



**LONG BEACH
CALIFORNIA
June 16-20, 2019**



Welcome to CVPR 2019


General chairs



Larry Davis
Univ. of Maryland



Philip Torr
Univ. of Oxford



Song-Chun Zhu
UCLA

Program chairs



Abhinav Gupta
CMU



Derek Hoiem
UIUC



Gang Hua
Wormpex



Zhuowen Tu
UCSD

Finance Chairs



Brandon Rothrock
Jet Propulsion Laboratory



Walter Scheirer
University of Notre Dame

Workshop Chairs



Sanja Fidler
University of Toronto



Andrea Vedaldi
University of Oxford

Corporate Relations Chairs



Mei Han
Ping An Technology



Florent Perronnin
Naver Labs Europe

Website Chair



Ming-Ming Cheng
Nankai University

AC Meeting Chairs



Manmohan Chandraker
University of California,
San Diego



Hao Su
University of California,
San Diego

Tutorials Chairs



Ali Farhadi
University of Washington



M. Alex O. Vasilescu
University of California,
Los Angeles

Doctoral Consortium Chairs



Olga Russakovsky
Princeton University




Nathan Jacobs
University of Kentucky

Volunteer chair



Yixin Zhu
University of California,
Los Angeles

Publications Chairs



William Brendel
Snap Inc.



Mohamed R. Amer
SRI International




Eric Mortensen
Lucidyne Technologies

Demo & Exhibition Chairs



Andrew D. Bagdanov
Università degli Studi di
Firenze



Yong Jae Lee
University of California,
Davis

Presentation Chairs



Maria Zontak
Microsoft



Tianfu Wu
NC State University

Ethics chairs and ombudsmen

- Ethics chairs: contact by email with concerns about unwelcoming, offensive, or unethical behavior at the conference



Derek Hoiem

Alex Vasilescu

- PAMI TC Ombudsmen: long-term role to ensure continued propriety of the conference and its organization
 - David Forsyth (University of Illinois)
 - Linda Shapiro (University of Washington)

Thank you sponsors!

284 Sponsors

104 Exhibitors

\$3.1M in

Sponsorship &
Exhibition




CVPR'19 by the Numbers

5,160 Paper Submissions


1,294 Papers Accepted

9,227 Registered Attendees

CVPR Attendance Trend




~9,200 attendees from 68 countries




Wednesday Night Event

Attendees **MUST** wear badges and only **Full Passport Registrations** (blue lanyards) are able to attend.

Drink tickets were provided with badge.




- All catered functions (including Wednesday) are outside – dress accordingly
- Breaks in Exhibit Hall
- Meals picked up in Pacific Ballroom
 - Seating in patios and Hall A
- Thursday meals picked up in Hall A




CVPR Submitted and Accepted Papers


56% yearly growth with 26% acceleration → 10.8B submitted papers in 2028




14,104 authors submitted 5,160 papers



2,887 reviewers provided 15,000+ reviews




Reviewers are experienced, 2+ papers authored



132 Area Chairs: balance topics, demographics

- 20 from institutions in Asia
- 22 women
- ~20% first time ACs
- Good topic distribution (matches submissions)



1,294 papers in CVPR'19 (25.2% acceptance rate)

Acceptance rate is roughly even across topics

action adaptation adversarial attention based clouds convolutional
data deep depth detection domain efficient estimation face
feature generative graph human image instance joint
learning local matching model motion network
neural object person point pose prediction recognition reconstruction
representation robust scene segmentation semantic shape
single structure supervised tracking transfer unsupervised video visual


New in 2019: short orals

288 short (5 min) oral presentations

- Groups of 3, organized by topic, followed by 3 min questions

New tool to find talks/posters of interest

<https://public.tableau.com/views/CVPR2019all/DashboardCVPR>



Thanks to
GeorgiaTech
Communications
Team

CVPR 2019 Awards

CVPR 2019 Best Paper Award Committee



Terry Boulton



Michael Brown



Kristin Dana



Sven Dickinson



Yasuyuki Matsushita



Greg Mori
(Chair)



Ian Reid



Lihi Zelnik-Manor

Thank you for your service!

Best Paper Honorable Mention

A Style-Based Generator Architecture for Generative Adversarial Networks

Tero Karras, Samuli Laine, Timo Aila

“Advances in generative models have come at a break-neck pace in recent years. This paper continues this push toward generation of highly realistic face images, particularly providing advances in controllable generation and interpolation across attributes.”

- Award Committee

Best Paper Honorable Mention

Learning the Depths of Moving People by Watching Frozen People

**Zhengqi Li, Tali Dekel, Forrester Cole, Richard Tucker,
Ce Liu, Bill Freeman, Noah Snavely**

“This paper examines 3d reconstruction of scenes with people from monocular video. The paper shows insightful creation of a dataset for this task (Mannequin challenge), in addition to a solid execution of a state of the art algorithm. It has strong potential to be an impactful paper in this area, facilitating future work in outdoor reconstruction with complex moving scenes.” – Award Committee

Best Student Paper Award

Reinforced Cross-Modal Matching and Self-Supervised Imitation Learning for Vision-Language Navigation

**Xin Wang, Qiuyuan Huang, Asli Celikyilmaz, Jianfeng Gao,
Dinghan Shen, Yuan-Fang Wang, William Yang Wang, Lei Zhang**

“Visual navigation is an important area of computer vision -- this paper makes advances in vision-language navigation. Building on previous work in this area, this paper demonstrates exciting results based on self-imitation learning within a cross-modal setting.”

-- Award Committee

Best Paper Award

A Theory of Fermat Paths for Non-Line-of-Sight Shape Reconstruction

**Shumian Xin, Sotiris Nousias, Kyros Kutulakos,
Aswin Sankaranarayanan, Srinivasa G. Narasimhan, Ioannis Gkioulekas**

“This paper makes significant advances in non-line-of-sight reconstruction -- in essence the ability to see around corners. It is a beautiful paper theoretically, as well as inspiring. It continues to push the boundaries of what is possible in computer vision.”

-- Award Committee

Best paper finalists

“Probabilistic Permutation Synchronization using the Riemannian Structure of the Birkhoff Polytope,” by Tolga Birdal; Umut Simsekli

“DeepCO³ : Deep Instance Co-segmentation by Co-peak Search and Co-saliency Detection,” by Kuang-Jui Hsu; Yen-Yu Lin; Yung-Yu Chuang

“SelfFlow: Self-Supervised Learning of Optical Flow,” by Pengpeng Liu, Michael Lyu, Irwin King, Jia Xu

“SDRSAC: Semidefinite-Based Randomized Approach for Robust Point Cloud Registration without Correspondences,” by Huu Le; Thanh-Toan Do; Tuan NA Hoang; Ngai-Man Cheung

“Deep Tree Learning for Zero-shot Face Anti-Spoofing,” by Yaojie Liu; Joel Stehouwer; Amin Jourabloo; Xiaoming Liu

“Neural RGB -> D Sensing: Depth and Uncertainty from a Video Camera,” by Chao Liu; Jinwei Gu; Kihwan Kim; Srinivasa G Narasimhan; Jan Kautz

“Self-supervised 3D hand pose estimation through training by fitting,” by Chengde Wan; Thomas Probst; Luc Van Gool; Angela Yao

“Neural Illumination: Lighting Prediction for Indoor Environments,” by Shuran Song; Thomas Funkhouser

“Shapes and Context: In-the-wild Image Synthesis & Manipulation,” Aayush Bansal; Yaser Sheikh; Deva Ramanan

“SiCloPe: Silhouette-based Clothed People,” by Ryota Natsume; Shunsuke Saito; Zeng Huang; Weikai Chen; Chongyang Ma; Shigeo Morishima; Hao Li

“A General and Adaptive Robust Loss Function,” by Jonathan T Barron

“2.5D Visual Sound,” by Ruohan Gao; Kristen Grauman

“Incremental Object Learning from Contiguous Views,” by Stefan Stojanov; Samarth Mishra; Ngoc Anh Thai; Nikhil Dhanda; Ahmad Humayun; Linda Smith; Chen Yu; James Rehg

“Text2Scene: Generating Compositional Scenes from Textual Descriptions,” by Fuwen Tan; Song Feng; Vicente Ordonez

“Relational Action Forecasting,” by Chen Sun; Abhinav Shrivastava; Carl Vondrick; Rahul Sukthankar; Kevin Murphy; Cordelia Schmid

“Shifting More Attention to Video Salient Object Detection,” by Deng-Ping Fan; Wenguan Wang; Ming-Ming Cheng; Jianbing Shen

“GA-Net: Guided Aggregation Net for End-to-end Stereo Matching,” by Feihu Zhang; Victor Prisacariu; Yang Ruigang; Philip Torr

“A Skeleton-bridged Deep Learning Approach for Generating Meshes of Complex Topologies from Single RGB Images,” by Jiapeng Tang; Xiaoguang Han; Junyi Pan; Kui Jia; Xin Tong

“Semantic Image Synthesis with Spatially-Adaptive Normalization,” by Taesung Park; Ming-Yu Liu; Ting-Chun Wang; Jun-Yan Zhu

“ContactDB: Analyzing and Predicting Grasp Contact via Thermal Imaging,” by Samarth Brahmbhatt; Cusuh Ham; Charlie Kemp; James Hays

“Revealing Scenes by Inverting Structure from Motion Reconstructions,” by Francesco Pittaluga; Sanjeev J Koppal; Sing Bing Kang; Sudipta Sinha

“A Theory of Fermat Paths for Non-Line-of-Sight Shape Reconstruction,” by Shumian Xin; Sotiris Nousias; Kyros Kutulakos; Aswin Sankaranarayanan; Srinivasa G Narasimhan; Ioannis Gkioulekas

“Relation-Shape Convolutional Neural Network for Point Cloud Analysis,” by Yongcheng Liu; Bin Fan; Shiming Xiang; Chunhong Pan

“BubbleNets: Learning to Select the Guidance Frame in Video Object Segmentation by Deep Sorting Frames,” by Brent Griffin; Jason J Corso

Best paper finalists

“Image Deformation Meta-Networks for One-Shot Learning,” by Zitian Chen; Yanwei Fu; Yu-Xiong Wang; Lin Ma; Wei Liu; Martial Hebert

“Estimating 3D Motion and Forces of Person-Object Interactions from Monocular Video,” by Zongmian Li; Jiri Sedlar; Justin Carpentier; Ivan Laptev; Nicolas Mansard; Josef Sivic

“A Style-Based Generator Architecture for Generative Adversarial Networks,” by Tero Karras; Samuli Laine; Timo Aila

“Unsupervised Part-Based Disentangling of Object Shape and Appearance,” by Dominik Lorenz; Leonard Bereska; Timo Milbich; Bjorn Ommer

“Pushing the Boundaries of View Extrapolation with Multiplane Images,” Pratul Srinivasan; Richard Tucker; Jonathan T Barron; Ravi Ramamoorthi; Ren Ng; Noah Snavely

“Path-Invariant Map Networks,” Zaiwei Zhang; Zhenxiao Liang; Lemeng Wu; Xiaowei Zhou; Qixing Huang

“Learning the Depths of Moving People by Watching Frozen People,” Zhengqi Li; Tali Dekel; Forrester Cole; Richard Tucker; Ce Liu; Bill Freeman; Noah Snavely

“Efficient Online Multi-Person 2D Pose Tracking with Recurrent Spatio-Temporal Affinity Fields,” Yaadhav Raaj; Haroon Idrees; Gines Hidalgo Martinez; Yaser Sheikh

“Learning to Compose Dynamic Tree Structures for Visual Contexts,” Kaihua Tang; Hanwang Zhang; Baoyuan Wu; Wenhan Luo; Wei Liu

“Cascaded Projection: End-to-End Network Compression and Acceleration,” Breton L Minnehan; Andreas Savakis

“Taking a Deeper Look at the Inverse Compositional Algorithm,” Zhaoyang Lv; Frank Dellaert; James Rehg; Andreas Geiger

“Occupancy Networks: Learning 3D Reconstruction in Function Space,” Lars M Mescheder; Michael Oechsle; Michael Niemeyer; Sebastian Nowozin (Google AI Berlin); Andreas Geiger

“Geometry-Consistent Generative Adversarial Networks for One-Sided Unsupervised Domain Mapping,” by Huan Fu; Mingming Gong; Chaohui Wang; Kayhan Batmanghelich; Kun Zhang; Dacheng Tao

“Convolutional Mesh Regression for Single-Image Human Shape Reconstruction,” Nikos Kolotouros; Georgios Pavlakos; Kostas Daniilidis

“Neural Rerendering in the Wild,” Moustafa Meshry; Ricardo Martin-Brualla; Noah Snavely; Hugues Hoppe; Sameh Khamis; Rohit Pandey; Dan B Goldman

“Content Authentication for Neural Imaging Pipelines: End-to-end Optimization of Photo Provenance in Complex Distribution Channels,” Pawel Korus; Nasir Memon

“Reinforced Cross-Modal Matching and Self-Supervised Imitation Learning for Vision-Language Navigation,” Xin Wang; Qiuyuan Huang; Asli Celikyilmaz; Jianfeng Gao; Dinghan Shen; Yuan-Fang Wang; William Yang Wang; Lei Zhang

“FilterReg: Robust and Efficient Probabilistic Point-Set Registration using Gaussian Filter and Twist Parameterization,” Wei Gao; Russ Tedrake

“Locating Objects Without Bounding Boxes,” Javier Ribera; David Güera; Yuhao Chen; Edward Delp

“DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation,” Jeong Joon Park; Peter R Florence; Julian Straub; Richard Newcombe; Steven Lovegrove

“CollaGAN: Collaborative GAN for Missing Image Data Imputation,” Dongwook Lee; Junyoung Kim; Won-Jin Moon; Jong Chul Ye

PAMI TC Awards

PAMI Longuet-Higgins Prize

Retrospective Most Impactful Paper from CVPR 2009

Awards Committee:

- Cordelia Schmid (chair)
- Horst Bischof
- Jitendra Malik
- Josef Sivic

PAMI Longuet-Higgins Prize

Retrospective Most Impactful Paper from CVPR 2009

ImageNet: A large-scale hierarchical image database

Jia Deng, Wei Dong, Richard Socher,
Li-Jia Li, Kai Li, and Li Fei-Fei

PAMI Young Researcher Award

Sponsored by *Image and Vision Computing* (Elsevier)

Awards Committee:

- Ramin Zabih (chair)
- Andrew Fitzgibbon
- Kristen Grauman
- Maja Pantic
- Nikos Paragios
- Long Quan

PAMI Young Researcher Award

Sponsored by *Image and Vision Computing* (Elsevier)

Karen Simonyan

IEEE Awards



2019 COMPUTER PIONEER AWARD RECIPIENT

Jitendra Malik

For a leading role in developing computer vision into a thriving discipline through pioneering research, leadership, and mentorship.



IEEE
COMPUTER
SOCIETY



IEEE



**LONG BEACH
CALIFORNIA
June 16-20, 2019**



Enjoy the Conference!