

Proceedings

SIGGRAPH 2024 Conference Papers

Denver, Colorado

July 28 – August 1, 2024

Conference Chair

Andres Burbano

SIGGRAPH 2024 Technical Papers Chair

Denis Zorin

SIGGRAPH 2024 Technical Papers Assistant Chair

Wojciech Jarosz

Proceedings Production Editor

Stephen N. Spencer, University of Washington

Sponsored by ACM SIGGRAPH



Association for
Computing Machinery

Advancing Computing as a Science & Profession

The Association for Computing Machinery
1601 Broadway, 10th Floor
New York, New York 10019, USA

ACM COPYRIGHT NOTICE. Copyright © 2024 by the Association for Computing Machinery, Inc. Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Publications Dept., ACM, Inc., fax +1 (212) 869-0481, or permissions@acm.org.

For other copying of articles that carry a code at the bottom of the first or last page, copying is permitted provided that the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923,

+1-978-750-8400, +1-978-750-4470 (fax).

ACM ISBN: 979-8-4007-0525-0

Table of Contents

Preface	xv
---------------	----

Vector Graphics

QT-Font: High-efficiency Font Synthesis via Quadtree-based Diffusion Models	Article 1
<i>Yitian Liu, Zhouhui Lian</i>	

Minkowski Penalties: Robust Differentiable Constraint Enforcement for Vector Graphics	Article 2
<i>Jří Minarčík, Sam Estep, Wode Ni, Keenan Crane</i>	

Ciallo: GPU-Accelerated Rendering of Vector Brush Strokes.....	Article 3
<i>Shen Ciao, Zhongyue Guan, Qianxi Liu, Li-Yi Wei, Zeyu Wang</i>	

Material Texture Generation and Painting

MaPa: Text-driven Photorealistic Material Painting for 3D Shapes	Article 4
<i>Shangzhan Zhang, Sida Peng, Tao Xu, Yuanbo Yang, Tianrun Chen, Nan Xue, Yujun Shen, Hujun Bao, Ruizhen Hu, Xiaowei Zhou</i>	

Diffusion Texture Painting	Article 5
<i>Anita Hu, Nishkrit Desai, Hassan Abu Alhaija, Seung Wook Kim, Maria Shugrina</i>	

TexPainter: Generative Mesh Texturing with Multi-view Consistency	Article 6
<i>Hongkun Zhang, Zherong Pan, Congyi Zhang, Lifeng Zhu, Xifeng Gao</i>	

TexSliders: Diffusion-Based Texture Editing in CLIP Space	Article 7
<i>Julia Guerrero-Viu, Milos Hasan, Arthur Roullier, Midhun Harikumar, Yiwei Hu, Paul Guerrero, Diego Gutierrez, Belen Masia, Valentin Deschaintre</i>	

Monte Carlo for PDEs

Velocity-Based Monte Carlo Fluids	Article 8
<i>Ryusuke Sugimoto, Christopher Batty, Toshiya Hachisuka</i>	

Neural Monte Carlo Fluid Simulation	Article 9
<i>Pranav Jain, Ziyin Qu, Peter Yichen Chen, Oded Stein</i>	

Neural Control Variates with Automatic Integration.....	Article 10
<i>Zilu Li, Guandao Yang, Qingqing Zhao, Xi Deng, Leonidas Guibas, Bharath Hariharan, Gordon Wetzstein</i>	

A Differential Monte Carlo Solver for the Poisson Equation	Article 11
<i>Zihan Yu, Lifan Wu, Zhiqian Zhou, Shuang Zhao</i>	

* denotes equal contribution.

Table of Contents

Generative 3D Geometry

Coin3D: Controllable and Interactive 3D Assets Generation with Proxy-Guided Conditioning.....	Article 12
<i>Wenqi Dong*, Bangbang Yang*, Lin Ma, Xiao Liu, Liyuan Cui, Hujun Bao, Yuewen Ma, Zhaopeng Cui</i>	
DreamFont3D: Personalized Text-to-3D Artistic Font Generation.....	Article 13
<i>Xiang Li, Lei Meng, Lei Wu, Manyi Li, Xiangxu Meng</i>	
ThemeStation: Generating Theme-Aware 3D Assets from Few Exemplars	Article 14
<i>Zhenwei Wang, Tengfei Wang, Gerhard Hancke, Ziwei Liu, Rynson W.H. Lau</i>	
Spice-E: Structural Priors in 3D Diffusion using Cross-Entity Attention.....	Article 15
<i>Etai Sella*, Gal Fiebelman*, Noam Atia, Hadar Averbuch-Elor</i>	

3D Face Generator and Animation

HeadArtist: Text-conditioned 3D Head Generation with Self Score Distillation.....	Article 16
<i>Hongyu Liu, Xuang Wang, Ziyu Wan, Yujun Shen, Yibing Song, Jing Liao, Qifeng Chen</i>	
Toonify3D: StyleGAN-based 3D Stylized Face Generator	Article 17
<i>Wonjong Jang, Yucheol Jung, Hyomin Kim, Gwangjin Ju, Chaewon Son, Joeeun Son, Seungyong Lee</i>	
Media2Face: Co-speech Facial Animation Generation with Multi-Modality Guidance.....	Article 18
<i>Qingcheng Zhao*, Pengyu Long,* Qixuan Zhang, Dafei Qin, Han Liang, Longwen Zhang, Yingliang Zhang, Lan Xu, Jingyi Yu</i>	

Differentiable Rendering

Path-Space Differentiable Rendering of Implicit Surfaces.....	Article 19
<i>Siwei Zhou, Youngha Chang, Nobuhiko Mukai, Hiroaki Santo, Fumio Okura, Yasuyuki Matsushita, Shuang Zhao</i>	
Woven Fabric Capture with a Reflection-Transmission Photo Pair	Article 20
<i>Yingjie Tang, Zixuan Li, Miloš Hašan, Jian Yang, Beibei Wang</i>	

Geometry: Reconstruction

MVD ² : Efficient Multiview 3D Reconstruction for Multiview Diffusion	Article 21
<i>Xin-Yang Zheng, Hao Pan, Yu-Xiao Guo, Xin Tong, Yang Liu</i>	
High-quality Surface Reconstruction using Gaussian Surfels.....	Article 22
<i>Pinxuan Dai, Jiamin Xu, Wenxiang Xie, Xinguo Liu, Huamin Wang, Weiwei Xu</i>	
Reach for the Arcs: Reconstructing Surfaces from SDFs via Tangent Points	Article 23
<i>Silvia Sellán, Yingying Ren, Christopher Batty, Oded Stein</i>	
Part123: Part-aware 3D Reconstruction from a Single-view Image.....	Article 24
<i>Anran Liu*, Cheng Lin*, Yuan Liu, Xiaoxiao Long, Zhiyang Dou, Hao-Xiang Guo, Ping Luo, Wenping Wang</i>	

* denotes equal contribution.

Table of Contents

Consistent Text-to-Image

Subject-Diffusion: Open Domain Personalized Text-to-Image Generation without Test-time Fine-tuning..... Article 25
Jian Ma, Junhao Liang, Chen Chen, Haonan Lu

The Chosen One: Consistent Characters in Text-to-Image Diffusion Models Article 26
Omri Avrahami, Amir Hertz, Yael Vinker, Moab Arar, Shlomi Fruchter, Ohad Fried, Daniel Cohen-Or, Dani Lischinski

Streetscapes: Large-scale Consistent Street View Generation Using Autoregressive Video Diffusion Article 27
Boyang Deng, Richard Tucker, Zhengqi Li, Leonidas Guibas, Noah Snaveley, Gordon Wetzstein**

Blue Noise for Diffusion Models Article 28
Xingchang Huang, Corentin Salaün, Cristina Vasconcelos, Christian Theobalt, Cengiz Öztireli, Gurprit Singh

Geometry: Mappings and Fields

Neural Geometry Fields for Meshes Article 29
Venkataram Sivaram, Ravi Ramamoorthi, Tzu-Mao Li

Fast Radiance Fields

RTG-SLAM: Real-time 3D Reconstruction at Scale Using Gaussian Splatting Article 30
Zhexi Peng, Tianjia Shao, Yong Liu, Jingke Zhou, Yin Yang, Jingdong Wang, Kun Zhou

BoostMVSNeRFs: Boosting MVS-based NeRFs to Generalizable View Synthesis in Large-scale Scenes Article 31
Chih-Hai Su, Chih-Yao Hu*, Shr-Ruei Tsai*, Jie-Ying Lee*, Chin-Yang Lin, Yu-Lun Liu*

2D Gaussian Splatting for Geometrically Accurate Radiance Fields..... Article 32
Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, Shenghua Gao

VR, Eye Tracking, Perception

Saccade-Contingent Rendering..... Article 33
Yuna Kwak, Eric Penner, Xuan Wang, Mohammad R. Saeedpour-Parizi, Olivier Mercier, Xiuyun Wu, T. Scott Murdison, Phillip Guan

Perceptual Evaluation of Steered Retinal Projection Article 34
Seungjae Lee, Seung-Woo Nam*, Kevin Rio, Renate Landig, Hsien-Hui Cheng, Lu Lu, Barry Silverstein*

Simulating Nature

Real-time Wing Deformation Simulations for Flying Insects Article 35
Qiang Chen, Zhigang Deng, Feng Li, Yuming Fang, Tingsong Lu, Yang Tong, Yifan Zuo

Modelling a Feather as a Strongly Anisotropic Elastic Shell..... Article 36
Jean Jouve, Victor Romero, Rahul Narain, Laurence Boissieux, Theodore Kim, Florence Bertails-Descoubes

* denotes equal contribution.

Table of Contents

Clothing Geometry

LayGA: Layered Gaussian Avatars for Animatable Clothing Transfer..... Article 37
Siyu Lin, Zhe Li, Zhaoqi Su, Zerong Zheng, Hongwen Zhang, Yebin Liu

Singular Foliations for Knit Graph Design Article 38
Rahul Mitra, Erick Jimenez Berumen, Megan Hofmann, Edward Chien

NeRFs and Lighting

NeRF as a Non-Distant Environment Emitter in Physics-based Inverse Rendering Article 39
Jingwang Ling, Ruihan Yu, Feng Xu, Chun Du, Shuang Zhao

3D Gaussian Splatting with Deferred Reflection Article 40
Keyang Ye, Qiming Hou, Kun Zhou

Lite2Relight: 3D-aware Single Image Portrait Relighting Article 41
Pramod Rao, Gereon Fox, Abhimitra Meka, Mallikarjun B R, Fangneng Zhan, Tim Weyrich, Bernd Bickel, Hanspeter Pfister, Wojciech Matusik, Mohamed Elgharib, Christian Theobalt

EyeIR: Single Eye Image Inverse Rendering in the Wild Article 42
Shijun Liang, Haofei Wang, Feng Lu

Fluids

Lagrangian Covector Fluid with Free Surface Article 43
Zhiqi Li, Barnabás Börcsök, Duowen Chen, Yutong Sun, Bo Zhu, Greg Turk

Fluid Control with Laplacian Eigenfunctions Article 44
Yixin Chen, David I. W. Levin, Timothy R. Langlois

3D People and their Habitats

Physics-based Scene Layout Generation from Human Motion..... Article 45
Jianan Li, Tao Huang, Qingxu Zhu, Tien-Tsin Wong

VRMM: A Volumetric Relightable Morphable Head Model..... Article 46
Haotian Yang, Mingwu Zheng, Chongyang Ma, Yu-Kun Lai, Pengfei Wan, Haibin Huang

Compressed Skinning for Facial Blendshapes..... Article 47
Ladislav Kavan, John Doublestein, Martin Prazak, Matthew Cioffi, Doug Roble

3D Fabrication

Modal Folding: Discovering Smooth Folding Patterns for Sheet Materials using Strain-Space Modes..... Article 48
Pengbin Tang, Ronan Hinchet, Roi Poranne, Bernhard Thomaszewski, Stelian Coros

* denotes equal contribution.

Table of Contents

Motion Capture

Towards Unstructured Unlabeled Optical Mocap: *A Video Helps!*..... Article 49
Nicholas Milef, John Keyser, Shu Kong

Physical Non-inertial Poser (PNP): Modeling Non-inertial Effects in Sparse-inertial Human Motion Capture Article 50
Xinyu Yi, Yuxiao Zhou, Feng Xu

Ultra Inertial Poser: Scalable Motion Capture and Tracking from Sparse Inertial Sensors and
Ultra-Wideband Ranging..... Article 51
Rayan Armani, Changlin Qian, Jiayi Jiang, Christian Holz

Hand-Object Interaction Controller (HOIC): Deep Reinforcement Learning for Reconstructing
Interactions with Physics..... Article 52
Haoyu Hu, Xinyu Yi, Zhe Cao, Jun-Hai Yong, Feng Xu

Physics-Informed Learning of Characteristic Trajectories for Smoke Reconstruction Article 53
Yiming Wang, Siyu Tang, Mengyu Chu

Shape Analysis

Consistent Point Orientation for Manifold Surfaces via Boundary Integration..... Article 54
Weizhou Liu, Xingce Wang, Haichuan Zhao, Xingfei Xue, Zhongke Wu, Xuequan Lu, Ying He

A Linear Method to Consistently Orient Normals of a 3D Point Cloud..... Article 55
Craig Gotsman, Kai Horman

Navigation-Driven Approximate Convex Decomposition Article 56
James Andrews

Into the Portal: Directable Fractal Self-Similarity..... Article 57
Alexa Schor, Theodore Kim

3D Head Avatars From Data

MonoGaussianAvatar: Monocular Gaussian Point-based Head Avatar Article 58
Yufan Chen, Lizhen Wang, Qijing Li, Hongjiang Xiao, Shengping Zhang, Hongxun Yao, Yebin Liu

InvertAvatar: Incremental GAN Inversion for Generalized Head Avatars Article 59
Xiaochen Zhao, Jingxiang Sun*, Lizhen Wang, Jinli Suo, Yebin Liu*

3D Gaussian Blendshapes for Head Avatar Animation Article 60
Shengjie Ma, Yanlin Weng, Tianjia Shao, Kun Zhou

* denotes equal contribution.

Table of Contents

Real-Time Rendering: Hair, Fabrics, and Super-Resolution

Real-Time Hair Rendering with Hair Meshes.....	Article 61
<i>Gaurav Bhokare, Eisen Montalvo, Elie Diaz, Cem Yuksel</i>	
Modeling Hair Strands with Roving Capsules.....	Article 62
<i>Alexander Reshetov, David Hart</i>	
Real-time Neural Woven Fabric Rendering.....	Article 63
<i>Xiang Chen, Lu Wang, Beibei Wang</i>	
Mob-FGSR: Frame Generation and Super Resolution for Mobile Real-Time Rendering.....	Article 64
<i>Sipeng Yang, Qingchuan Zhu, Junhao Zhuge, Qiang Qiu, Chen Li, Yuzhong Yan, Huihui Xu, Ling-Qi Yan, Xiaogang Jin</i>	
Deep Fourier-based Arbitrary-scale Super-resolution for Real-time Rendering	Article 65
<i>Haonan Zhang, Jie Guo, Jiawei Zhang, Haoyu Qin, Zesen Feng, Ming Yang, Yanwen Guo</i>	

Character Animation from Data

LGTM: Local-to-Global Text-Driven Human Motion Diffusion Model.....	Article 66
<i>Haowen Sun, Ruikun Zheng, Haibin Huang, Chongyang Ma, Hui Huang, Ruizhen Hu</i>	
Taming Diffusion Probabilistic Models for Character Control.....	Article 67
<i>Rui Chen, Mingyi Shi, Shaoli Huang, Ping Tan, Taku Komura, Xuelin Chen</i>	
TEDi: Temporally-Entangled Diffusion for Long-Term Motion Synthesis.....	Article 68
<i>Zihan Zhang, Richard Liu, Kfir Aberman, Rana Hanocka</i>	
Flexible Motion In-betweening with Diffusion Models.....	Article 69
<i>Setareh Cohan, Guy Tevet, Daniele Reda, Xue Bin Peng, Michiel van de Panne</i>	
WalkTheDog: Cross-Morphology Motion Alignment via Phase Manifolds.....	Article 70
<i>Peizhuo Li, Sebastian Starke, Yuting Ye, Olga Sorkine-Hornung</i>	
Iterative Motion Editing with Natural Language	Article 71
<i>Purvi Goel, Kuan-Chieh Wang, C. Karen Liu, Kayvon Fatahalian</i>	

Rendering, Sampling and Tracing

Quad-Optimized Low-Discrepancy Sequences	Article 72
<i>Victor Ostromoukhov, Nicolas Bonneel, David Coeurjolly, Jean-Claude Iehl</i>	

* denotes equal contribution.

Table of Contents

Lighting and Matting with Image Generation

DiLightNet: Fine-grained Lighting Control for Diffusion-based Image Generation Article 73
Chong Zeng, Yue Dong, Pieter Peers, Youkang Kong, Hongzhi Wu, Xin Tong

IntrinsicDiffusion: Joint Intrinsic Layers from Latent Diffusion Models Article 74
Jundan Luo, Duygu Ceylan, Jae Shin Yoon, Nanxuan Zhao, Julien Philip, Anna Frühstück, Wenbin Li, Christian Richardt, Tuanfeng Y. Wang

RGB $\leftrightarrow$ X: Image Decomposition and Synthesis using Material- and Lighting-aware Diffusion Models Article 75
Zheng Zeng, Valentin Deschaintre, Iliyan Georgiev, Yannick Hold-Geoffroy, Yiwei Hu, Fujun Luan, Ling-Qi Yan, Miloš Hašan

Matting by Generation..... Article 76
Zhixiang Wang, Baiang Li, Jian Wang, Yu-Lun Liu, Jinwei Gu, Yung-Yu Chuang, Shin'ichi Satoh

Virtual Interaction and Real Devices

Dragon's Path: Synthesizing User-Centered Flying Creature Animation Paths for
Outdoor Augmented Reality Experiences Article 77
Minyoung Kim, Rawan Alghofaili, Changyang Li, Lap-Fai Yu

VR-GS: A Physical Dynamics-Aware Interactive Gaussian Splatting System in Virtual Reality..... Article 78
Ying Jiang, Chang Yu*, Tianyi Xie*, Xuan Li*, Yutao Feng, Huamin Wang, Minchen Li, Henry Lau, Feng Gao, Yin Yang, Chenfanfu Jiang*

Temporal Acoustic Point Holography Article 79
Giorgos Christopoulos, Lei Gao, Diego Martinez Plasencia, Marta Betcke, Ryuji Hirayama, Sriram Subramanian

Cloth Simulation

Neural-Assisted Homogenization of Yarn-Level Cloth..... Article 80
Xudong Feng, Huamin Wang, Yin Yang, Weiwei Xu

ContourCraft: Learning to Resolve Intersections in Neural Multi-Garment Simulations..... Article 81
Artur Grigorev, Giorgio Becherini, Michael J. Black, Otmar Hilliges, Bernhard Thomaszewski

3D Shape Analysis

DAE-Net: Deforming Auto-Encoder for Fine-grained Shape Co-segmentation Article 82
Zhiqin Chen, Qimin Chen, Hang Zhou, Hao Zhang

* denotes equal contribution.

Table of Contents

Dynamic Radiance Fields

ST-4DGS: Spatial-Temporally Consistent 4D Gaussian Splatting for Efficient Dynamic Scene Rendering.....	Article 83
<i>Deqi Li, Shi-Sheng Huang, Zhiyuan Lu, Xinran Duan, Hua Huang</i>	
GaussianPrediction: Dynamic 3D Gaussian Prediction for Motion Extrapolation and Free View Synthesis.....	Article 84
<i>Boming Zhao, * Yuan Li*, Ziyu Sun, Lin Zeng, Yujun Shen, Rui Ma, Yinda Zhang, Hujun Bao, Zhaopeng Cui</i>	
Factorized Motion Fields for Fast Sparse Input Dynamic View Synthesis	Article 85
<i>Nagabhushan Somraj, Kapil Choudhary, Sai Harsha Mupparaju, Rajiv Soundararajan</i>	
Modeling Ambient Scene Dynamics for Free-view Synthesis.....	Article 86
<i>Meng-Li Shih, Jia-Bin Huang, Changil Kim, Rajvi Shah, Johannes Kopf, Chen Gao</i>	
4D-Rotor Gaussian Splatting: Towards Efficient Novel View Synthesis for Dynamic Scenes	Article 87
<i>Yuanxing Duan*, Fangyin Wei*, Qiyu Dai, Yuhang He, Wenzheng Chen, Baoquan Chen</i>	
Controllable Neural Style Transfer for Dynamic Meshes	Article 88
<i>Guilherme G. Haetinge, Jingwei Tang, Raphael Ortiz, Paul Kanyuk, Vinicius C. Azevedo</i>	

Appearance Models

A Realistic Multi-scale Surface-based Cloth Appearance Model	Article 89
<i>Junqiu Zhu, Christophe Hery, Lukas Bode, Carlos Aliaga, Adrian Jarabo, Ling-Qi Yan, Matt Jen-Yuan Chiang</i>	

Simulating Deformation

Stabler Neo-Hookean Simulation: Absolute Eigenvalue Filtering for Projected Newton	Article 90
<i>Honglin Chen, Hsueh-Ti Derek Liu, David I.W. Levin, Changxi Zheng, Alec Jacobson</i>	
Efficient Position-Based Deformable Colon Modeling for Endoscopic Procedures Simulation.....	Article 91
<i>Marcelo Martins, Lucas Morais, Luciana Nedel, Rafael Torchelsen, Anderson Maciel</i>	

Generative 3D Geometry and Editing

GEM3D: GEnenerative Medial Abstractions for 3D Shape Synthesis	Article 92
<i>Dmitry Petrov, Pradyumn Goyal, Vikas Thamizharasan, Vladimir G. Kim, Matheus Gadelha, Melinos Averkiou, Siddhartha Chaudhuri, Evangelos Kalogerakis</i>	

Rendering, Denoising & Path Guiding

Practical Error Estimation for Denoised Monte Carlo Image Synthesis	Article 93
<i>Arthur Firmino, Ravi Ramamoorthi, Jeppe Revall Frisvad, Henrik Wann Jensen</i>	

Perception, Image, Video

Versatile Vision Foundation Model for Image and Video Colorization	Article 94
<i>Vukasin Bozic, Abdelaziz Djelouah, Yang Zhang, Radu Timofte, Markus Gross, Christopher Schroers</i>	

* denotes equal contribution.

Table of Contents

Simulation with Contact

Primal–Dual Non-Smooth Friction for Rigid Body Animation..... Article 95
Yi-Lu Chen, Mickaël Ly, Chris Wojtan

Preconditioned Nonlinear Conjugate Gradient Method for Real-time Interior-point Hyperelasticity Article 96
Xing Shen, Runyuan Cai, Mengxiao Bi, Tangjie Lv

A Dynamic Duo of Finite Elements and Material Points Article 97
Xuan Li, Minchen Li, Xuchen Han, Huamin Wang, Yin Yang, Chenfanfu Jiang

Spatial Data Structures

Neural Bounding Article 98
Stephanie Wenxin Liu, Michael Fischer, Paul D. Yoo, Tobias Ritschel

N-BVH: Neural Ray Queries with Bounding Volume Hierarchies Article 99
Philippe Weier, Alexander Rath, Élie Michel, Iliyan Georgiev, Philipp Slusallek, Tamy Boubekeur

ReFiNe: Recursive Field Networks for Cross-Modal Multi-Scene Representation Article 100
Sergey Zakharov, Katherine Liu, Adrien Gaidon, Rareş Ambruş

Controllable Image Generation and Completion

Filter-Guided Diffusion for Controllable Image Generation Article 101
Zeqi Gu, Ethan Yang*, Abe Davis*

LOOSECONTROL: Lifting ControlNet for Generalized Depth Conditioning Article 102
Shariq Farooq Bhat, Niloy J. Mitra, Peter Wonka

Separate-and-Enhance: Compositional Finetuning for Text2Image Diffusion Models Article 103
Zipeng Bao, Yijun Li, Krishna Kumar Singh, Yu-Xiong Wang, Martial Hebert

Object-level Scene Deocclusion Article 104
Zhengzhe Liu, Qing Liu, Chirui Chang, Jianming Zhang, Daniil Pakhomov, Haitian Zheng, Zhe Lin, Daniel Cohen-Or, Chi-Wing Fu

Character Animation: 2D, 3D, Robot

Text-Guided Synthesis of Crowd Animation..... Article 105
Xuebo Ji, Zherong Pan, Xifeng Gao, Jia Pan

Soft Pneumatic Actuator Design using Differentiable Simulation Article 106
Arvi Gjoka, Espen Knoop, Moritz Bäcker, Denis Zorin, Daniele Panozzo

SMEAR: Stylized Motion Exaggeration with ARt-direction Article 107
Jean Basset, Pierre Bénard, Pascal Barla

* denotes equal contribution.

Table of Contents

Geometry: Editing and Deformation

Semantic Shape Editing with Parametric Implicit Templates	Article 108
<i>Uday Kusupati, Mathieu Gaillard, Jean-Marc Thiery, Adrien Kaiser</i>	
A Unified Differentiable Boolean Operator with Fuzzy Logic	Article 109
<i>Hsueh-Ti Derek Liu, Maneesh Agrawala, Cem Yuksel, Tim Omernick, Vinith Misra, Stefano Corazza, Morgan McGuire, Victor Zordan</i>	
CNS-Edit: 3D Shape Editing via Coupled Neural Shape Optimization.....	Article 110
<i>Jingyu Hu, Ka-Hei Hui, Zhengzhe Liu, Hao Zhang, Chi-Wing Fu</i>	

Video Generation

Motion-I2V: Consistent and Controllable Image-to-Video Generation with Explicit Motion Modeling.....	Article 111
<i>Xiaoyu Shi*, Zhaoyang Huang*, Fu-Yun Wang*, Weikang Bian*, Dasong Li, Yi Zhang, Manyuan Zhang, Ka Chun Cheung, Simon See, Hongwei Qin, Jifeng Dai, Hongsheng Li</i>	
I2V-Adapter: A General Image-to-Video Adapter for Diffusion Models.....	Article 112
<i>Xun Guo, Mingwu Zheng, Liang Hou, Yuan Gao, Yufan Deng, Pengfei Wan, Di Zhang, Yufan Liu, Weiming Hu, Zhengjun Zha, Haibin Huang, Chongyang Ma</i>	
Direct-a-Video: Customized Video Generation with User-Directed Camera Movement and Object Motion.....	Article 113
<i>Shiyuan Yang, Liang Hou, Haibin Huang, Chongyang Ma, Pengfei Wan, Di Zhang, Xiaodong Chen, Jing Liao</i>	
MotionCtrl: A Unified and Flexible Motion Controller for Video Generation	Article 114
<i>Zhouxia Wang, Ziyang Yuan, Xintao Wang, Yaowei Li, Tianshui Chen, Menghan Xia, Ping Luo, Ying Shan</i>	
X-Portrait: Expressive Portrait Animation with Hierarchical Motion Attention	Article 115
<i>You Xie, Hongyi Xu, Guoxian Song, Chao Wang, Yichun Shi, Linjie Luo</i>	

Computational Cameras and Displays

Tele-Aloha: A Low-budget and High-authenticity Telepresence System Using Sparse RGB Cameras.....	Article 116
<i>Hanzhang Tu, Ruizhi Shao, Xue Dong, Shunyuan Zheng, Hao Zhang, Lili Chen, Meili Wang, Wenyu Li, Siyan Ma, Shengping Zhang, Boyao Zhou, Yebin Liu</i>	
Aperture-Aware Lens Design.....	Article 117
<i>Arjun Teh, Ioannis Gkioulekas, Matthew O'Toole</i>	
Scale-Invariant Monocular Depth Estimation via SSI Depth	Article 118
<i>S. Mahdi H. Miangoleh, Mahesh Reddy, Yağız Aksoy</i>	
Deep Hybrid Camera Deblurring for Smartphone Cameras	Article 119
<i>Jaesung Rim, Junyong Lee, Heemin Yang, Sunghyun Cho</i>	
Self-Supervised Video Defocus Deblurring with Atlas Learning	Article 120
<i>Lingyan Ruan, Martin Bálint, Mojtaba Bermana, Krzysztof Wolski, Hans-Peter Seidel, Karol Myszkowski, Bin Chen</i>	

* denotes equal contribution.

Table of Contents

Character Control

SuperPADL: Scaling Language-Directed Physics-Based Control with Progressive Supervised Distillation Article 121
Jordan Juravsky, Yunrong Guo, Sanja Fidler, Xue Bin Peng

Strategy and Skill Learning for Physics-based Table Tennis Animation..... Article 122
Jiashun Wang, Jessica Hodgins, Jungdam Won

Procedural Geometry

Recompose Grammars for Procedural Architecture Article 123
Niklaus Houska, Cheryl Lau, Matthias Specht

Radiance Field Processing

A Construct-Optimize Approach to Sparse View Synthesis without Camera Pose..... Article 124
Kaiwen Jiang, Yang Fu, Mukund Varma T, Yash Belhe, Xiaolong Wang, Hao Su, Ravi Ramamoorthi

Rip-NeRF: Anti-aliasing Radiance Fields with Ripmap-Encoded Platonic Solids..... Article 125
Junchen Liu, Wenbo Hu*, Zhuo Yang*, Jianteng Chen, Guoliang Wang, Xiaoxue Chen, Yantong Cai, Huan-ang Gao, Hao Zhao*

N-Dimensional Gaussians for Fitting of High Dimensional Functions Article 126
Stavros Diolatzis, Tobias Zirr, Alexander Kuznetsov, Georgios Kopanas, Anton Kaplanyan

Sound, Light, Radiofrequency

AONeuS: A Neural Rendering Framework for Acoustic-Optical Sensor Fusion..... Article 127
Mohamad Qadri, Kevin Zhang*, Akshay Hinduja, Michael Kaess, Adithya Pediredla, Christopher A. Metzler*

DIFFSOUND: Differentiable Modal Sound Rendering and Inverse Rendering for Diverse Inference Tasks Article 128
Xutong Jin, Chenxi Xu*, Ruohan Gao, Jiajun Wu, Guoping Wang, Sheng Li*

Accelerating Saccadic Response through Spatial and Temporal Cross-Modal Misalignments..... Article 129
Daniel Jiménez-Navarro, Xi Peng, Yunxiang Zhang, Karol Myszkowski, Hans-Peter Seidel, Qi Sun, Ana Serrano

Radar Fields: Frequency-Space Neural Scene Representations for FMCW Radar..... Article 130
David Borts, Erich Liang, Tim Brödermann, Andrea Ramazzina, Stefanie Walz, Edoardo Palladin, Jipeng Sun, David Bruggemann, Christos Sakaridis, Luc Van Gool, Mario Bijelic, Felix Heide

* denotes equal contribution.

Table of Contents

Art, Illusion, Fabrication

Diffusion Illusions: Hiding Images in Plain Sight.....	Article 131
<i>Ryan Burgert, Xiang Li, Abe Leite, Kanchana Ranasinghe, Michael S. Ryoo</i>	
Cross-Image Attention for Zero-Shot Appearance Transfer	Article 132
<i>Yuval Alaluf*, Daniel Garibi*, Or Patashnik, Hadar Averbuch-Elor, Daniel Cohen-Or</i>	
Generative Escher Meshes.....	Article 133
<i>Noam Aigerman, Thibault Groueix</i>	
Fabricable 3D Wire Art	Article 134
<i>Kenji Tojo, Ariel Shamir, Bernd Bickel, Nobuyuki Umetani</i>	
SIGGRAPH 2024 Conference Committee.....	xvii
SIGGRAPH 2024 Technical Papers Committee and Reviewers	xviii

* denotes equal contribution.

Preface

We are pleased to introduce the special issue of TOG and SIGGRAPH Conference Proceedings which include all papers accepted to the SIGGRAPH 2024 conference. This year, we received an extraordinary 844 submissions, a 36% increase over last year. The Technical Papers Committee (PC), composed of 121 experts from academia and industry, accepted 252 of the submissions: 118 as Journal papers and 134 as Conference papers, resulting in an overall acceptance rate of 30%. As in previous years, papers were submitted either as Journal-only (no page limit) or dual-track (page limit of 7 pages, with 2 additional pages for figures). There were 192 journal-only and 652 dual-track submissions. In addition to these papers, the Technical Papers program will include 27 papers accepted to other issues of TOG, totaling 279 presentations—double the number in 2019.

The themes of the papers presented at SIGGRAPH continue to evolve, showcasing exciting advances in well-established computer graphics research directions and exploration of novel applications, in particular, based on generative AI. The advances in AI have expanded the range of topics shared between computer vision and computer graphics, accounting for a significant portion of the growth in submissions.

This extraordinary growth, while promising for the future of SIGGRAPH, comes with challenges: assembling a 121-person committee is difficult and leaves fewer people available for the SIGGRAPH Asia committee. Consequently, the next SIGGRAPH and SIGGRAPH Asia chairs and the Papers Advisory Group (PAG) decided to make some modifications to the review process for the future. This year, the review process largely remained the same—well-tuned and thorough—with some small enhancements. However, this is likely the last year in the near future when we can afford not to reduce the number of reviews per paper.

This year, we introduced an additional pre-review meeting of the Technical Papers Committee, held around the time of the sorting committee meeting, to review and discuss the review criteria and policies before the reviews began. Another new feature is a more definitive policy on comparison requests in reviews and the definition of prior work.

As in previous years, the process remained double-blind, with five reviews per paper. We switched to a single score (from strong reject to strong accept) for dual-track papers, with a separate question about conference vs. dual paper acceptance. A refined description of the differences between conference and journal papers was included in the CFP and committee guidance. Nevertheless, many discussions focused on this distinction, sometimes distracting from the merits of the papers themselves.

Most experienced members of the committee agreed to serve as Moderators, who helped facilitate post-rebuttal reviewer discussions and provide feedback on review summaries. This year, the moderator role was particularly important: the committee meeting duration (3 days) is fixed, but decisions need to be made on an increasing number of papers, making it crucial to converge on decisions before the meeting. To improve the quality of feedback provided to authors, we worked to provide detailed instructions on summary content and structure, particularly for rejected papers.

Since 2020, the Technical Papers Committee meetings have been virtual, and this year was no exception. Although we briefly considered returning to in-person meetings, it was not feasible given the committee size. Virtual meetings have their own set of challenges. This year, the videoconferencing system used in 2021-2023 (Ohayay) was terminated by its owner, Snap, so we switched to a Zoom-based system developed by Kayvon Fatahalian.

Preface

Serving as the technical papers chairs gave us a new appreciation for the scale of the effort and the incredible work done by all volunteers and staff members. Our sincere gratitude goes to the many participants in the process:

- * Leona Caffey for patiently guiding us through the process and providing unparalleled support at every step, handling countless aspects of the program.
- * Leona's colleagues at SmithBucklin, Brenda Dreier, Mandie Aadland, and Emma Siemsen, for their help during the meeting.
- * Mark Montague for implementing new features, extracting needed information from Linklings on short notice, and coming up with creative solutions to various problems.
- * Adam Finkelstein for making numerous changes to HepCat to accommodate the new meeting format and working with Kayvon to integrate it with the new Zoom app.
- * Kayvon Fatahalian for building the Ohay replacement, figuring out the endless idiosyncrasies of the Zoom API, and running it during the committee meeting.
- * The paper's video team, Derek Bradley, Jim McCann, Enrique Rosales, and Nicholas Vining, for putting the trailer together, and the student volunteers (Ryan Capouellez, Arvi Gjoka, Juhyeon Kim, Abhishek Madan, Jonathan Panuelos, Ty Mason Trusty, Zhecheng Wang) for selecting the videos.
- * The 15 sorters who assigned all 844 submissions to 121 committee members in 2 days.
- * The 121 papers committee members who worked hard to find reviewers for all submissions, which is becoming increasingly difficult as the number of submissions grows, and for making decisions on all papers in a 3-day meeting.
- * The 29 Conflict of Interest Coordinators who had to do more work than expected.
- * The 1416 tertiary reviewers who, along with committee members, wrote 3531 reviews.
- * The Advisory Board and the Papers Advisory Group and its Chair, George Drettakis, for providing timely input on a range of topics.
- * The Test-of-Time Awards Committee and the Best Paper Awards Committee, and Sylvain Paris and Richard Zhang for chairing these committees.

In conclusion, we are very proud of the quality and breadth of the SIGGRAPH 2024 Technical Papers program. Despite the challenges posed by our growth, the dedication and hard work of the Technical Papers Committee, reviewers, and all involved have ensured that SIGGRAPH continues to be a premier venue for groundbreaking research. We are optimistic about the future of SIGGRAPH, confident that it will continue to thrive and adapt, setting new standards for excellence in the years to come.

Denis Zorin
Technical Papers Chair

Wojciech Jarosz
Technical Papers Assistant Chair