2024 International Conference on 3D Vision (3DV 2024)

Davos, Switzerland
18-21 March 2024

Pages 1-548

IEEE Catalog Number: CFP24203-POD
ISBN: 979-8-3503-6246-6

Copyright © 2024 by the Institute of Electrical and Electronics Engineers, Inc.
All Rights Reserved

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

For other copying, reprint or republication permission, write to IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854. All rights reserved.

****** This is a print representation of what appears in the IEEE Digital Library. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number:	CFP24203-POD
ISBN (Print-On-Demand):	979-8-3503-6246-6
ISBN (Online):	979-8-3503-6245-9
ISSN:	2378-3826

Additional Copies of This Publication Are Available From:

Curran Associates, Inc
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: (845) 758-0400
Fax: (845) 758-2633
E-mail: curran@proceedings.com
Web: www.proceedings.com

2024 International Conference on 3D Vision (3DV)

3DV 2024

Table of Contents

Welcome Message ..... **xxiv**
Conference Committee ..... **xxvi**
Reviewers ..... **xxviii**

Orals

Oral 1

FoVA-Depth: Field-of-View Agnostic Depth Estimation for Cross-Dataset Generalization 1
Daniel Lichy (University of Maryland, USA; NVIDIA), Hang Su (NVIDIA), Abhishek Badki (NVIDIA), Jan Kautz (NVIDIA), and Orazio Gallo (NVIDIA)

LFM-3D: Learnable Feature Matching Across Wide Baselines Using 3D Signals 11
Arjun Karpur (Google Research), Guilherme Perrotta (University of Campinas), Ricardo Martin-Brualla (Google Research), Howard Zhou (Google Research), and André Araujo (Google Research)

SCENES: Subpixel Correspondence Estimation With Epipolar Supervision 21
Dominik Alexander Kloepfer (University of Oxford), João F. Henriques (University of Oxford), and Dylan Campbell (The Australian National University)

Oral 2

Few-View Object Reconstruction with Unknown Categories and Camera Poses 31
Hanwen Jiang (UT Austin), Zhenyu Jiang (UT Austin), Kristen Grauman (UT Austin), and Yuke Zhu (UT Austin)

NICER-SLAM: Neural Implicit Scene Encoding for RGB SLAM 42
Zihan Zhu (ETH Zurich), Songyou Peng (ETH Zurich; MPI for Intelligent Systems, Tubingen), Viktor Larsson (Lund University, Sweden), Zhaopeng Cui (Zhejiang University), Martin R. Oswald (University of Amsterdam; ETH Zurich), Andreas Geiger (University of Tübingen, Tübingen AI Center), and Marc Pollefeys (ETH Zurich; Microsoft)

RaNeuS: Ray-Adaptive Neural Surface Reconstruction 53
Yida Wang (Technical Universität of Münich), David Joseph Tan (Google), Nassir Navab (Technical Universität of Münich), and Federico Tombari (Technical Universität of Münich; Google)

SlimmeRF: Slimmable Radiance Fields 64
Shiran Yuan (Tsinghua University, China; Duke Kunshan University, China; Duke University, USA) and Hao Zhao (Tsinghua University, China)

Oral 3

Fusing Directions and Displacements in Translation Averaging 75
Lalit Manam (Indian Institute of Science, India) and Venu Madhav Govindu (Indian Institute of Science, India)

POCO: 3D Pose and Shape Estimation with Confidence 85
Sai Kumar Dwivedi (Max Planck Institute for Intelligent Systems, Germany), Cordelia Schmid (PSL Research University, France), Hongwei Yi (Max Planck Institute for Intelligent Systems, Germany), Michael J. Black (Max Planck Institute for Intelligent Systems, Germany), and Dimitrios Tzionas (University of Amsterdam, the Netherlands)

Relative Pose for Nonrigid Multi-Perspective Cameras: The Static Case 96
Min Li (ShanghaiTech University, China), Jiaqi Yang (ShanghaiTech University, China), and Laurent Kneip (ShanghaiTech University, China)

RelPose++: Recovering 6D Poses from Sparse-View Observations 106
Amy Lin (Carnegie Mellon University, USA), Jason Y. Zhang (Carnegie Mellon University, USA), Deva Ramanan (Carnegie Mellon University, USA), and Shubham Tulsiani (Carnegie Mellon University, USA)

Oral 4

GAPS: Geometry-Aware, Physics-Based, Self-Supervised Neural Garment Draping 116
Ruochen Chen (École Centrale de Lyon, France; LIRIS, France), Liming Chen (École Centrale de Lyon, France; LIRIS, France), and Shaifali Parashar (INSA-Lyon, France)

Multi-Body Neural Scene Flow 126
Kavisha Vidanapathirana (Queensland University of Technology; CSIRO Robotics), Shin-Fang Chng (The University of Adelaide), Xueqian Li (The University of Adelaide), and Simon Lucey (The University of Adelaide)

Objects With Lighting: A Real-World Dataset for Evaluating Reconstruction and Rendering for Object Relighting 137
Benjamin Ummenhofer (Intel Labs, Germany), Sanskar Agrawal (Preimage, India), Rene Sepulveda (Intel Labs, USA), Yixing Lao (The University of Hong Kong, China), Kai Zhang (Adobe, USA), Tianhang Cheng (University of Illinois Urbana-Champaign, USA), Stephan Richter (Intel Labs, Germany), Shenlong Wang (University of Illinois Urbana-Champaign, USA), and German Ros (NVIDIA, USA)

Oral 5

DeDoDe: Detect, Don't Describe — Describe, Don't Detect for Local Feature Matching 148
Johan Edstedt (Linköping University), Georg Bökman (Chalmers University of Technology), Mårten Wadenbäck (Linköping University), and Michael Felsberg (Linköping University)

Enhancing Generalizability of Representation Learning for Data-Efficient 3D Scene Understanding 158
Yunsong Wang (National University of Singapore, Singapore), Na Zhao (Singapore University of Technology and Desig), and Gim Hee Lee (National University of Singapore, Singapore)

OPDMulti: Openable Part Detection for Multiple Objects 169
Xiaohao Sun (Simon Fraser University, Canada), Hanxiao Jiang (Simon Fraser University, Canada), Manolis Savva (Simon Fraser University, Canada), and Angel Chang (Simon Fraser University, Canada; Alberta Machine Intelligence Institute (Amii), Canada)

Scalable 3D Panoptic Segmentation As Superpoint Graph Clustering 179
Damien ROBERT (CSAI, ENGIE Lab CRIGEN, France; Univ Gustave Eiffel, France), Hugo Raguet (Univ. de Tours, France), and Loic Landrieu (Univ. Gustave Eiffel, France; Univ Gustave Eiffel, France)

Oral 6

A Local Appearance Model for Volumetric Capture of Diverse Hairstyles 190
Ziyan Wang (Meta Reality Labs Research; Reality Labs Research), Giljoo Nam (Reality Labs Research), Aljaz Bozic (Reality Labs Research), Chen Cao (Reality Labs Research), Jason Saragih (Reality Labs Research), Michael Zollhöfer (Reality Labs Research), and Jessica Hodgins (Reality Labs Research)

ContactArt: Learning 3D Interaction Priors for Category-Level Articulated Object and Hand Poses Estimation 201
Zehao Zhu (University of Texas at Austin), Jiashun Wang (Carnegie Mellon University), Yuzhe Qin (UC San Diego), Deqing Sun (Google Research), Varun Jampani (Google Research), and Xiaolong Wang (UC San Diego)

SPHEAR: Spherical Head Registration for Complete Statistical 3D Modeling 213
Eduard Gabriel Bazavan (Google Research), Andrei Zanfir (Google Research), Teodor Alexandru Szente (Google Research), Mihai Zanfir (Google Research), Thiemo Alldieck (Google Research), and Cristian Sminchisescu (Google Research)

Spotlights

Spotlight 1

3DRef: 3D Dataset and Benchmark for Reflection Detection in RGB and Lidar Data 225
Xiting Zhao (ShanghaiTech University, China) and Sören Schwertfeger (ShanghaiTech University, China)

ArtiGrasp: Physically Plausible Synthesis of Bi-Manual Dexterous Grasping and Articulation..... 235
Hui Zhang (ETH Zurich, Switzerland; Max Planck Institute for Intelligent Systems, Germany), Sammy Christen (ETH Zurich, Switzerland), Zicong Fan (ETH Zurich, Switzerland; Max Planck Institute for Intelligent Systems, Germany), Luocheng Zheng (ETH Zurich, Switzerland), Jemin Hwangbo (Korea Advanced Institute of Science and Technology, South Korea), Jie Song (ETH Zurich, Switzerland), and Otmar Hilliges (ETH Zurich, Switzerland)

DehazeNeRF: Multi-Image Haze Removal and 3D Shape Reconstruction using Neural Radiance Fields 247
Wei-Ting Chen (National Taiwan University; Stanford University), Yifan Wang (Stanford University), Sy-Yen Kuo (National Taiwan University), and Gordon Wetzstein (Stanford University)

Diffusion-HPC: Synthetic Data Generation for Human Mesh Recovery in Challenging Domains 257
Zhenzhen Weng (Stanford University), Laura Bravo-Sánchez (Stanford University), and Serena Yeung-Levy (Stanford University)

Efficient 3D Articulated Human Generation with Layered Surface Volumes 268
Yinghao Xu (Stanford University), Wang Yifan (Stanford Unviersity), Alexander W. Bergman (Stanford University), Menglei Chai (Google), Bolei Zhou (UCLA), and Gordon Wetzstein (Stanford University)

FastHuman: Reconstructing High-Quality Clothed Human in Minutes 280
Lixiang Lin (ZheJiang University), Songyou Peng (ETH Zurich), Qijun Gan (ZheJiang University), and Jianke Zhu (ZheJiang University)

Spotlight 2

3D Pose Estimation of Two Interacting Hands from a Monocular Event Camera 291
Christen Millerdurai (Max Planck Institute for Informatics; Saarland University), Diogo Luvizon (Max Planck Institute for Informatics), Viktor Rudnev (Max Planck Institute for Informatics; Saarland University), André Jonas (RPTU Kaiserslautern-Landau), Jiayi Wang (Max Planck Institute for Informatics), Christian Theobalt (Max Planck Institute for Informatics), and Vladislav Golyanik (Max Planck Institute for Informatics)

LumiGAN: Unconditional Generation of Relightable 3D Human Faces 302
Boyang Deng (Stanford University), Yifan Wang (Stanford University), and Gordon Wetzstein (Stanford University)

PathFusion: Path-Consistent Lidar-Camera Deep Feature Fusion 313
Lemeng Wu (Meta Reality Labs, United States), Dilin Wang (Meta Reality Labs, United States), Meng Li (Peking University, China), Yunyang Xiong (Meta Reality Lab, United States), Raghuraman Krishnamoorthi (Meta Reality Labs, United States), Qiang Liu (University of Texas at Austin, United States), and Vikas Chandra (Meta Reality Labs, United States)

PRAGO: Differentiable Multi-View Pose Optimization From Objectness Detections 324
Matteo Taiana (Pattern Analysis and Computer Vision (PAVIS), Istituto Italiano di Tecnologia (IIT), Italy), Matteo Toso (Pattern Analysis and Computer Vision (PAVIS), Istituto Italiano di Tecnologia (IIT), Italy), Stuart James (Pattern Analysis and Computer Vision (PAVIS), Istituto Italiano di Tecnologia (IIT), Italy), and Alessio Del Bue (Pattern Analysis and Computer Vision (PAVIS), Istituto Italiano di Tecnologia (IIT), Italy)

Spotlight 3

LabelMaker: Automatic Semantic Label Generation from RGB-D Trajectories 334
Silvan Weder (ETH Zurich), Hermann Blum (ETH Zurich), Francis Engelmann (ETH Zurich; ETH AI Center; Google), and Marc Pollefeys (ETH Zurich, Microsoft)

MC-Stereo: Multi-Peak Lookup and Cascade Search Range for Stereo Matching 344
Miaojie Feng (Huazhong University of Science and Technology), Junda Cheng (Huazhong University of Science and Technology), Hao Jia (Huazhong University of Science and Technology), Longliang Liu (Huazhong University of Science and Technology), Gangwei Xu (Huazhong University of Science and Technology), and Xin Yang (Huazhong University of Science and Technology)

MELON: NeRF with Unposed Images in SO(3) 354
Axel Levy (Stanford University), Mark Matthews (Google), Matan Sela (Google), Gordon Wetzstein (Stanford University), and Dmitry Lagun (Google)

Ray-Patch: An Efficient Querying for Light Field Transformers 365
Tomás Berriel Martins (Univesity of Zaragoza, Spain) and Javier Civera (University of Zaragoza, Spain)

Spotlight 4

Handbook on Leveraging Lines for Two-View Relative Pose Estimation 376
Petr Hruby (Department of Computer Science, ETH Zurich, Switzerland), Shaohui Liu (Department of Computer Science, ETH Zurich, Switzerland), Remi Pautrat (Department of Computer Science, ETH Zurich, Switzerland), Marc Pollefeys (Department of Computer Science, ETH Zurich, Switzerland; Microsoft Mixed Reality and AI Zurich lab), and Daniel Barath (ETH Zurich, Switzerland)

NCRF: Neural Contact Radiance Fields for Free-Viewpoint Rendering of Hand-Object Interaction 387
Zhongqun Zhang (University of Birmingham, UK; Huawei, Noah's Ark Lab), Jifei Song (Huawei, Noah's Ark Lab), Eduardo Pérez-Pellitero (Huawei, Noah's Ark Lab), Yiren Zhou (Huawei, Noah's Ark Lab), HyungJin Chang (University of Birmingham, UK), and Aleš Leonardis (University of Birmingham, UK; Huawei, Noah's Ark Lab)

PACE: Human and Camera Motion Estimation from in-the-Wild Videos 397
Muhammed Kocabas (NVIDIA; Max Planck Institute for Intelligent Systems, Germany; ETH Zurich, Switzerland), Ye Yuan (NVIDIA), Pavlo Molchanov (NVIDIA), Yunrong Guo (NVIDIA), Michael J. Black (Max Planck Institute for Intelligent Systems, Germany), Otmar Hilliges (ETH Zurich, Switzerland), Jan Kautz (NVIDIA), and Umar Iqbal (NVIDIA)

S4C: Self-Supervised Semantic Scene Completion With Neural Fields 409
Adrian Hayler (Technical University of Munich, Germany; MCML), Felix Wimbauer (Technical University of Munich, Germany; MCML), Dominik Muhle (Technical University of Munich, Germany; MCML), Christian Rupprecht (University of Oxford, United Kingdom), and Daniel Cremers (Technical University of Munich, Germany; MCML)

SALUDA: Surface-Based Automotive LiDAR Unsupervised Domain Adaptation 421
Björn Michele (Valeo.ai, France; Univ. Bretagne Sud, France), Alexandre Boulch (Valeo.ai, France), Gilles Puy (Valeo.ai, France), Tuan-Hung Vu (Valeo.ai, France), Renaud Marlet (Valeo.ai, France; Univ Gustave Eiffel, France), and Nicolas Courty (Univ. Bretagne Sud, France)

Sparse 3D Reconstruction via Object-Centric Ray Sampling 432
Llukman Cerkezi (University of Bern, Switzerland) and Paolo Favaro (University of Bern, Switzerland)

Spotlight 5

Out of the Room: Generalizing Event-Based Dynamic Motion Segmentation for Complex Scenes ... 442
Stamatios Georgoulis (Huawei Technologies, Zurich Research Center), Weining Ren (ETH Zurich), Alfredo Bochicchio (Huawei Technologies, Zurich Research Center), Daniel Eckert (Huawei Technologies, Zurich Research Center), Yuanyou Li (Huawei Technologies, Zurich Research Center), and Abel Gawel (Huawei Technologies, Zurich Research Center)

Partial-View Object View Synthesis via Filtering Inversion 453
Fan-Yun Sun (Stanford University; Nvidia), Jonathan Tremblay (Nvidia), Valts Blukis (Nvidia), Kevin Lin (Stanford University), Danfei Xu (Georgie Institute of Technology; Nvidia), Boris Ivanovic (Nvidia), Peter Karkus (Nvidia), Stan Birchfield (Nvidia), Dieter Fox (Nvidia), Ruohan Zhang (Stanford University), Yunzhu Li (University of Illinois Urbana-Champaign), Jiajun Wu (Stanford University), Marco Pavone (Nvidia; Stanford University), and Nick Haber (Stanford University)

Physically Plausible Full-Body Hand-Object Interaction Synthesis 464
Jona Braun (ETH Zurich, Switzerland), Sammy Christen (ETH Zurich, Switzerland), Muhammed Kocabas (ETH Zurich, Switzerland; Max Planck Institute for Intelligent Systems, Germany), Emre Aksan (ETH Zurich, Switzerland), and Otmar Hilliges (ETH Zurich, Switzerland)

Select-Sliced Wasserstein Distance for Point Cloud Learning 474
Bang Du (University of California San Diego, USA), Kunyao Chen (University of California San Diego, USA), Haochen Zhang (University of California San Diego, USA), and Truong Nguyen (University of California San Diego, USA)

Self-Supervised Learning of Neural Implicit Feature Fields for Camera Pose Refinement 484
Maxime Pietrantoni (Czech Technical University in Prague), Gabriela Csurka (Naver Labs Europe), Martin Humenberger (Naver Labs Europe), and Torsten Sattler (CIIRC CVUT)

Zero-BEV: Zero-Shot Projection of Any First-Person Modality to BEV Maps 495
Gianluca Monaci (NAVER LABS Europe, France), Leonid Antsfeld (NAVER LABS Europe, France), Boris Chidlovskii (NAVER LABS Europe, France), and Christian Wolf (NAVER LABS Europe, France)

3DV 2024

3D-TexSeg: Unsupervised Segmentation of 3D Texture Using Mutual Transformer Learning 506
Iyyakutti Iyappan Ganapathi (Khalifa University, UAE), Fayaz Ali (Khalifa University, UAE), Sajid Javed (Khalifa University, UAE), Syed Sadaf Ali (Khalifa University, UAE), and Naoufel Werghi (\Khalifa University, UAE)

A Benchmark Grocery Dataset of Realworld Point Clouds From Single View 516
Shivanand Venkanna Sheshappanavar (University of Delaware, USA), Tejas Anvekar (University of Delaware, USA), Shivanand Kundargi (University of Delaware, USA), Yufan Wang (University of Delaware, USA), and Chandra Kambhamettu (University of Delaware, USA)

A Cross Branch Fusion-Based Contrastive Learning Framework for Point Cloud Self-Supervised Learning 528
Chengzhi Wu (Karlsruhe Institute of Technology, Germany), Qianliang Huang (Karlsruhe Institute of Technology, Germany), Kun Jin (Tsinghua University, China), Julius Pfrommer (Fraunhofer Institute of Optronics, System Technologies and Image Exploitation IOSB, Germany), and Jürgen Beyerer (Karlsruhe Institute of Technology, Germany; Fraunhofer Institute of Optronics, System Technologies and Image Exploitation IOSB, Germany)

ActiveNeuS: Neural Signed Distance Fields for Active Stereo 539
Kazuto Ichimaru (Kyushu University, Japan), Takaki Ikeda (Kyushu University, Japan), Diego Thomas (Kyushu University, Japan), Takafumi Iwaguchi (Kyushu University, Japan), and Hiroshi Kawasaki (Kyushu University, Japan)

Addressing Low-Shot MVS by Detecting and Completing Planar Surfaces 549
Rajbir Kataria (University of Illinois Urbana-Champaign), Zhizhong Li (Amazon), Joseph DeGol (Microsoft), and Derek Hoiem (University of Illinois Urbana-Champaign)

BEVContrast: Self-Supervision in BEV Space for Automotive Lidar Point Clouds 559
Corentin Sautier (ENPC, France; valeo.ai, France), Gilles Puy (valeo.ai, France), Alexandre Boulch (valeo.ai, France), Renaud Marlet (ENPC, France; valeo.ai, France), and Vincent Lepetit (ENPC, France)

BLiSS: Bootstrapped Linear Shape Space 569
Sanjeev Muralikrishnan (University College London), Chun-Hao Paul Huang (Adobe Research), Duygu Ceylan (Adobe Research), and Niloy J. Mitra (University College London; Adobe Research)

Classical Photometric Stereo in Point Lighting Environments: Error Analysis and Mitigation..... 581
Simon Brenner (TU Wien, Austria) and Robert Sablatnig (TU Wien, Austria)

CloSe: A 3D Clothing Segmentation Dataset and Model 591
Dimitrije Antić (University of Amsterdam, Netherlands), Garvita Tiwari (University of Tübingen, Germany; Tübingen AI Center, Germany; Max Planck Institute for Informatics, Germany), Batuhan Ozcomlekci (University of Tübingen, Germany), Riccardo Marin (University of Tübingen, Germany; Tübingen AI Center, Germany), and Gerard Pons-Moll (University of Tübingen, Germany; Tübingen AI Center, Germany; Max Planck Institute for Informatics, Germany)

Cloth2Tex: A Customized Cloth Texture Generation Pipeline for 3D Virtual Try-On 602
Daiheng Gao (Alibaba Group, China), Xu Chen (ETH Zurich, Switzerland; Max Planck Institute for Intelligent Systems), Xindi Zhang (Alibaba Group, China), Qi Wang (Alibaba Group, China), Ke Sun (Alibaba Group, China), Bang Zhang (Alibaba Group, China), Liefeng Bo (Alibaba Group, China), and Qixing Huang (UT Austin, USA)

CoARF: Controllable 3D Artistic Style Transfer for Radiance Fields 612
Deheng Zhang (ETH Zurich, Switzerland), Clara Fernandez-Labrador (DisneyResearch, Studios, Switzerland), and Christopher Schroers (DisneyResearch, Studios, Switzerland)

Coherent Enhancement of Depth Images and Normal Maps Using Second-Order Geometric Models on Weighted Finite Graphs 623
Andreas Görlitz (University of Siegen, Germany), Michael Möller (University of Siegen, Germany), and Andreas Kolb (University of Siegen, Germany)

Color-NeuS: Reconstructing Neural Implicit Surfaces with Color 631
Licheng Zhong (Shanghai Jiao Tong University), Lixin Yang (Shanghai Jiao Tong University; Shanghai Qi Zhi Institute), Kailin Li (Shanghai Jiao Tong University), Haoyu Zhen (Shanghai Jiao Tong University), Mei Han (National University of Singapore), and Cewu Lu (Shanghai Jiao Tong University; Shanghai Qi Zhi Institute)

CombiNeRF: A Combination of Regularization Techniques for Few-Shot Neural Radiance Field View Synthesis 641
Matteo Bonotto (University of Padova, Italy; FlexSight, Italy), Luigi Sarrocco (University of Padova, Italy; FlexSight, Italy), Daniele Evangelista (University of Padova, Italy; FlexSight, Italy), Marco Imperoli (FlexSight, Italy), and Alberto Pretto (University of Padova, Italy)

Compositional 3D Scene Generation using Locally Conditioned Diffusion 651
Ryan Po (Stanford University) and Gordon Wetzstein (Stanford University)

Consistent-1-to-3: Consistent Image to 3D View Synthesis via Geometry-Aware Diffusion Models 664
Jianglong Ye (UC San Diego), Peng Wang (ByteDance), Kejie Li (ByteDance), Yichun Shi (ByteDance), and Heng Wang (ByteDance)

Continuous Cost Aggregation for Dual-Pixel Disparity Extraction 675
Sagi Monin (Technion - Israel Institute of Technology, Israel; Snap Inc., Austria), Sagi Katz (Snap Inc., Austria), and Georgios Evangelidis (Snap Inc., Austria)

Control3Diff: Learning Controllable 3D Diffusion Models from Single-View Images 685
Jiatao Gu (Apple), Qingzhe Gao (Shandong University), Shuangfei Zhai (Apple), Baoquan Chen (Peking University), Lingjie Liu (University of Pennsylvania), and Josh Susskind (Apple)

Controllable Dynamic Appearance for Neural 3D Portraits 697
ShahRukh Athar (Stony Brook University), Zhixin Shu (Adobe Research), Zexiang Xu (Adobe Research), Fujun Luan (Adobe Research), Sai Bi (Adobe Research), Kalyan Sunkavalli (Adobe Research), and Dimitris Samaras (Stony Brook University)

Correspondence-Free Online Human Motion Retargeting 707
Rim Rekik (Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, France), Mathieu Marsot (Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, France), Anne-Hélène Olivier (Univ. Rennes, Inria, CNRS, IRISA, M2S, France), Jean-Sébastien Franco (Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, France), and Stefanie Wuhrer (Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, France)

Cross3DVG: Cross-Dataset 3D Visual Grounding on Different RGB-D Scans 717
Taiki Miyanishi (Advanced Telecommunications Research Institute International; RIKEN AIP), Daichi Azuma (Kyoto University), Shuhei Kurita (RIKEN AIP), and Motoaki Kawanabe (Advanced Telecommunications Research Institute International)

DAC: Detector-Agnostic Spatial Covariances for Deep Local Features 728
Javier Tirado-Garín (University of Zaragoza), Frederik Warburg (Technical University of Denmark), and Javier Civera (University of Zaragoza)

Deep Event Visual Odometry 739
Simon Klenk (Technical University of Munich, Germany; Munich Center for Machine Learning), Marvin Motzet (Technical University of Munich, Germany; Munich Center for Machine Learning), Lukas Koestler (Technical University of Munich, Germany; Munich Center for Machine Learning), and Daniel Cremers (Technical University of Munich, Germany; Munich Center for Machine Learning)

DeepDR: Deep Structure-Aware RGB-D Inpainting for Diminished Reality 750
Christina Gsaxner (Graz University of Technology), Shohei Mori (Graz University of Technology), Dieter Schmalstieg (Graz University of Technology; University of Stuttgart), Jan Egger (Graz University of Technology; University of Duisburg-Essen), Gerhard Paar (Joanneum Research), Werner Bailer (Joanneum Research), and Denis Kalkofen (Graz University of Technology; Flinders University of South Australia)

DeepShaRM: Multi-View Shape and Reflectance Map Recovery Under Unknown Lighting 761
Kohei Yamashita (Kyoto University, Japan), Shohei Nobuhara (Kyoto University, Japan), and Ko Nishino (Kyoto University, Japan)

Depth Reconstruction with Neural Signed Distance Fields in Structured Light Systems 770
Rukun Qiao (Peking University, China), Hiroshi Kawasaki (Kyushu University, Japan), and Hongbin Zha (Peking University, China)

Developability Approximation for Neural Implicits Through Rank Minimization 780
Pratheba Selvaraju (University of Massachusetts, USA)

Diffusion Shape Prior for Wrinkle-Accurate Cloth Registration 790
Jingfan Guo (University of Minnesota), Fabian Prada (Codec Avatars Lab, Meta), Donglai Xiang (Carnegie Mellon University), Javier Romero (Codec Avatars Lab, Meta), Chenglei Wu (Codec Avatars Lab, Meta), Hyun Soo Park (University of Minnesota), Takaaki Shiratori (Codec Avatars Lab, Meta), and Shunsuke Saito (Codec Avatars Lab, Meta)

Dynamic 3D Gaussians: Tracking by Persistent Dynamic View Synthesis 800
Jonathon Luiten (Carnegie Mellon University, USA; RWTH Aachen University, Germany), Georgios Kopanas (Inria & Université Côte d'Azur, France), Bastian Leibe (RWTH Aachen University, Germany), and Deva Ramanan (Carnegie Mellon University, USA)

Dynamic Prototype Adaptation with Distillation for Few-Shot Point Cloud Segmentation 810
Jie Liu (University of Amsterdam, Netherlands), Wenzhe Yin (University of Amsterdam, Netherlands), Haochen Wang (University of Amsterdam, Netherlands), Yunlu Chen (Carnegie Mellon University, USA), Jan-Jakob Sonke (The Netherlands Cancer Institute, Netherlands, Netherlands), and Efstratios Gavves (University of Amsterdam, Netherlands)

DynamicSurf: Dynamic Neural RGB-D Surface Reconstruction With an Optimizable Feature Grid . 820
Mirgahney Mohamed (University College London) and Lourdes Agapito (University College London)

Event-Based Visual Odometry on Non-Holonomic Ground Vehicles 831
Wanting Xu (Mobile Perception Lab, ShanghaiTech University, China), Si'ao Zhang (Mobile Perception Lab, ShanghaiTech University, China), Li Cui (Mobile Perception Lab, ShanghaiTech University, China), Xin Peng (Mobile Perception Lab, ShanghaiTech University, China), and Laurent Kneip (Mobile Perception Lab, ShanghaiTech University, China)

Exploit Spatiotemporal Contextual Information for 3D Single Object Tracking via Memory Networks 842
Jongwon Ra (Zhejiang University), MengMeng Wang (Zhejiang University), Jianbiao Mei (Zhejiang University), Shanqi Liu (Zhejiang University), Yu Yang (Zhejiang University), and Yong Liu (Zhejiang University)

Farm3D: Learning Articulated 3D Animals by Distilling 2D Diffusion 852
Tomas Jakab (University of Oxford, United Kingdom), Ruining Li (University of Oxford, United Kingdom), Shangzhe Wu (University of Oxford, United Kingdom), Christian Rupprecht (University of Oxford, United Kingdom), and Andrea Vedaldi (University of Oxford, United Kingdom)

Fast High Dynamic Range Radiance Fields for Dynamic Scenes 862
Guanjun Wu (Huazhong University of Science and Technology), Taoran Yi (Huazhong University of Science and Technology), Jiemin Fang (Huazhong University of Science and Technology), Wenyu Liu (Huazhong University of Science and Technology), and Xinggang Wang (Huazhong University of Science and Technology)

Fast Relative Pose Estimation using Relative Depth 873
Jonathan Astermark (Lund University, Sweden), Yaqing Ding (Lund University, Sweden; Czech Technical University in Prague), Viktor Larsson (Lund University, Sweden), and Anders Heyden (Lund University, Sweden)

GAN-Avatar: Controllable Personalized GAN-Based Human Head Avatar 882
Berna Kabadayi (Max Planck Institute for Intelligent Systems, Germany), Wojciech Zielonka (Max Planck Institute for Intelligent Systems, Germany), Bharat Lal Bhatnagar (University of Tübingen; Max Planck Institute for Informatics, Germany), Gerard Pons-Moll (University of Tübingen; Tübingen AI Center; Max Planck Institute for Informatics, Germany), and Justus Thies (Max Planck Institute for Intelligent Systems, Germany; Technical University of Darmstadt)

Generalizing Single-View 3D Shape Retrieval to Occlusions and Unseen Objects 893
Qirui Wu (Simon Fraser University), Daniel Ritchie (Brown University), Manolis Savva (Simon Fraser University), and Angel X. Chang (Simon Fraser University)

Generating Continual Human Motion in Diverse 3D Scenes 903
Aymen Mir (University of Tübingen, Germany; Max Planck Institute for Informatics, Saarland Informatics Campus, Germany), Xavier Puig (FAIR at Meta), Angjoo Kanazawa (University of California, Berkeley), and Gerard Pons-Moll (University of Tübingen, Germany; Max Planck Institute for Informatics, Saarland Informatics Campus, Germany)

Geometrically Consistent Partial Shape Matching 914
Viktoria Ehm (Technical University of Munich, Germany; Munich Center for Machine Learning, Germany), Paul Roetzer (University of Bonn, Germany), Marvin Eisenberger (Technical University of Munich, Germany; Munich Center for Machine Learning, Germany), Maolin Gao (Technical University of Munich, Germany; Munich Center for Machine Learning, Germany), Florian Bernard (University of Bonn, Germany), and Daniel Cremers (Technical University of Munich, Germany; Munich Center for Machine Learning, Germany)

GHuNeRF: Generalizable Human NeRF from a Monocular Video 923
Chen Li (National University of Singapore, Singapore), Jiahao Lin (National University of Singapore, Singapore), and Gim Hee Lee (National University of Singapore, Singapore)

GRIP: Generating Interaction Poses Using Spatial Cues and Latent Consistency 933
Omid Taheri (Max Planck Institute for Intelligent Systems, Germany), Yi Zhou (Adobe Research), Dimitrios Tzionas (University of Amsterdam), Yang Zhou (Adobe Research), Duygu Ceylan (Adobe Research), Soren Pirk (Kiel University), and Michael J. Black (Max Planck Institute for Intelligent Systems, Germany)

HOC-Search: Efficient CAD Model and Pose Retrieval From RGB-D Scans 944
Stefan Ainetter (Graz University of Technology, Austria), Sinisa Stekovic (Graz University of Technology, Austria), Friedrich Fraundorfer (Graz University of Technology, Austria), and Vincent Lepetit (Ecole des Ponts, France)

HumanReg: Self-Supervised Non-Rigid Registration of Human Point Cloud 954
Yifan Chen (Tsinghua University, China), Zhiyu Pan (Tsinghua University, China), Zhicheng Zhong (Tsinghua University, China), Wenxuan Guo (Tsinghua University, China), Jianjiang Feng (Tsinghua University, China), and Jie Zhou (Tsinghua University, China)

Hyper-SNBRDF: Hypernetwork for Neural BRDF Using Sinusoidal Activation 965
Zhiqiang Li (Beihang University, China; Beihang University Yunnan Innovation Institute, China), Xukun Shen (Beihang University, China), Xueyang Zhou (Beihang University, China), Yong Hu (Beihang University, China; Beihang University Yunnan Innovation Institute, China), and Bowen Li (Beihang University, China)

Improved Scene Landmark Detection for Camera Localization 975
Tien Do (Tesla, USA) and Sudipta N. Sinha (Microsoft, USA)

Incorporating Rotation Invariance with Non-Invariant Networks for Point Clouds 985
Jiajun Fei (Tsinghua University, China) and Zhidong Deng (Tsinghua University, China)

InstantAvatar: Efficient 3D Head Reconstruction via Surface Rendering 995
Antonio Canela (Crisalix Labs, Spain; Universitat Politècnica de Catalunya; Institut de Robòtica i Informàtica Industrial (CSIC-UPC), Spain), Pol Caselles (Crisalix Labs, Spain; Universitat Politècnica de Catalunya; Institut de Robòtica i Informàtica Industrial (CSIC-UPC), Spain), Ibrar Malik (Crisalix Labs, Spain; Universitat Politècnica de Catalunya; Institut de Robòtica i Informàtica Industrial (CSIC-UPC), Spain), Eduard Ramon (Crisalix Labs, Spain; Universitat Politècnica de Catalunya), Jaime García (Crisalix Labs, Spain), Jordi Sànchez-Riera (Institut de Robòtica i Informàtica Industrial (CSIC-UPC), Spain), Gil Triginer (Crisalix Labs, Spain), and Francesc Moreno-Noguer (Institut de Robòtica i Informàtica Industrial (CSIC-UPC), Spain)

Interaction Replica: Tracking Human–Object Interaction and Scene Changes From Human Motion..... 1006
Vladimir Guzov (University of Tübingen, Germany; Max Planck Institute for Informatics, Germany), Julian Chibane (University of Tübingen, Germany; Max Planck Institute for Informatics, Germany), Riccardo Marin (University of Tübingen, Germany), Yannan He (University of Tübingen, Germany), Yunus Saracoglu (University of Tübingen, Germany), Torsten Sattler (Czech Technical University in Prague, Czech Republic), and Gerard Pons-Moll (University of Tübingen, Germany; Max Planck Institute for Informatics, Germany)

IS-NEAR: Implicit Semantic Neural Engine and Multi-Sensor Data Rendering With 3D Global Feature 1017
Tiecheng Sun (Huawei, China), Wei Zhang (University Of Stuttgart, Germany; Audiovisual Lab, Huawei Munich Research Center), Xingliang Dong (Huawei, China), and Tao Lin (Huawei, China)

Joint Spatial-Temporal Calibration for Camera and Global Pose Sensor 1027
Junlin Song (University of Luxembourg, Luxembourg), Antoine Richard (University of Luxembourg, Luxembourg), and Miguel Olivares-Mendez (University of Luxembourg, Luxembourg)

Lang3DSG: Language-Based Contrastive Pre-Training for 3D Scene Graph Prediction 1037
Sebastian Koch (Bosch Center for Artificial Intelligence; Robert Bosch Corporate Research; University of Ulm), Pedro Hermosilla (TU Wien), Narunas Vaskevicius (Bosch Center for Artificial Intelligence; Robert Bosch Corporate Research), Mirco Colosi (Robert Bosch Corporate Research), and Timo Ropinski (University of Ulm)

Learning Based Infinite Terrain Generation with Level of Detailing 1048
Aryamaan Jain (IIIT Hyderabad, India; Inria, Université Côte d'Azur, France), Avinash Sharma (IIIT Hyderabad, India; IIT Jodhpur, India), and K S Rajan (IIIT Hyderabad, India)

Learning to Estimate 6DoF Pose from Limited Data: A Few-Shot, Generalizable Approach using RGB Images 1059
Panwang Pan (bytedance), Zhiwen Fan (The University of Texas as Austin), Brandon Y. Feng (MIT), Peihao Wang (The University of Texas as Austin), Chenxin Li (The Chinese University of Hong Kong), and Zhangyang Wang (The University of Texas as Austin)

LocPoseNet: Robust Location Prior for Unseen Object Pose Estimation 1072
Chen Zhao (EPFL-CVLab), Yinlin Hu (Magic Leap), and Mathieu Salzmann (ClearSpace SA; EPFL-CVLab)

MACS: Mass Conditioned 3D Hand and Object Motion Synthesis 1082
Soshi Shimada (Max-Planck-Institut für Informatik, VIA Research Center), Franziska Mueller (Google Inc.), Jan Bednarik (Google Inc.), Bardia Doosti (Google Inc.), Bernd Bickel (Google Inc.), Danhang Tang (Google Inc.), Vladislav Golyanik (Max-Planck-Institut für Informatik), Jonathan Taylor (Google Inc.), Christian Theobalt (Max-Planck-Institut für Informatik, VIA Research Center), and Thabo Beeler (Google Inc.)

Mirror-Aware Neural Humans 1092
Daniel Ajisafe (The University of British Columbia), James Tang (The University of British Columbia), Shih-Yang Su (The University of British Columbia), Bastian Wandt (The University of British Columbia; Linkoping University), and Helge Rhodin (The University of British Columbia)

Mixing-Denoising Generalizable Occupancy Networks 1103
Amine Ouasfi (Inria, Univ. Rennes, France) and Adnane Boukhayma (CNRS, IRISA, M2S, France)

MixRT: Mixed Neural Representations For Real-Time NeRF Rendering 1115
Chaojian Li (Georgia Tech), Bichen Wu (Gen AI, Meta), Peter Vajda (Gen AI, Meta), and Yingyan Celine Lin (Georgia Tech)

MonoLSS: Learnable Sample Selection For Monocular 3D Detection 1125
Zhenjia Li (Baidu Inc., China), Jinrang Jia (Baidu Inc., China), and Yifeng Shi (Baidu Inc., China)

MuVieCAST: Multi-View Consistent Artistic Style Transfer 1136
Nail Ibrahimli (Delft University of Technology, The Netherlands), Julian F.P. Kooij (Delft University of Technology, The Netherlands), and Liangliang Nan (Delft University of Technology, The Netherlands)

NeRF-Feat: 6D Object Pose Estimation using Feature Rendering 1146
Shishir Reddy Vutukur (Technical University of Munich; Siemens AG), Heike Brock (Siemens AG), Benjamin Busam (Technical University of Munich), Tolga Birdal (Imperial College London), Andreas Hutter (Siemens AG), and Slobodan Ilic (Siemens AG)

NeRFMeshing: Distilling Neural Radiance Fields into Geometrically-Accurate 3D Meshes 1156
Marie-Julie Rakotosaona (Google, Switzerland), Fabian Manhardt (Google, Switzerland), Diego Martin Arroyo (Google, Switzerland), Michael Niemeyer (Google, Switzerland), Abhijit Kundu (Google, United States), and Federico Tombari (Google, Switzerland; TU Munich, Germany)

Neural Field Regularization by Denoising for 3D Sparse-View X-Ray Computed Tomography 1166
Romain Vo (Université Paris-Saclay, France; PSL University, France), Julie Escoda (Université Paris-Saclay, France), Caroline Vienne (Université Paris-Saclay, France), and Étienne Decencière (PSL University, France)

NeVRF: Neural Video-Based Radiance Fields for Long-Duration Sequences 1177
Minye Wu (KU Leuven) and Tinne Tuytelaars (KU Leuven)

OCBEV: Object-Centric BEV Transformer for Multi-View 3D Object Detection 1188
Zhangyang Qi (The University of Hong Kong), Jiaqi Wang (Shanghai Artificial Intelligence Laboratory), Xiaoyang Wu (The University of Hong Kong), and Hengshuang Zhao (The University of Hong Kong)

Occlusion Resilient 3D Human Pose Estimation 1198
Soumava Kumar Roy (Computer Vision Lab, EPFL, Switzerland), Ilia Badanin (Machine Learning and Optimization Lab, EPFL, Switzerland), Sina Honari (Samsung AI Center Toronto), and Pascal Fua (Computer Vision Lab, EPFL, Switzerland)

Oriented-Grid Encoder for 3D Implicit Representations 1208
Arihant Gaur (Mitsubishi Electric Research Labs (MERL)), G. Dias Pais (Mitsubishi Electric Research Labs (MERL); Instituto Superior Técnico, Lisboa), and Pedro Miraldo (Mitsubishi Electric Research Labs (MERL))

PanoSSC: Exploring Monocular Panoptic 3D Scene Reconstruction for Autonomous Driving 1219
Yining Shi (Tsinghua University, China; DiDi Chuxing, China), Jiusi Li (Tsinghua University, China), Kun Jiang (Tsinghua University, China), Ke Wang (Kargobot, China), Yunlong Wang (Tsinghua University, China), Mengmeng Yang (Tsinghua University, China), and Diange Yang (Tsinghua University, China)

PhoMoH: Implicit Photorealistic 3D Models of Human Heads 1229
Mihai Zanfir (Google Research), Thiemo Alldieck (Google Research), and Cristian Sminchisescu (Google Research)

Photometric Visibility Matrix for the Automatic Selection of Optimal Viewpoints 1240
Vanessa Staderini (AIT Austrian Institute of Technology GmbH, Austria), Tobias Glück (AIT Austrian Institute of Technology GmbH, Austria), Roberto Mecca (AIT Austrian Institute of Technology GmbH, Austria), Petra Gospodnetic (Fraunhofer ITWM, Germany), Philipp Schneider (AIT Austrian Institute of Technology GmbH, Austria), and Andreas Kugi (AIT Austrian Institute of Technology GmbH, Austria)

Physics-Based Indirect Illumination for Inverse Rendering 1249
Youming Deng (Cornell University, USA), Xueting Li (NVIDIA, USA), Sifei Liu (NVIDIA, USA), and Ming-Hsuan Yang (University of California, Merced, USA)

Physics-Based Rigid Body Object Tracking and Friction Filtering From RGB-D Videos 1259
Rama Krishna Kandukuri (Max Planck Institute for Intelligent Systems, Germany), Michael Strecke (Max Planck Institute for Intelligent Systems, Germany), and Joerg Stueckler (Max Planck Institute for Intelligent Systems, Germany)

PIVOT-Net: Heterogeneous Point-Voxel-Tree-Based Framework for Point Cloud Compression ... 1270
Jiahao Pang (InterDigital, USA), Kevin Bui (InterDigital, USA), and Dong Tian (InterDigital, USA)

Pix4Point: Image Pretrained Standard Transformers for 3D Point Cloud Understanding 1280
Guocheng Qian (King Abdullah University of Science and Technology (KAUST)), Abdullah Hamdi (King Abdullah University of Science and Technology (KAUST)), Xingdi Zhang (King Abdullah University of Science and Technology (KAUST)), and Bernard Ghanem (King Abdullah University of Science and Technology (KAUST))

PlaNeRF: SVD Unsupervised 3D Plane Regularization for NeRF Large-Scale Urban Scene Reconstruction 1291
Fusang Wang (Huawei Noah's Ark Lab), Arnaud Louys (Huawei Noah's Ark Lab), Nathan Piasco (Huawei Noah's Ark Lab), Moussab Bennehar (Huawei Noah's Ark Lab), Luis Roldao (Huawei Noah's Ark Lab), and Dzmitry Tsishkou (Huawei Noah's Ark Lab)

Practical Measurement and Neural Encoding of Hyperspectral Skin Reflectance 1301
Xiaohui Li (Imperial College London; Lumirithmic Ltd.), Giuseppe Claudio Guarnera (Lumirithmic Ltd.; University of York), Arvin Lin (Imperial College London; Lumirithmic Ltd.), and Abhijeet Ghosh (Imperial College London; Lumirithmic Ltd.)

Purposer: Putting Human Motion Generation in Context 1310
Nicolas Ugrinovic (Institut de Robobotica i Informatica Indutrial, Spain), Thomas Lucas (NAVER LABS Europe, France), Fabien Baradel (NAVER LABS Europe, France), Philippe Weinzaepfel (NAVER LABS Europe, France), Grégory Rogez (NAVER LABS Europe, France), and Francesc Moreno-Noguer (Institut de Robobotica i Informatica Indutrial, Spain)

PU-SDF: Arbitrary-Scale Uniformly Upsampling Point Clouds via Signed Distance Functions 1320
Shaohui Pan (South China University of Technology, China), Yong Xu (South China University of Technology, China; Peng Cheng Laboratory, China; Communication and Computer Network Laboratory of Guangdong, China), and Ruotao Xu (South China University of Technology, China)

Q-REG: End-to-End Trainable Point Cloud Registration with Surface Curvature 1330
Shengze Jin (ETH Zurich, Switzerland), Daniel Barath (ETH Zurich, Switzerland), Marc Pollefeys (ETH Zurich, Switzerland; Microsoft), and Iro Armeni (ETH Zurich, Switzerland; Stanford University)

Quantum-Hybrid Stereo Matching With Nonlinear Regularization and Spatial Pyramids 1340
Cameron Braunstein (Saarland University, Germany), Eddy Ilg (Saarland University, Germany), and Vladislav Golyanik (MPI for Informatics, Germany)

Range-Agnostic Multi-View Depth Estimation with Keyframe Selection 1350
Andrea Conti (Università di Bologna, Italy), Matteo Poggi (Università di Bologna, Italy), Valerio Cambareri (Sony Depthsensing Solutions, Italy), and Stefano Mattoccia (Università di Bologna, Italy)

Revisiting Depth Completion from a Stereo Matching Perspective for Cross-Domain Generalization 1360
Luca Bartolomei (University of Bologna, Italy), Matteo Poggi (University of Bologna, Italy), Andrea Conti (University of Bologna, Italy), Fabio Tosi (University of Bologna, Italy), and Stefano Mattoccia (University of Bologna, Italy)

Revisiting Map Relations for Unsupervised Non-Rigid Shape Matching 1371
Dongliang Cao (University of Bonn), Paul Roetzer (University of Bonn), and Florian Bernard (University of Bonn)

RIVQ-VAE: Discrete Rotation-Invariant 3D Representation Learning 1382
Mariem Mezghanni (LIX, Ecole Polytechnique, France), Malika Boulkenafed (LIX, Ecole Polytechnique, France), and Maks Ovsjanikov (LIX, Ecole Polytechnique, France)

ROAM: Robust and Object-Aware Motion Generation Using Neural Pose Descriptors 1392
Wanyue Zhang (Max Planck Institute for Informatics, Germany; Saarbrücken Research Center for Visual Computing, Interaction and AI), Rishabh Dabral (Max Planck Institute for Informatics, Germany; Saarbrücken Research Center for Visual Computing, Interaction and AI), Thomas Leimkühler (Max Planck Institute for Informatics, Germany), Vladislav Golyanik (Max Planck Institute for Informatics, Germany), Marc Habermann (Max Planck Institute for Informatics, Germany; Saarbrücken Research Center for Visual Computing, Interaction and AI), and Christian Theobalt (Max Planck Institute for Informatics, Germany; Saarbrücken Research Center for Visual Computing, Interaction and AI)

Robust Point Cloud Processing Through Positional Embedding 1403
Jianqiao Zheng (University of Adelaide), Xueqian Li (University of Adelaide), Sameera Ramasinghe (Amazon), and Simon Lucey (University of Adelaide)

RoomDesigner: Encoding Anchor-Latents for Style-Consistent and Shape-Compatible Indoor Scene Generation 1413
Yiqun Zhao (ShanghaiTech University), Zibo Zhao (ShanghaiTech University), Jing Li (Xiaohongshu Inc), Sixun Dong (ShanghaiTech University), and Shenghua Gao (ShanghaiTech University)

SceNeRFlow: Time-Consistent Reconstruction of General Dynamic Scenes 1424
Edith Tretschk (Max Planck Institute for Informatics, Germany), Vladislav Golyanik (Max Planck Institute for Informatics, Germany), Michael Zollhöfer (Meta Reality Labs Research, USA), Aljaž Božič (Meta Reality Labs Research, Switzerland), Christoph Lassner (Meta Reality Labs Research, USA), and Christian Theobalt (Max Planck Institute for Informatics, Germany)

Self-Supervised Learning of Skeleton-Aware Morphological Representation for 3D Neuron Segments 1436
Daiyi Zhu (National Engineering Laboratory for Brain-inspired Intelligence Technology and Application, University of Science and Technology of China, China. Institute of Artificial Intelligence, Hefei Comprehensive National Science Center, China), Qihua Chen (National Engineering Laboratory for Brain-inspired Intelligence Technology and Application, University of Science and Technology of China, China. Institute of Artificial Intelligence, Hefei Comprehensive National Science Center, China), and Xuejin Chen (National Engineering Laboratory for Brain-inspired Intelligence Technology and Application, University of Science and Technology of China, China. Institute of Artificial Intelligence, Hefei Comprehensive National Science Center, China)

SimpleEgo: Predicting Probabilistic Body Pose from Egocentric Cameras 1446
Hanz Cuevas Velasquez (Max Planck Institute for Intelligent Systems, Germany), Charlie Hewitt (Microsoft, UK), Sadegh Aliakbarian (Microsoft, UK), and Tadas Baltrušaitis (Microsoft, UK)

Single-View 3D Scene Reconstruction with High-Fidelity Shape and Texture 1456
Yixin Chen (National Key Laboratory of General Artificial Intelligence, BIGAI), Junfeng Ni (Tsinghua University), Nan Jiang (Peking University), Yaowei Zhang (National Key Laboratory of General Artificial Intelligence, BIGAI), Yixin Zhu (Peking University), and Siyuan Huang (National Key Laboratory of General Artificial Intelligence, BIGAI)

Split, Merge, and Refine: Fitting Tight Bounding Boxes via Over-Segmentation and Iterative Search 1468
Chanhyeok Park (KAIST, South Korea) and Minhyuk Sung (KAIST, South Korea)

Stable Surface Regularization for Fast Few-Shot NeRF 1478
Byeongin Joung (KAIST), Byeong-Uk Lee (KRAFTON), Jaesung Choe (KAIST), Ukcheol Shin (CMU), Minjun Kang (KAIST), Taeyeop Lee (KAIST), In So Kweon (KAIST), and Kuk-Jin Yoon (KAIST)

SUCRe: Leveraging Scene Structure for Underwater Color Restoration 1488
Clémentin Boittiaux (Ifremer, France; Université de Toulon, France), Ricard Marxer (Université de Toulon, Aix Marseille Univ, France), Claire Dune (Université de Toulon, France), Aurélien Arnaubec (Ifremer, France), Maxime Ferrera (Ifremer, France), and Vincent Hugel (Université de Toulon, France)

Synthesizing Physically Plausible Human Motions in 3D Scenes 1498
Liang Pan (Southeast University), Jingbo Wang (Shanghai AI Laboratory), Buzhen Huang (Southeast University), Junyu Zhang (Southeast University), Haofan Wang (Xiaohongshu Inc.), Xu Tang (Xiaohongshu Inc.), and Yangang Wang (Southeast University)

TADA! Text to Animatable Digital Avatars 1508
Tingting Liao (Mohamed bin Zayed University of Artificial Intelligence), Hongwei Yi (Max Planck Institute for Intelligent Systems), Yuliang Xiu (Max Planck Institute for Intelligent Systems), Jiaxiang Tang (Peking University), Yangyi Huang (State Key Lab of CAD & CG, Zhejiang University), Justus Thies (Max Planck Institute for Intelligent Systems), and Michael J. Black (Max Planck Institute for Intelligent Systems)

TECA: Text-Guided Generation and Editing of Compositional 3D Avatars 1520
Hao Zhang (Max Planck Institute for Intelligent Systems; Tsinghua University; RWTH Aachen University), Yao Feng (Max Planck Institute for Intelligent System; ETH Zurich), Peter Kulits (Max Planck Institute for Intelligent Systems), Yandong Wen (Max Planck Institute for Intelligent Systems), Justus Thies (Max Planck Institute for Intelligent Systems), and Michael J. Black (Max Planck Institute for Intelligent Systems)

TeCH: Text-Guided Reconstruction of Lifelike Clothed Humans 1531
Yangyi Huang (State Key Lab of CAD&CG, Zhejiang University), Hongwei Yi (Max Planck Institute for Intelligent Systems), Yuliang Xiu (Max Planck Institute for Intelligent Systems), Tingting Liao (Mohamed bin Zayed University of Artificial Intelligence), Jiaxiang Tang (Peking University), Deng Cai (State Key Lab of CAD&CG, Zhejiang University), and Justus Thies (Max Planck Institute for Intelligent Systems)

Test-Time Augmentation for 3D Point Cloud Classification and Segmentation 1543
Tuan-Anh Vu (The Hong Kong University of Science and Technology), Srinjay Sarkar (VinAI Research), Zhiyuan Zhang (Singapore Management University), Binh-Son Hua (Trinity College Dublin), and Sai-Kit Yeung (The Hong Kong University of Science and Technology)

TextMesh: Generation of Realistic 3D Meshes From Text Prompts 1554
Christina Tsalicoglou (ETH Zurich; Google), Fabian Manhardt (Google), Alessio Tonioni (Google), Michael Niemeyer (Google), and Federico Tombari (Google; Technical University of Munich)

Towards Learning Monocular 3D Object Localization From 2D Labels Using the Physical Laws of Motion 1564
Daniel Kienzle (University of Augsburg, Germany), Katja Ludwig (University of Augsburg, Germany), Julian Lorenz (University of Augsburg, Germany), and Rainer Lienhart (University of Augsburg, Germany)

UAVD4L: A Large-Scale Dataset for UAV 6-DoF Localization 1574
Rouwan Wu (National University of Defense Technology, China), Xiaoya Cheng (National University of Defense Technology, China), Juelin Zhu (National University of Defense Technology, China), Yuxiang Liu (National University of Defense Technology, China), Maojun Zhang (National University of Defense Technology, China), and Shen Yan (National University of Defense Technology, China)

Unsupervised 3D Keypoint Discovery with Multi-View Geometry 1584
Sina Honari (Samsung AI Center, Canada), Chen Zhao (Computer Vision Lab, EPFL, Switzerland), Mathieu Salzmann (Computer Vision Lab, EPFL, Switzerland), and Pascal Fua (Computer Vision Lab, EPFL, Switzerland)

Unsupervised Representation Learning for Diverse Deformable Shape Collections 1594
Sara Hahner (Fraunhofer SCAI and University of Bonn, Germany), Souhaib Attaiki (École Polytechnique, France), Jochen Garcke (Fraunhofer SCAI and University of Bonn, Germany), and Maks Ovsjanikov (École Polytechnique, France)

Visual Tomography: Physically Faithful Volumetric Models of Partially Translucent Objects 1605
David Nakath (University of Kiel, Germany), Xiangyu Weng (GEOMAR – Helmholtz Centre for Ocean Research Kiel, Germany), Mengkun She (University of Kiel, Germany), and Kevin Köser (University of Kiel, Germany)

YOLO-6D-Pose: Enhancing YOLO for Single-Stage Monocular Multi-Object 6D Pose Estimation .. 1616
Debapriya Maji (Texas Instruments, India), Soyeb Nagori (Texas Instruments, India), Manu Mathew (Texas Instruments, India), and Deepak Poddar (Texas Instruments, India)

Author Index