2021 ICCV OCTOBER 11-17 VIRTUAL

ICCV 2021 Prizes

Dima Damen
University of Bristol

Tal Hassner
Facebook AI

Christopher Pal
Ecole Polytechnique
de Montreal

Yoichi Sato
University of Tokyo

ICCV 2021 Prizes

- Outstanding Reviewers

- Marr Prize - ICCV 2021 Best Paper
  - and honourable mentions

- Best Student Paper

Outstanding Reviewers

- Reviewers for which there was overwhelming evidence for outstanding contribution to reviewing.
- Ratings excluded papers withdrawn during the rebuttal phase
- 18% of all reviews were ranked as exceeding expectations
- Multiple “exceed expectations” ratings
- 5% of all experienced reviewers & 5% of all student reviewers

Outstanding Reviewers - ICCV 2021

ICCV2021 @ICCV_2021 · Sep 1

[PCs Update] We acknowledge 210 outstanding reviewers (top 5% experienced and top 5% student reviewers) online at:
iccv2021.thecvf.com/outstanding-re...
amongst the many amazing reviewers this year.
We also ack. generous emergency reviewers:
iccv2021.thecvf.com/emergency-revi...

5 34 120

100 free registations

Outstanding Reviewers - ICCV 2021

Abdullah Abuolaim
Mahmoud Afifi
Samuel Albanie
Cenek Albl
Jose Alvarez
Relja Arandjelović
Nikita Araslanov
Pablo Arbelaez
Muhammad Asad
Yuki Asano
Nicolas Audebert
Melinos Averkiou
Angelica Aviles-Rivero
Yannis Avrithis
Jonathan Barron
Miguel Angel Bautista
Jens Behley
Assia Benbihi
Alexander Bergman
Amit Bermano
Timo Bolkart
Amine Bourki
Ivaylo Boyadzhiev
Eric Brachmann
Robert-Jan Bruintjes
Zhongang Cai
Luca Carlone
Arantxa Casanova
Fabio Cermelli
Ayan Chakrabarti
Kelvin C.K. Chan
Min-Hung Chen
Dongdong Chen
Jiefeng Chen
Yixin Chen
Ming-Ming Cheng
Myungsub Choi
Daniel Cremers
Gabriela Csurka
Martin Danelljan
Abir Das
Adrien Deliege
Boyang Deng
Xiaoyi Dong
Bardia Doosti
Hazel Doughty
Michal Drozdzal
Amanda Duarte
Mohamed El Banani
Ismail Elezi
Qianli Feng
Victoria Fernandez Abrevaya
Claudio Ferrari
Chi-Wing Fu
Cheng-Yang Fu
Yasuhisa Fujii
Antonino Furnari
Akshay Gadi Patil
Rinon Gal
Orazio Gallo
Elena Garces
Pablo Garrido
Ioannis Gkioulekas
Ankit Goyal
Colin Graber
Fatma Guney
Ankush Gupta
Bumsub Ham
Xintong Han
Kai Han
Yana Hasson
Michael Hofmann
Yicong Hong
Xiaolin Hu
Qingyong Hu
Junhwa Hur
Jaedong Hwang
Mona Jalal
Vazquez-Corral Javier
Vicky Kalogeiton
Corentin Kervadec
Boris Knyazev
Praveen Krishnan
Alexander Krull
Jean-Francois Lalonde
Vuong Le
Xiang Li
Chongyi Li
Yunzhu Li
Tianye Li
Zhengqin Li
Kevin Liang
Jun Hao Liew
Zhe Lin
Yonghuai Liu
Yu Liu
Juncheng Liu
Juwei Lu
Oisin Mac Aodha
Massimiliano Mancini
Kevis-Kokitsi Maninis
Renaud Marlet
Stefan Mathe
Minesh Mathew
Yusuke Matsui
Efi Mavroudi
Juhong Min
Anand Mishra
Gaurav Mittal
Philippos Mordohai
Francesc Moreno
Jonathan Munro
Ana Murillo
Seungjun Nah
Seonghyeon Nam
Sanath Narayan
Simon Niklaus
Yulei Niu
David Novotny
Anton Obukhov
Jihyong Oh
Utkarsh Ojha
Mohamed Omran
Jose Oramas
Aljosa Osep
Cheng Ouyang
Liang Pan
Liyuan Pan
Andreas Panteli
Seonwook Park
JaeYoo Park
Sujoy Paul
Georgios Pavlakos
Adithya Pediredla
Xingchao Peng
Juan Perez
Stavros Petridis
Khoi Pham
Silvia Pintea
Soeren Pirk
Bryan Plummer
Adria Recasens
Zhongzheng Ren
Christian Richardt
Daniel Ritchie
Marcus Rohrbach
Rob Romijnders
Adria Ruiz
Christian Rupprecht
Bryan Russell
Aniruddha Saha
Samuele Salti
Enrique Sanchez
Nikolaos Sarafianos
Saquib Sarfraz
Paul-Edouard Sarlin
Hanno Scharr
Jie Shen
Liyue Shen
Yucong Shen
Assaf Shocher
Abhinav Shrivastava
Leonid Sigal
Jeany Son
Jin Sun
Anshuman Suri
Ajinkya Tejankar
Christopher Thomas
Federico Tombari
Nergis Tomen
Tatiana Tommasi
Fabio Tosi
Eduard Trulls
Yi-Hsuan Tsai
Aggeliki Tsoli
Daniyar Turmukhambetov
Jack Valmadre
Nanne van Noord
Sai Vemprala
Luisa Verdoliva
Tomas Vojir
Konstantinos Vougioukas
Guangrun Wang
Wenguan Wang
Dongdong Wang
Guangting Wang
Anne Wannenwetsch
Luca Weihs
Martin Weinmann
Davis Wertheimer
Olivia Wiles
Michael Wray
Donglai Xiang
Chaochao Yan
Jingkang Yang
Quanming Yao
Kwang Moo Yi
Xin Yu
Sangdoo Yun
Xingyu ZENG
Kuo-Hao Zeng
ZEKUN ZHANG
Yifan Zhao
Chenglong Zhao
Kaiyang Zhou
Mo Zhou
Yuqian Zhou
Jun-Yan Zhu
Maria Zontak
Maria A. Zuluaga

Marr Prize

- 13 papers were selected
- Nominated by two reviewers and 1 area chair
- Committee selected with no conflict, focusing on diversity

ICCV 2021 Marr Prize Committee

- Greg Mori (chair, Simon Fraser University, Canada)
- Tinne Tuytelaars (KU Leuven, Belgium)
- Kyong Mu Lee (Seoul National University, South Korea)
- Richa Singh (IIT Jodhpur, India)
- Xiaoming Liu (Michigan State University, US)
- Kosta Derpanis (York University, Canada)
- Barbara Caputo (Politechnico di Torino, Italy)

Nominated papers

- 4 papers for honourable mention
- 1 paper for Best Student Prize
- 1 Marr Prize - Best Paper Prize

Honorable Mentions

Mip-NeRF: A Multiscale Representation for Anti-Aliasing Neural Radiance Fields

Jonathan T Barron, Ben Mildenhall (Google Research). Matthew Tancik (UC Berkeley), Peter Hedman (Google Research), Ricardo Martin-Brualla (Google), Pratul Srinivasan (Google Research)

Session 5 (A/B)

Mip-NeRF: A Multiscale Representation for Anti-Aliasing Neural Radiance Fields

Jonathan T. Barron[1] Ben Mildenhall[1] Matthew Tancik[2]
Peter Hedman[1] Ricardo Martin-Brualla[1] Pratul P. Srinivasan[1]
[1]Google [2]UC Berkeley

Abstract

The rendering procedure used by neural radiance fields (NeRF) samples a scene with a single ray per pixel and may therefore produce renderings that are excessively blurred or aliased when training or testing images observe scene content at different resolutions. The straightforward solution of supersampling by rendering with multiple rays per pixel is impractical for NeRF, because rendering each ray requires querying a multilayer perceptron hundreds of times. Our solution, which we call "mip-NeRF" (à la "mipmap"), extends NeRF to represent the scene at a continuously-valued scale. By efficiently rendering anti-aliased conical frustums instead of rays, mip-NeRF reduces objectionable aliasing artifacts and significantly improves NeRF's ability to represent fine details, while also being 7% faster than NeRF and half the size. Compared to NeRF, mip-NeRF reduces aver-

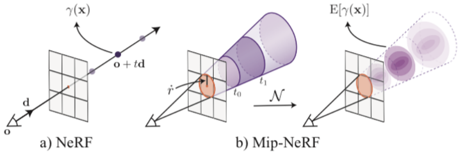

Figure 1: NeRF (a) samples points $\mathbf{x}$ along rays that are traced from the camera center of projection through each pixel, then encodes those points with a positional encoding (PE) γ to produce a feature $\gamma(\mathbf{x})$. Mip-NeRF (b) instead reasons about the 3D *conical frustum* defined by a camera pixel. These conical frustums are then featurized with our integrated positional encoding (IPE), which works by approximating the frustum with a multivariate Gaussian and then computing the (closed form) integral $E[\gamma(\mathbf{x})]$ over the positional encodings of the coordinates within the Gaussian.

2021 ICCV OCTOBER 11-17 VIRTUAL

Honorable Mentions

OpenGAN: Open-Set Recognition via Open Data Generation

Shu Kong, Deva Ramanan (Carnegie Mellon University)

Session 1 (A/B)

OpenGAN: **Open-Set Recognition via Open Data Generation**

Shu Kong*, Deva Ramanan*,†
*Carnegie Mellon University †Argo AI
{shuk, deva}@andrew.cmu.edu

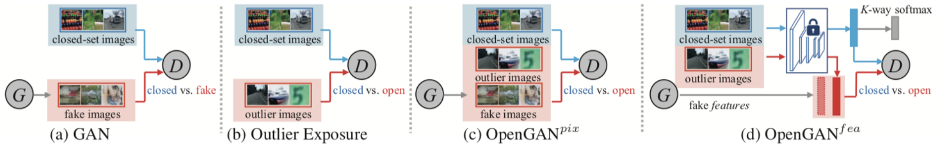

Figure 1: We explore open-set recognition, which requires the ability to discriminate open-set test examples outside K classes of interest. **(a)** Past work has suggested that GAN discriminators can serve as open-set likelihood functions, but this does not work well due to instable training of GANs [47, 44, 39, 56, 30]. **(b)** Outlier Exposure [25] exploits some outlier data to learn a binary discriminator D for open-set discrimination. Because outliers observed during training will not exhaustively span the open-world, the discriminator D tends to generalize poorly to diverse open data [48]. **(c)** We introduce OpenGAN, which augments training outliers with *fake* open data synthesized by a generator G trained to fool the discriminator D. Importantly, we find that a small number of outliers stabilizes training by enabling effective model selection of the discriminator D. **(d)** Because we are concerned with accurate discrimination rather than realistic pixel generation, we find it more efficient to generate (and discriminate) *features* from the off-the-shelf K-way classification network. This allows OpenGAN to be implemented via a lightweight discriminator head built on top of an existing K-way network, enabling closed-world systems to be readily modified for open-set recognition.

Honorable Mentions

Viewing Graph Solvability via Cycle Consistency

Federica Arrigoni (University of Trento), Andrea Fusiello (UNIUD), Elisa Ricci (University of Trento), Tomas Pajdla (Czech Technical University in Prague)

Session 5 (A/B)

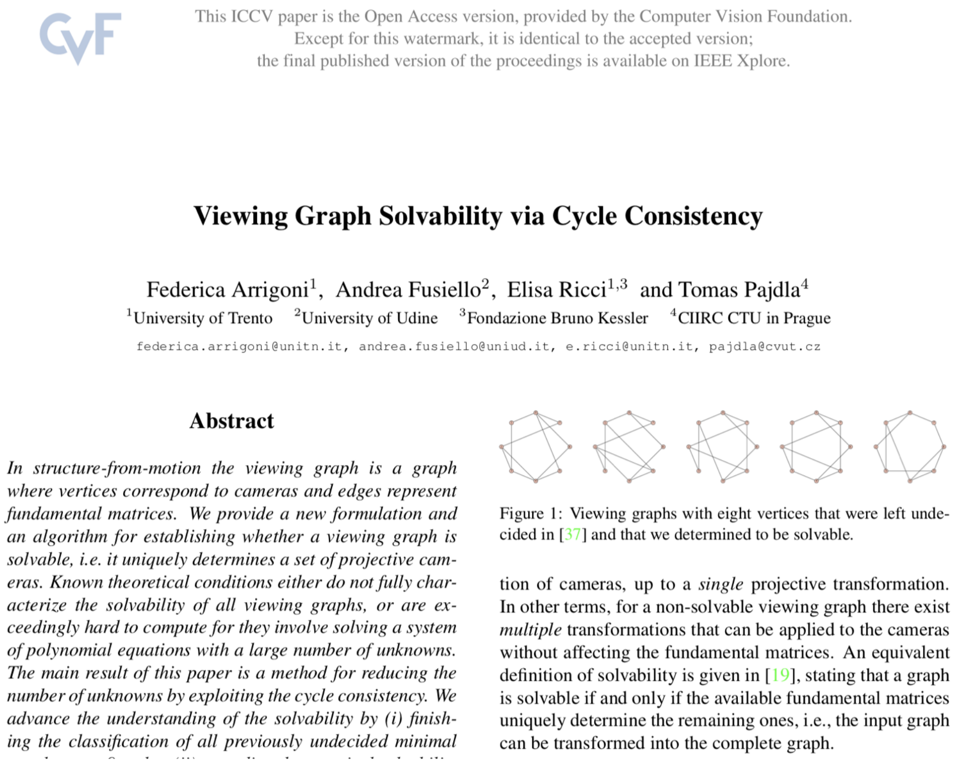

Viewing Graph Solvability via Cycle Consistency

Federica Arrigoni[1], Andrea Fusiello[2], Elisa Ricci[1,3] and Tomas Pajdla[4]
[1]University of Trento [2]University of Trento [3]Fondazione Bruno Kessler [4]CIIRC CTU in Prague
federica.arrigoni@unitn.it, andrea.fusiello@uniud.it, e.ricci@unitn.it, pajdla@cvut.cz

Abstract

In structure-from-motion the viewing graph is a graph where vertices correspond to cameras and edges represent fundamental matrices. We provide a new formulation and an algorithm for establishing whether a viewing graph is solvable, i.e. it uniquely determines a set of projective cameras. Known theoretical conditions either do not fully characterize the solvability of all viewing graphs, or are exceedingly hard to compute for they involve solving a system of polynomial equations with a large number of unknowns. The main result of this paper is a method for reducing the number of unknowns by exploiting the cycle consistency. We advance the understanding of the solvability by (i) finishing the classification of all previously undecided minimal

Figure 1: Viewing graphs with eight vertices that were left undecided in [37] and that we determined to be solvable.

tion of cameras, up to a *single* projective transformation. In other terms, for a non-solvable viewing graph there exist *multiple* transformations that can be applied to the cameras without affecting the fundamental matrices. An equivalent definition of solvability is given in [19], stating that a graph is solvable if and only if the available fundamental matrices uniquely determine the remaining ones, i.e., the input graph can be transformed into the complete graph.

Honorable Mentions

Common Objects in 3D: Large-Scale Learning and Evaluation of Real-life 3D Category Reconstruction

Jeremy Reizenstein (Facebook AI Research), Philipp Henzler (University College London), Roman Shapovalov, Luca Sbordone, Patrick Labatut, David Novotny (Facebook AI Research)

Session 8 (A/B)

Common Objects in 3D:
Large-Scale Learning and Evaluation of Real-life 3D Category Reconstruction

Jeremy Reizenstein[1] Roman Shapovalov[1] Philipp Henzler[2] Luca Sbordone[1]
Patrick Labatut[1] David Novotny[1]

{reizenstein,romansh,lsbordone,plabatut,dnovotny}@fb.com {p.henzler}@cs.ucl.ac.uk

[1]Facebook AI Research [2]University College London

https://github.com/facebookresearch/co3d

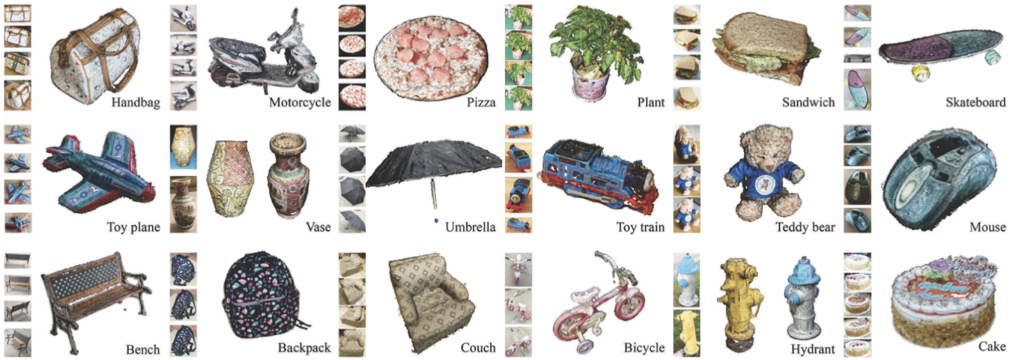

Figure 1: We introduce the **Common Objects in 3D (CO3D)** dataset comprising 1.5 million multi-view images of almost 19k objects from 50 MS-COCO categories annotated with accurate cameras and 3D point clouds (visualized above).

Best Student Paper Award

Pixel-Perfect Structure-from-Motion with Featuremetric Refinement

Philipp Lindenberger, Paul-Edouard Sarlin, Viktor Larsson (ETH Zurich),
Marc Pollefeys (ETH Zurich / Microsoft)

Session 5 (A/B)

2021 ICCV OCTOBER 11-17 VIRTUAL

Pixel-Perfect Structure-from-Motion with Featuremetric Refinement

Philipp Lindenberger[1*] Paul-Edouard Sarlin[2*] Viktor Larsson[2] Marc Pollefeys[2,3]
Departments of [1]Mathematics and [2]Computer Science, ETH Zurich [3]Microsoft

Abstract

Finding local features that are repeatable across multiple views is a cornerstone of sparse 3D reconstruction. The classical image matching paradigm detects keypoints per-image once and for all, which can yield poorly-localized features and propagate large errors to the final geometry. In this paper, we refine two key steps of structure-from-motion by a direct alignment of low-level image information from multiple views: we first adjust the initial keypoint locations prior to any geometric estimation, and subsequently refine points and camera poses as a post-processing. This refinement is robust to large detection noise and appearance changes, as it optimizes a featuremetric error based on dense features

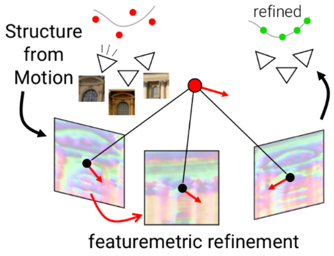

Marr Prize

Swin Transformer: Hierarchical Vision Transformer using Shifted Windows

Ze Liu (USTC), Yutong Lin (Xi'an Jiaotong University),
Yue Cao (Microsoft Research), Han Hu (Microsoft Research Asia),
Yixuan Wei (Tsinghua University), Zheng Zhang (MSRA, Huazhong University of Science and Technolog), Stephen Lin (Microsoft Research),
Baining Guo (MSR Asia)

Session 8 (A/B)


Congratulations