sun, Dilip Krishnan, Feng Yang, Liang-Chieh Chen, Michael Rubinstein, Rodrigo Benenson, Timnit Gebru, Thomas Funkhouser, Varun Jampani, Vittorio Ferrari, William Freeman**

Sunday, June 14

Workshops

3rd Workshop and Challenge on Learned Image Compression

George Toderici, *Wenzhe Shi*, **Eirikur Agustsson**,
Radu Timofte, **Lucas Theis**, **Johannes Ballé**,
Nick Johnston, *Fabian Mentzer*

June 14; Full-Day

[Workshop URL](#)

CLVISION 1st Workshop on Continual Learning in Computer Vision

Pau Rodriguez, German I. Parisi, David Vazquez, Vincenzo Lomonaco, Nikhil Churamani,
Zhiyuan (Brett) Chen, *Marc Pickett*

June 14; Full-Day

[Workshop URL](#)

Embodied AI

Aniruddha Kembhavi, Roozbeh Mottaghi, Dustin Schwenk, Eric Kolve, Ali Farhadi, **Alexander Toshev**, Roberto Martín-Martín, Amir Zamir, **Jie Tan**, **Aleksandra Faust**, **Anelia Angelova**, Erik Wijmans, Samyak Datta, Abhishek Kadian, Oleksandr Maksymets, Julia Straub, Jitendra Malik, Vladlen Koltun, Stefan Lee, Angel X. Chang, Manolis Savva, Devi Parikh, Dhruv Batra, German Ros, Jose M. Alvarez, Jose A. Iglesias-Guitian, Antonio Lopez, Yongjoon Lee, Vangelis Kokkevis, Philipp Slusallek

June 14; Full-Day

June 15; Full-Day

[Workshop URL](#)

The 1st International Workshop and Prize Challenge on Agriculture-Vision: Challenges & Opportunities for Computer Vision in Agriculture

Humphrey Honghui Shi, Naira Hovakimyan, Thomas S. Huang, **Zhen Li**, Adrian Tudor, Robert E. Paull, **Jim Yuan**

June 14; Full-Day

[Workshop URL](#)

Tutorials

How to be a good reviewer ... as part of The System

Vittorio Ferrari

June 14; Afternoon

[Tutorial URL](#)

Disentangled 3D Representations for Relightable Performance Capture of Humans

Sean Fanello, **Christoph Rhemann**, **Jonathan Taylor**, **Sofien Bouaziz**, **Adarsh Kowdle**, **Rohit Pandey**, **Sergio Orts-Escolano**, **Paul Debevec**, **Shahram Izadi**

June 14; Full-Day

[Tutorial URL](#)

Learning Representations via Graph-Structured Networks

Xiaolong Wang, Sifei Liu, Saining Xie, Shubham Tulsiani, **Chen Sun**, Han Hu, Jan Kautz, **Ming-Hsuan Yang**, Abhinav Gupta, Trevor Darrell

June 14; Afternoon

[Tutorial URL](#)

Novel View Synthesis: From Depth-Based Warping to Multi-Plane Images and Beyond

Orazio Gallo, Alejandro Troccoli, **Varun Jampani**

June 14; Full-Day

[Tutorial URL](#)

Monday, June 15

Workshops

Embodied AI

Aniruddha Kembhavi, Roozbeh Mottaghi, Dustin Schwenk, Eric Kolve, Ali Farhadi, **Alexander Toshev**, Roberto Martín-Martín, Amir Zamir, **Jie Tan**, **Aleksandra Faust**, **Anelia Angelova**, Erik Wijmans, Samyak Datta, Abhishek Kadian, Oleksandr Maksymets, Julia Straub, Jitendra Malik, Vladlen Koltun, Stefan Lee, Angel X. Chang, Manolis Savva, Devi Parikh, Dhruv Batra, German Ros, Jose M. Alvarez, Jose A. Iglesias-Guitian, Antonio Lopez, Yongjoon Lee, Vangelis Kokkevis, Philipp Slusallek

June 14; Full-Day

June 15; Full-Day

[Workshop URL](#)

New Trends in Image Restoration and Enhancement workshop and challenges on image and video restoration and enhancement (NTIRE)

Talk: "Sky Optimization: Semantically aware image processing of skies in low-light photography"

Orly Liba, **Longqi Cai**, **Yun-Ta Tsai**, **Elad Eban**, **Yair Movshovitz-Attias**, **Yael Pritch**, **Huizhong Chen**, **Jonathan Barron**

June 15; 12.50 - 1.00 PM

[Workshop URL](#)

The End-of-End-to-End A Video Understanding Pentathlon

Samuel Albanie, Yang Liu, Arsha Nagrani, Antoine Miech, Ernesto Coto, Ivan Laptev, **Rahul Sukthankar**, Bernard Ghanem, Andrew Zisserman

June 15; Afternoon

[Workshop URL](#)

4th Workshop on Media Forensics

Arslan Basharat, **Christoph Bregler**, Cristian Canton, Davide Cozzolino, Edward Delp, Sergio Escalera, Hany Farid, Sam Gregory, Guodong Guo, Isabelle Guyon, Hugo Jair, Laura Leal-Taixé, Ser-Nam Lim, Scott McCloskey, Matthias Nießner, Andreas Rössler, Christa Scharfenberg, Justus Thies, Luisa Verdoliva, Jun Wan

June 15; Full-Day

[Workshop URL](#)

4th Workshop on Visual Understanding by Learning from Web Data

Wen Li, Hilde Kuehne, Suman Saha, Qin Wang, Limin Wang, Wei Li, **Jesse Berent**, Abhinav Gupta, **Rahul Sukthankar**, Luc Van Gool

June 15; Full-Day

[Workshop URL](#)

AI for Content Creation

Deqing Sun, Ming-Yu Liu, **Lu Jiang**, James Tompkin, **Weilong Yang**, Kalyan Sunkavalli

June 15; Full-Day

[Workshop URL](#)

Fourth Workshop on Computer Vision for AR/VR

Fernando De la Torre, **Sofien Bouaziz**, Harald Haraldsson, Serge Belongie, Andrew Rabinovich, Matt Uyttendaele

June 15; Full-Day

[Workshop URL](#)

Low-Power Computer Vision Competition (LPCVC)

Hartwig Adam, Mohamed M. Sabry Aly, Alexander, C Berg, Yiran Chen, **Bo Chen**, Yen-Kuang Chen, **Andrew Howard**, Christian Keller, **Jaeyoun Kim**, Mark Liao, Yung-Hsiang Lu, An Luo, Ryan McBee, Andreas Moshovos, Ashish Sirasao, Joseph Spisak, Yueh-Hua Wu, Junsong Yuan

June 15; Full-Day

[Workshop URL](#)

Sight and Sound

Andrew Owens, Jiajun Wu, Jean-Charles Bazin, Andrew Zisserman, **William Freeman**, Antonio Torralba, Kristen Grauman

June 15; Full-Day

[Workshop URL](#)

Workshop on Efficient Deep Learning for Computer Vision

Li Liu, Yu Liu, Wanli Ouyang, Jiwen Lu, Matti Pietikäinen, Luc Van Gool, Chris Rowen, Peizhao Zhang, **Pete Warden**, Kurt Keutzer, Jonathan Dekhtiar, Wojciech Samek, Yingyan Lin, Werner Bailer, Peter Vajda

June 15; Full-Day

[Workshop URL](#)

Tutorials

Neural Rendering

Ayush Tewari, Ohad Fried, Justus Thies, Vincent Sitzmann, Stephen Lombardi, Kalyan Sunkavalli,

Ricardo Martin-Brualla, Tomas Simon, Jason Saragih, Matthias Nießner, **Rohit K. Pandey**, Sean

Fanello, Gordon Wetzstein, Jun-Yan Zhu, Christian

Theobalt, **Maneesh Agrawala**, Eli Shechtman, **Dan B.**

Goldman, Michael Zollhöfer

June 15; Full-Day

[Tutorial URL](#)

Tuesday, June 16

Orals

Evolving Losses for Unsupervised Video Representation Learning

AJ Piergiovanni, **Anelia Angelova**, **Michael Ryoo**

10:05 & 22:05

[Link to Paper](#)

CvxNet: Learnable Convex Decomposition

Boyang Deng, **Kyle Genova**, **Soroosh Yazdani**, **Sofien Bouaziz**, **Geoffrey Hinton**, **Andrea Tagliasacchi**

10:15 & 22:15

[Link to Paper](#)

Neural SDE: Stabilizing Neural ODE Networks with Stochastic Noise

Xuanqing Liu, Tesi Xiao, **Si Si**, **Qin Cao**, **Sanjiv Kumar**, Cho-Jui Hsieh

10:25 & 22:25

[Link to Paper](#)

Scalability in Perception for Autonomous Driving: Waymo Open Dataset

Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurélien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, **Benjamin Caine**, **Vijay Vasudevan**, **Wei Han**, **Jiquan Ngiam**, Hang Zhao, Aleksei Timofeev, Scott Ettinger, Maxim Krivokon, Amy Gao, Aditya Joshi, Sheng Zhao, Shuyang Chen, Yu Zhang, **Jon Shlens**, **Zhifeng Chen**, Dragomir Anguelov

12:00 & 0:00

[Link to Paper](#)

Deep Implicit Volume Compression

Saurabh Singh, Danhang Tang, Cem Keskin, Philip Chou, Christian Haene, Mingsong Dou, Sean Fanello, Jonathan Taylor, Andrea Tagliasacchi, Philip Davidson, Yinda Zhang, Onur Guleryuz, Shahram Izadi, Sofien Bouaziz

12:20 & 00:20

[Link to Paper](#)

Neural Networks Are More Productive Teachers Than Human Raters: Active Mixup for Data-Efficient Knowledge Distillation from a Blackbox Model

Dongdong Wan, Yandong Li, Liqiang Wang, and **Boqing Gong**

Gong

12:20 & 00:20

[Link to Paper](#)

Google Landmarks Dataset v2 - A Large-Scale Benchmark for Instance-Level Recognition and Retrieval

Tobias Weyand, Andre Araujo, Jack Sim, Bingyi Cao

14:00 & 2:00

[Link to Paper](#)

CycleISP: Real Image Restoration via Improved Data Synthesis

Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, **Ming-Hsuan Yang**, Ling Shao

14:00 & 2:00

[Link to Paper](#)

Dynamic Graph Message Passing Networks

Li Zhang, Dan Xu, **Anurag Arnab**, Philip Torr

16:10 and 04:10

[Link to paper](#)

Posters

Single-view view synthesis with multiplane images

Richard Tucker, Noah Snavely

10:00 & 22:00

[Link to Paper](#)

Adversarial Examples Improve Image Recognition

Cihang Xie, Mingxing Tan, Boqing Gong, Jiang Wang,

Alan Yuille, **Quoc V. Le**

10:00 & 22:00

[Link to Paper](#)

Adversarial Texture Optimization from RGB-D Scans

Jingwei Huang, Justus Thies, Angela Dai, Abhijit Kundu, Chiyu "Max" Jiang, Leonidas Guibas, Matthias Niessner, Thomas Funkhouser

12:00 & 0:00

[Link to Paper](#)

Single-Image HDR Reconstruction by Learning to Reverse the Camera Pipeline

Yu-Lun Liu, **Wei-Sheng Lai**, Yu-Sheng Chen, Yi-Lung Kao, **Ming-Hsuan Yang**, Yung-Yu Chuang, Jia-Bin Huang

12:00 & 00:00

[Link to Paper](#)

Collaborative Distillation for Ultra-Resolution Universal Style Transfer

Huan Wang, Yijun Li, Yuehai Wang, Haoji Hu, **Ming-Hsuan Yang**

12:00 & 00:00

[Link to Paper](#)

Learning to Autofocus

Charles Herrmann, Richard Strong Bowen, **Neal Wadhwa, Rahul Garg, Qiurui He, Jonathan T. Barron, Ramin Zabih**

12:00 & 00:00

[Link to Paper](#)

Multi-Scale Boosted Dehazing Network with Dense Feature Fusion

Hang Dong, Jinshan Pan, Lei Xiang, Zhe Hu, Xinyi Zhang, Fei Wang, **Ming-Hsuan Yang**

12:00 & 0:00

[Link to Paper](#)

Composing Good Shots by Exploiting Mutual Relations

Debang Li, Junge Zhang, Kaiqi Huang, **Ming-Hsuan Yang**

16:00 & 4:00

[Link to Paper](#)

PatchVAE: Learning Local Latent Codes for Recognition

Kamal Gupta, **Saurabh Singh**, Abhinav Shrivastava

16:00 & 4:00

[Link to Paper](#)

Wednesday, June 17

Orals

Local Deep Implicit Functions for 3D Shape

Kyle Genova, Forrester Cole, Avneesh Sud, Aaron Sarna, Thomas Funkhouser

10:00 & 22:00

[Link to Paper](#)

GHUM & GHUML: Generative 3D Human Shape and Articulated Pose Models

Hongyi Xu, Eduard Gabriel Bazavan, Andrei Zanfir, William Freeman, Rahul Sukthankar, Cristian Sminchisescu

12:35 & 00:35

[Link to Paper](#)

Search to Distill: Pearls are Everywhere but not the Eyes

Yu Liu, Xuhui Jia, Mingxing Tan, Raviteja Vemulapalli, Yukun Zhu, Bradley Green, Xiaogang Wang

14:05 & 2:05

[Link to Paper](#)

Semantic Pyramid for Image Generation

Assaf Shocher, Yossi Gandelsman, Inbar Mosseri, Michal Yarom, Michal Irani, William Freeman, Tali Dekel

14:10 & 2:10

[Link to Paper](#)

Flow Contrastive Estimation of Energy-Based Models

Ruiqi Gao, Erik Nijkamp, Diederik Kingma, Zhen Xu, Andrew Dai, Ying Nian Wu

14:40 & 02:40

[Link to Paper](#)

Rethinking Class-Balanced Methods for Long-Tailed Visual Recognition from A Domain Adaptation Perspective

Muhammad Abdullah Jamal, Matthew Brown, Ming-Hsuan Yang, Liqiang Wang, Boqing Gong

14:40 & 02:40

[Link to Paper](#)

Category-Level Articulated Object Pose Estimation

Xiaolong Li, He Wang, Li Yi, Leonidas Guibas, Amos Abbott, Shuran Song

16:00 & 4:00

[Link to Paper](#)

AdaCoSeg: Adaptive Shape Co-Segmentation with Group Consistency Loss

Chenyang Zhu, Kai Xu, Siddhartha Chaudhuri, Li Yi, Leonidas Guibas, Hao Zhang

16:10 & 4:10

[Link to Paper](#)

Posters

Neural Voxel Renderer: Learning an Accurate and Controllable Rendering Tool

Konstantinos Rematas, Vittorio Ferrari

10:00 & 22:00

[Link to Paper](#)

Local Implicit Grid Representations for 3D Scenes

Chiyu "Max" Jiang, Avneesh Sud, Ameesh Makadia, Jingwei Huang, Matthias Niessner, Thomas Funkhouser

10:00 & 22:00

[Link to Paper](#)

Large Scale Video Representation Learning via Relational Graph Clustering

Hyodong Lee, Joonseok Lee, Joe Yue-Hei Ng, Apostol (Paul) Natsev

12:00 & 00:00

[Link to Paper](#)

Deep Homography Estimation for Dynamic Scenes

Hoang Le, Feng Liu, Shu Zhang, Aseem Agarwala

14:00 & 2:00

[Link to Paper](#)

C-Flow: Conditional Generative Flow Models for Images and 3D Point Clouds

Albert Pumarola, Stefan Popov, Francesc Moreno-Noguer, Vittorio Ferrari

14:00 & 2:00

[Link to Paper](#)

Lighthouse: Predicting Lighting Volumes for Spatially-Coherent Illumination

Pratul Srinivasan, Ben Mildenhall, Matthew Tancik, Jonathan T. Barron, Richard Tucker, Noah Snavely

14:00 & 2:00

[Link to Paper](#)

Scale-space flow for end-to-end optimized video compression

Eirikur Agustsson, David Minnen, Nick Johnston, Johannes Ballé, Sung Jin Hwang, George Toderici

14:00 & 2:00

[Link to Paper](#)

StructEdit: Learning Structural Shape Variations

Kaichun Mo, Paul Guerrero, Li Yi, Hao Su, Peter Wonka, Niloy Mitra, Leonidas Guibas

16:00 & 4:00

[Link to Paper](#)

3D-MPA: Multi Proposal Aggregation for 3D Semantic Instance Segmentation

Francis Engelmann, Martin Bokeloh, Alireza Fathi, Bastian Leibe, Matthias Niessner

16:00 4:00

[Link to Paper](#)

Sequential mastery of multiple tasks: Networks naturally learn to learn and forget to forget

Guy Davidson, Michael C. Mozer

16:00 & 4:00

[Link to Paper](#)

Distilling Effective Supervision from Severe Label Noise

Zizhao Zhang, Han Zhang, Sercan Ö. Arik, Honglak Lee, Tomas Pfister

16:00 & 4:00

[Link to Paper](#)

ViewAL: Active Learning With Viewpoint Entropy for Semantic Segmentation

Yawar Siddiqui, Julien Valentin, Matthias Niessner

16:00 & 4:00

[Link to Paper](#)

Attribution in Scale and Space

Shawn Xu, Subhashini Venugopalan, Mukund Sundararajan

16:00 & 4:00

[Link to Paper](#)

Weakly-Supervised Semantic Segmentation via Sub-category Exploration

Yu-Ting Chang, Qiaosong Wang, Wei-Chih Hung, Robinson Piramuthu, Yi-Hsuan Tsai, Ming-Hsuan Yang

16:00 & 4:00

[Link to Paper](#)

Thursday, June 18

Orals

SpeedNet: Learning the Speediness in Videos

Sagie Benaim, Ariel Ephrat, Oran Lang, Inbar Mosseri, William Freeman, Michael Rubinstein, Michal Irani, Tali Dekel

9:50 & 21:50

[Link to Paper](#)

BSP-Net: Generating Compact Meshes via Binary Space Partitioning

Zhiqin Chen, Andrea Tagliasacchi, Hao Zhang

10:20 & 22:20

[Link to Paper](#)

SAPIEN: A SimulATED Part-based Interactive ENvironment

Fanbo Xiang, Yuzhe Qin, Kaichun Mo, Yikuan Xia, Hao Zhu, Fangchen Liu, Minghua Liu, Hanxiao Jiang, Yifu Yuan, He Wang, Li Yi, Angel Chang, Leonidas Guibas, Hao Su

11:00 & 23:00

[Link to Paper](#)

SurfelGAN: Synthesizing Realistic Sensor Data for Autonomous Driving

Zhenpei Yang, Yuning Chai, Dragomir Anguelov, Yin Zhou, Pei Sun, Dumitru Erhan, Sean Rafferty, Henrik Kretzschmar

11:10 & 23:10

[Link to Paper](#)

Filter Response Normalization Layer: Eliminating Batch Dependence in the Training of Deep Neural Networks

Saurabh Singh, Shankar Krishnan

11:20 & 23:20

[Link to Paper](#)

RL-CycleGAN: Reinforcement Learning Aware Simulation-To-Real

Kanishka Rao, Chris Harris, Alex Irpan, Sergey Levine, Julian Ibarz, Mohi Khansari

11:30 & 23:30

[Link to Paper](#)

Open Compound Domain Adaptation

Ziwei Liu, Zhongqi Miao, Xingang Pan, Xiaohang Zhan, Dahua Lin, Stella X. Yu, and **Boqing Gong**

15:25 & 03:25

[Link to Paper](#)

Posters

Speech2Action: Cross-modal Supervision for Action Recognition

Arsha Nagrani, Chen Sun, David Ross, Rahul Sukthankar, Cordelia Schmid, Andrew Zisserman

9:00 & 21:00

[Link to Paper](#)

Counting Out Time: Class Agnostic Video Repetition Counting in the Wild

Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, Andrew Zisserman

9:00 & 21:00

[Link to Paper](#)

The Garden of Forking Paths: Towards Multi-Future Trajectory Prediction

Junwei Liang, Lu Jiang, Kevin Murphy, Ting Yu, Alexander Hauptmann

9:00 & 21:00

[Link to Paper](#)

Self-training with Noisy Student improves ImageNet classification

Qizhe Xie, Minh-Thang Luong, Eduard Hovy, Quoc V. Le

9:00 & 21:00

[Link to Paper](#)

EfficientDet: Scalable and Efficient Object Detection

Mingxing Tan, Ruoming Pang, Quoc Le

9:00 & 21:00

[Link to Paper](#)

ACNe: Attentive Context Normalization for Robust Permutation-Equivariant Learning

Weiwei Sun, Wei Jiang, **Eduard Trulls, Andrea Tagliasacchi**, Kwang Moo Yi

11:00 & 23:00

[Link to Paper](#)

VectorNet: Encoding HD Maps and Agent Dynamics from Vectorized Representation

Jiyang Gao, **Chen Sun**, Hang Zhao, Yi Shen, Dragomir Anguelov, **Cordelia Schmid**, Congcong Li

11:00 & 23:00

[Link to Paper](#)

SpineNet: Learning Scale-Permuted Backbone for Recognition and Localization

Xianzhi Du, Tsung-Yi Lin, Pengchong Jin, Golnaz Ghiasi, Mingxing Tan, Yin Cui, Quoc Le, Xiaodan Song

11:00 & 23:00

[Link to Paper](#)

KeyPose: Multi-View 3D Labeling and Keypoint Estimation for Transparent Objects

Xingyu Liu, Rico Jonschkowski, Anelia Angelova, Kurt Konolige

11:00 & 23:00

[Link to Paper](#)

Structured Multi-Hashing for Model Compression

Elad Eban, Yair Movshovitz-Attias, Hao Wu, Mark Sandler, Andrew Poon, Yerlan Idelbayev, Miguel A. Carreira-Perpinan

11:00 & 23:00

[Link to Paper](#)

DOPS: Learning to Detect 3D Objects and Predict their 3D Shapes

Mahyar Najibi, **Guangda Lai, Abhijit Kundu, Zhichao Lu, Vivek Rathod, Tom Funkhouser, Caroline Pantofaru, David Ross, Larry Davis, Alireza Fathi**

11:00 & 23:00

[Link to Paper](#)

Panoptic-DeepLab: A Simple, Strong, and Fast Baseline for Bottom-Up Panoptic Segmentation

Bowen Cheng, Maxwell Collins, Yukun Zhu, Ting Liu, Thomas S. Huang, Hartwig Adam, Liang-Chieh Chen

15:00 & 3:00

[Link to Paper](#)

Context R-CNN: Long Term Temporal Context for Per-Camera Object Detection

Sara Beery, Guanhang Wu, Vivek Rathod, Ronny Votel, Jonathan Huang

15:00 & 3:00

[Link to Paper](#)

Distortion Agnostic Deep Watermarking

Xiyang Luo, Ruohan Zhan, Huiwen Chang, Feng Yang, Peyman Milanfar

17:00 & 5:00

[Link to Paper](#)

Can weight sharing outperform random architecture search? An investigation with TuNAS

Gabriel Bender, Hanxiao Liu, Bo Chen, Grace Chu, Shuyang Cheng, Pieter-Jan Kindermans, Quoc Le

17:00 & 5:00

[Link to Paper](#)

GIFnets: Differentiable GIF Encoding Framework

Innfarn Yoo, Xiyang Luo, Yilin Wang, Feng Yang, Peyman Milanfar

17:00 & 5:00

6

Your Local GAN: Designing Two Dimensional Local Attention Mechanisms for Generative Models

Giannis Daras, Augustus Odena, Han Zhang, Alex Dimakis

17:00 & 5:00

[Link to Paper](#)

Fast Sparse ConvNets

Erich Eisen, Marat Dukhan, Trevor Gale, Karen Simonyan

17:00 & 5:00

[Link to Paper](#)

RetinaTrack: Online Single Stage Joint Detection and Tracking

Zhichao Lu, Vivek Rathod, Ronny Votel, Jonathan Huang

17:00 & 5:00

[Link to Paper](#)

Learning to See Through Obstructions

Yu-Lun Liu, Wei-Sheng Lai, Ming-Hsuan Yang, Yung-Yu Chuang, Jia-Bin Huang

17:00 & 5:00

[Link to Paper](#)

Self-Supervised Learning of Video-Induced Visual Invariances

Michael Tschannen, Josip Djolonga, Marvin Ritter, Aravindh Mahendran, Neil Houlsby, Sylvain Gelly, Mario Lucic

17:00 & 5:00

[Link to Paper](#)

Friday, June 19

Workshops

Extreme classification in computer vision

Ramin Zabih, Zhen Li, Manik Varma

June 19; Morning

[Workshop URL](#)

Image Matching: Local Features and Beyond

Vassileios Balntas, Vincent Lepetit, Jiri Matas, Dmytro Mishkin, Johannes Schoenberger, Eduard

Trulls, Kwang Moo Yi

June 19; Morning

[Workshop URL](#)

The DAVIS Challenge on Video Object Segmentation

Sergi Caelles, Alberto Montes, Federico Perazzi, Jordi

Pont-Tuset, Kevis-Kokitsi Maninis, Luc Van Gool

June 19; Morning

[Workshop URL](#)

2nd Workshop on Precognition: Seeing through the Future

Khoa Luu, Kris Kitani, Minh Hoai Nguyen, Hien Van Nguyen, Nemanja Djuric, Utsav Prabhu

June 19; Afternoon

[Workshop URL](#)

Computational Cameras and Displays (CCD)

Invited talk: "Computational Photography in Very Low Light" by **Orly Liba**

June 19; 03:00 - 04:30

[Workshop URL](#)

2nd Workshop on Learning from Unlabeled Videos (LUV)

Yale Song, Carl Vondrick, Katerina Fragkiadaki, Honglak Lee, Rahul Sukthankar

June 19: Full-Day

[Workshop URL](#)

7th Workshop on Fine Grained Visual Categorization (FGVC7)

*Oisin Mac Aodha, Ryan Farrell, Subhransu Maji, **Christine
Kaeser-Chen**, Grant Van Horn, **Serge Belongie***

June 19; Full-Day

[Workshop URL](#)

Language & Vision with applications to Video Understanding

*Qi Wu, Xin Wang, Chenxi Liu, Licheng Yu, **Lu Jiang**, Yan
Huang, Ting Yao, Qin Jin, William Wang, Anton van den
Hengel, Andrei Barbu, Siddharth N., Dan Gutfreund, Philip
Torr*

June 19; Full-Day

[Workshop URL](#)

Neural Architecture Search and Beyond for Representation Learning

*Rameswar Panda, Saining Xie, **Barret Zoph**, AJ
Piergiovanni, Bishwaranjan Bhattacharjee, Rogerio S. Feris,
Tianfu Wu, Ying Wu, Sinisa Todorovic, Ying Nian Wu, Arnav
Jhala, Caiming Xiong, David Xianfeng Gu, Song-Chun Zhu*

June 19; Full-Day

[Workshop URL](#)

Tutorials

Fairness Accountability Transparency and Ethics
and Computer Vision

Timnit Gebru, Emily Denton

June 19; Afternoon (Half-day)

[Tutorial URL](#)



@GoogleAI



ai.googleblog.com

Google Research