







# acm Transactions on Graphics


July 2024  
Volume 43 Number 4





# acm Transactions on Graphics

July 2024  
Volume 43 Number 4



The Association for Computing Machinery, Inc.  
1601 Broadway, 10th Floor  
New York, New York 10019-7434

Copyright © 2024 by the Association for Computing Machinery, Inc (ACM). Permission to make digital or hard copies of portions of this work for personal or classroom use is granted without fee provided that the 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. Copyright 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 permission to republish from Publications Department, ACM, Inc. Fax +1-212-869-0481 or e-mail [permissions@acm.org](mailto: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.

#### Notice to Past Authors of ACM-Published Articles

ACM intends to create a complete electronic archive of all articles and/or other material previously published by ACM. If you have written a work that was previously published by ACM in any journal or conference proceedings prior to 1978, or any SIG newsletter at any time, and you do NOT want this work to appear in the ACM Digital Library, please inform [permissions@acm.org](mailto:permissions@acm.org), stating the title of the work, the author(s), and where and when published.

ACM ISSN 0730-0301  
ACM Order Number 428030

Additional copies may be ordered from ACM.

ACM  
1601 Broadway, 10th Floor  
New York, NY 10019-7434  
+1-212-869-7440  
+1-212-869-0481 (fax)

Articles in this journal issue are paginated by article number and page number within the article, rather than by consecutive page numbers from the start of the first issue of the journal's current volume. The table of contents, author index and reference format all use this article-based pagination system.

ACM is transitioning to an article-based, "online first" content publishing system and all ACM journals are undergoing a similar transition.

# Table of Contents

Preface..... xiii

## Vector Graphics

Text-to-Vector Generation with Neural Path Representation..... Article 36  
*Peiyang Zhang, Nanxuan Zhao, Jing Liao*

Deep Sketch Vectorization via Implicit Surface Extraction..... Article 37  
*Chuan Yan, Yong Li, Deepali Aneja, Matthew Fisher, Edgar Simo-Serra, Yotam Gingold*

Robust Containment Queries over Collections of Rational Parametric Curves via Generalized Winding Numbers ..... Article 38  
*Jacob Spainhour, David Gunderman, Kenneth Weiss*

## Material Texture Generation and Painting

DreamMat: High-quality PBR Material Generation with Geometry- and Light-aware Diffusion Models ..... Article 39  
*Yuqing Zhang\*, Yuan Liu\*, Zhiyu Xie, Lei Yang, Zhongyuan Liu, Mengzhou Yang, Runze Zhang, Qilong Kou, Cheng Lin, Wenping Wang, Xiaogang Jin*

EASI-Text: Edge-Aware Mesh Texturing from Single Image..... Article 40  
*Sai Raj Kishore Perla, Yizhi Wang, Ali Mahdavi-Amiri, Hao Zhang*

## Monte Carlo for PDEs

Walkin' Robin: Walk on Stars with Robin Boundary Conditions..... Article 41  
*Bailey Miller\*, Rohan Sawhney\*, Keenan Crane\*\*, Ioannis Gkioulekas\*\**

Stochastic Computation of Barycentric Coordinates ..... Article 42  
*Fernando de Goes, Mathieu Desbrun*

## Generative 3D Geometry

BlockFusion: Expandable 3D Scene Generation using Latent Tri-plane Extrapolation ..... Article 43  
*Zhennan Wu, Yang Li, Han Yan, Taizhang Shang, Weixuan Sun, Senbo Wang, Ruikai Cui, Weizhe Liu, Hiroyuki Sato, Hongdong Li, Pan Ji*

SketchDream: Sketch-based Text-to-3D Generation and Editing..... Article 44  
*Feng-Lin Liu, Hongbo Fu, Yu-Kun Lai, Lin Gao*

## 3D Face Generator and Animation

Portrait3D: Text-Guided High-Quality 3D Portrait Generation Using Pyramid Representation and GANs Prior ..... Article 45  
*Yiqian Wu, Hao Xu, Xiangjun Tang, Xien Chen, Siyu Tang, Zhebin Zhang, Chen Li, Xiaogang Jin*

DiffPoseTalk: Speech-Driven Stylistic 3D Facial Animation and Head Pose Generation via Diffusion Models..... Article 46  
*Zhiyao Sun, Tian Lv, Sheng Ye, Matthieu Lin, Jenny Sheng, Yu-Hui Wen, Minjing Yu, Yong-Jin Liu*

S3: Speech, Script and Scene driven Head and Eye Animation ..... Article 47  
*Yifang Pan, Rishabh Agrawal, Karan Singh*

\* and \*\* denote equal contribution

# Table of Contents

## Differentiable Rendering

Conditional Mixture Path Guiding for Differentiable Rendering ..... Article 48  
*Zhimin Fan, Pengcheng Shi, Mufan Guo, Ruoyu Fu, Yanwen Guo, Jie Guo*

ZeroGrads: Learning Local Surrogates for Non-Differentiable Graphics ..... Article 49  
*Michael Fischer, Tobias Ritschel*

## Geometry: Reconstruction

NeuralTO: Neural Reconstruction and View Synthesis of Translucent Objects ..... Article 50  
*Yuxiang Cai, Jiaxiong Qiu, Zhong Li, Bo Ren*

NeurCADRecon: Neural Representation for Reconstructing CAD Surfaces by Enforcing Zero Gaussian Curvature.... Article 51  
*Qiuji Dong, Rui Xu, Pengfei Wang, Shuangmin Chen, Shiqing Xin, Xiaohong Jia, Wenping Wang, Changhe Tu*

## Consistent Text-to-Image

Training-Free Consistent Text-to-Image Generation..... Article 52  
*Yoad Tewel, Omri Kaduri, Rinon Gal, Yoni Kasten, Lior Wolf, Gal Chechik, Yuval Atzmon*

## Fluids and Flows

A Vortex Particle-on-Mesh Method for Soap Film Simulation ..... Article 53  
*Ningxiao Tao, Liangwang Ruan, Yitong Deng, Bo Zhu, Bin Wang, Baoquan Chen*

Multi-Material Mesh-Based Surface Tracking with Implicit Topology Changes ..... Article 54  
*Peter Heiss-Synak\*, Aleksei Kalinov\*, Malina Strugaru, Arian Etemadi, Huidong Yang, Chris Wojtan*

Kinetic Simulation of Turbulent Multifluid Flows..... Article 55  
*Wei Li, Kui Wu, Mathieu Desbrun*

An Induce-on-Boundary Magnetostatic Solver for Grid-Based Ferrofluids..... Article 56  
*Xingyu Ni\*, Ruicheng Wang\*, Bin Wang, Baoquan Chen*

Going with the Flow..... Article 57  
*Yousuf Soliman, Marcel Padilla, Oliver Gross, Felix Knöppel, Ulrich Pinkall, Peter Schröder*

Efficient Debris-flow Simulation for Steep Terrain Erosion ..... Article 58  
*Aryamaan Jain, Bedrich Benes, Guillaume Cordonnier*

## Geometry: Mappings and Fluids

Smooth Bijective Projection in a High-order Shell..... Article 59  
*Shibo Liu, Yang Ji, Jia-Peng Guo, Ligang Liu, Xiao-Ming Fu*

Lifting Directional Fields to Minimal Sections ..... Article 60  
*David Palmer, Albert Chern, Justin Solomon*

Seamless Parametrization in Penner Coordinates ..... Article 61  
*Ryan Capouellez, Denis Zorin*

\* and \*\* denote equal contribution

# Table of Contents

## Fast Radiance Fields

A Hierarchical 3D Gaussian Representation for Real-Time Rendering of Very Large Datasets..... Article 62  
*Bernhard Kerbl\*, Andreas Meuleman\*, Georgios Kopanas, Michael Wimmer, Alexandre Lanvin, George Drettakis*

SMERF: Streamable Memory Efficient Radiance Fields for Real-Time Large-Scene Exploration..... Article 63  
*Daniel Duckworth\*, Peter Hedman\*, Christian Reiser, Peter Zhizhin, Jean-François Thibert, Mario Lučić, Richard Szeliski, Jonathan T. Barron*

StopThePop: Sorted Gaussian Splatting for View-Consistent Real-time Rendering ..... Article 64  
*Lukas Radl\*, Michael Steiner\*, Mathias Parger, Alexander Weinrauch, Bernhard Kerbl, Markus Steinberger*

## VR, Eye Tracking, Perception

Towards Motion Metamers for Foveated Rendering..... Article 65  
*Taimoor Tariq, Piotr Didyk*

Cybersickness Reduction via Gaze-Contingent Image Deformation ..... Article 66  
*Colin Groth, Marcus Magnor, Steve Grogorick, Martin Eisemann, Piotr Didyk*

PEA-PODs: Perceptual Evaluation of Algorithms for Power Optimization in XR Displays ..... Article 67  
*Kenneth Chen, Thomas Wan, Nathan Matsuda, Ajit Ninan, Alexandre Chapiro, Qi Sun*

Holographic Parallax Improves 3D Perceptual Realism ..... Article 68  
*Dongyeon Kim\*, Seung-Woo Nam\*, Suyeon Choi, Jong-Mo Seo, Gordon Wetzstein, Yoonchan Jeong*

## Simulating Nature

Differentiable Voronoi Diagrams for Simulation of Cell-Based Mechanical Systems ..... Article 69  
*Logan Numerow, Yue Li, Stelian Coros, Bernhard Thomaszewski*

Scintilla: Simulating Combustible Vegetation for Wildfires..... Article 70  
*Andrzej Kokosza, Helge Wrede, Daniel Gonzalez Esparza, Milosz Makowski, Daoming Liu, Dominik L. Michels, Sören Pirk, Wojtek Pałubicki*

Cyclogenesis: Simulating Hurricanes and Tornadoes ..... Article 71  
*Jorge Alejandro Amador Herrera, Jonathan Klein, Daoming Liu, Wojtek Pałubicki, Sören Pirk, Dominik L. Michels*

## Clothing Geometry

DressCode: Autoregressively Sewing and Generating Garments from Text Guidance ..... Article 72  
*Kai He, Kaixin Yao, Qixuan Zhang, Lingjie Liu, Jingyi Yu, Lan Xu*

Proxy Asset Generation for Cloth Simulation in Games ..... Article 73  
*Zhongtian Zheng, Tongtong Wang, Qijia Feng, Zherong Pan, Xifeng Gao, Kui Wu*

Automatic Digital Garment Initialization from Sewing Patterns ..... Article 74  
*Chen Liu, Weiwei Xu, Yin Yang, Huamin Wang*

\* and \*\* denote equal contribution

# Table of Contents

## NeRFs and Lighting

LightFormer: Light-Oriented Global Neural Rendering in Dynamic Scene ..... Article 75  
*Haocheng Ren, Yuchi Huo, Yifan Peng, Hongtao Sheng, Weidong Xue, Hongxiang Huang, Jingzhen Lan, Rui Wang, Hujun Bao*

Eulerian-Lagrangian Fluid Simulation on Particle Flow Maps..... Article 76  
*Junwei Zhou, Duowen Chen, Molin Deng, Yitong Deng, Yuchen Sun, Sinan Wang, Shiyong Xiong, Bo Zhu*

Lightning-fast Method of Fundamental Solutions ..... Article 77  
*Jiong Chen, Florian Schäfer, Mathieu Desbrun*

## Geometric Modeling & Analysis

Spin-It Faster: Quadrics Solve All Topology Optimization Problems That Depend Only On Mass Moments ..... Article 78  
*Christian Hafner, Mickaël Ly, Chris Wojtan*

X-SLAM: Scalable Dense SLAM for Task-aware Optimization using CSFD..... Article 79  
*Zhexi Peng, Yin Yang, Tianjia Shao, Chenfanfu Jiang, Kun Zhou*

CWF: Consolidating Weak Features in High-quality Mesh Simplification..... Article 80  
*Rui Xu\*, Longdu Liu\*, Ningna Wang, Shuangmin Chen, Shiqing Xin, Xiaohu Guo, Zichun Zhong, Taku Komura, Wenping Wang, Changhe Tu*

Alignment conditions for NURBS-based design of mixed tension-compression grid shells ..... Article 81  
*Masaaki Miki, Toby Mitchell*

Adaptive grid generation for discretizing implicit complexes..... Article 82  
*Yiwen Ju, Xingyi Du, Qingnan Zhou, Nathan Carr, Tao Ju*

## 3D People and Their Habitats

Spatial and Surface Correspondence Field for Interaction Transfer ..... Article 83  
*Zeyu Huang, Honghao Xu, Haibin Huang, Chongyang Ma, Hui Huang, Ruizhen Hu*

CharacterGen: Efficient 3D Character Generation from Single Images with Multi-View Pose Canonicalization..... Article 84  
*Hao-Yang Peng, Jia-Peng Zhang, Meng-Hao Guo, Yan-Pei Cao, Shi-Min Hu*

\* and \*\* denote equal contribution



# Table of Contents

## 3D Fabrication

Neural Slicer for Multi-Axis 3D Printing..... Article 85  
*Tao Liu\**, *Tianyu Zhang\**, *Yongxue Chen*, *Yuming Huang*, *Charlie C. L. Wang*

FlexScale: Modeling and Characterization of Flexible Scaled Sheets..... Article 86  
*Juan Montes Maestre\**, *Yinwei Du\**, *Ronan Hinchet*, *Stelian Coros*, *Bernhard Thomaszewski*

Computational Homogenization for Inverse Design of Surface-based Inflatables ..... Article 87  
*Yingying Ren*, *Julian Panetta*, *Seiichi Suzuki*, *Uday Kusupati*, *Florin Isvoranu*, *Mark Pauly*

Solid Knitting..... Article 88  
*Yuichi Hirose*, *Mark Gillespie*, *Angelica M. Bonilla Fominaya*, *James McCann*

Fabric Tessellation: Realizing Freeform Surfaces by Smocking ..... Article 89  
*Aviv Segall*, *Jing Ren*, *Amir Vaxman*, *Olga Sorkine-Hornung*

## Motion Capture

Audio Matters Too! Enhancing Markerless Motion Capture with Audio Signals for String Performance Capture..... Article 90  
*Yitong Jin\**, *Zhiping Qiu\**, *Yi Shi*, *Shuangpeng Sun*, *Chongwu Wang*, *Donghao Pan*, *Jiachen Zhao*,  
*Zhenghao Liang*, *Yuan Wang*, *Xiaobing Li*, *Feng Yu*, *Tao Yu*, *Qionghai Dai*

## Shape Analysis

Differentiable Geodesic Distance for Intrinsic Minimization on Triangle Meshes..... Article 91  
*Yue Li*, *Logan Numerow*, *Bernhard Thomaszewski*, *Stelian Coros*

A Heat Method for Generalized Signed Distance..... Article 92  
*Nicole Feng*, *Keenan Crane*

## 3D Head Avatars From Data

Universal Facial Encoding of Codec Avatars from VR Headsets ..... Article 93  
*Shaojie Bai*, *Te-Li Wang*, *Chenghui Li*, *Akshay Venkatesh*, *Tomas Simon*, *Chen Cao*, *Gabriel Schwartz*,  
*Jason Saragih*, *Yaser Sheikh*, *Shih-En Wei*

Learning a Generalized Physical Face Model From Data ..... Article 94  
*Lingchen Yang*, *Gaspard Zoss*, *Prashanth Chandran*, *Markus Gross*, *Barbara Solenthaler*,  
*Eftychios Sifakis*, *Derek Bradley*

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

Real-Time Physically Guided Hair Interpolation ..... Article 95  
*Jerry Hsu*, *Tongtong Wang*, *Zherong Pan*, *Xifeng Gao*, *Cem Yuksel*, *Kui Wu*

\* and \*\* denote equal contribution

# Table of Contents

## Rendering, Sampling and Tracing

Stylized Rendering as a Function of Expectation..... Article 96  
*Rex West, Sayan Mukherjee*

Proxy Tracing: Unbiased Reciprocal Estimation for Optimized Sampling in BDPT..... Article 97  
*Fujia Su\*, Bingxuan Li\*, QingYang Yin, Yanchen Zhang, Sheng Li*

Area ReSTIR: Resampling for Real-Time Defocus and Antialiasing..... Article 98  
*Song Zhang\*, Daqi Lin\*, Markus Kettunen, Cem Yuksel, Chris Wyman*

Ray Tracing Harmonic Functions..... Article 99  
*Mark Gillespie, Denise Yang, Mario Botsch, Keenan Crane*

## Lighting and Matting With Image Generation

Transparent Image Layer Diffusion using Latent Transparency..... Article 100  
*Lymin Zhang, Maneesh Agrawala*

## Virtual Interaction and Real Devices

3D-Layers: Bringing Layer-Based Color Editing to VR Painting..... Article 101  
*Emilie Yu, Fanny Chevalier, Karan Singh, Adrien Bousseau*

NICER: A New and Improved Consumed Endurance and Recovery Metric  
to Quantify Muscle Fatigue of Mid-Air Interactions..... Article 102  
*Yi Li, Benjamin Tag, Shaozhang Dai, Robert Crowther, Tim Dwyer, Pourang Irani, Barrett Ens*

Capacitive Touch Sensing on General 3D Surfaces..... Article 103  
*Gianpaolo Palma, Narges Pourjafarian, Jürgen Steimle, Paolo Cignoni*

## Cloth Simulation

Progressive Dynamics for Cloth and Shell Animation..... Article 104  
*Jiayi Eris Zhang, Doug L. James, Danny M. Kaufman*

Super-Resolution Cloth Animation with Spatial and Temporal Coherence..... Article 105  
*Jiawang Yu, Zhendong Wang*

\* and \*\* denote equal contribution

# Table of Contents

## 3D Shape Analysis

DiffCAD: Weakly-Supervised Probabilistic CAD Model Retrieval and Alignment from an RGB Image..... Article 106  
*Daoyi Gao, Dávid Rozenberszki, Stefan Leutenegger, Angela Dai*

3Doodle: Compact Abstraction of Objects with 3D Strokes..... Article 107  
*Changwoon Choi, Jaeah Lee, Jaesik Park, Young Min Kim*

Split-and-Fit: Learning B-Reps via Structure-Aware Voronoi Partitioning..... Article 108  
*Yilin Liu, Jiale Chen, Shanshan Pan, Daniel Cohen-Or, Hao Zhang, Hui Huang*

Variational Feature Extraction in Scientific Visualization..... Article 109  
*Nico Dassler, Tobias Günther*

Implicit Swept Volume SDF: Enabling Continuous Collision-Free Trajectory Generation for Arbitrary Shapes..... Article 110  
*Jingping Wang\*, Tingrui Zhang\*, Qixuan Zhang, Chuxiao Zeng, Jingyi Yu, Chao Xu, Lan Xu, Fei Gao*

## Appearance Models

A Fully-correlated Anisotropic Micrograin BSDF Model..... Article 111  
*Simon Lucas, Mickaël Ribardière, Romain Pacanowski, Pascal Barla*

From microfacets to participating media: A unified theory of light transport with stochastic geometry..... Article 112  
*Dario Seyb, Eugene d'Eon, Benedikt Bitterli, Wojciech Jarosz*

A Free-Space Diffraction BSDF..... Article 113  
*Shlomi Steinberg, Ravi Ramamoorthi, Benedikt Bitterli, Arshiya Mollazainali, Eugene d'Eon, Matt Pharr*

One Noise to Rule Them All: Learning a Unified Model of Spatially-Varying Noise Patterns..... Article 114  
*Arman Maesumi, Dylan Hu\*, Krishi Saripalli\*, Vladimir G. Kim, Matthew Fisher, Sören Pirk, Daniel Ritchie*

## Simulating Deformation

Position-Based Nonlinear Gauss-Seidel for Quasistatic Hyperelasticity..... Article 115  
*Yizhou Chen, Yushan Han, Jingyu Chen, Zhan Zhang, Alex McAdams, Joseph Teran*

Vertex Block Descent..... Article 116  
*Anka He Chen, Ziheng Liu, Yin Yang, Cem Yuksel*

Simplicits: Mesh-Free, Geometry-Agnostic, Elastic Simulation..... Article 117  
*Vismay Modi, Nicholas Sharp, Or Perel, Shinjiro Sueda, David I. W. Levin*

A Neural Network Model for Efficient Musculoskeletal-Driven Skin Deformation..... Article 118  
*Yushan Han, Yizhou Chen, Carmichael Ong, Jingyu Chen, Jennifer Hicks, Joseph Teran*

\* and \*\* denote equal contribution

# Table of Contents

## Generative 3D Geometry and Editing

*BrepGen: A B-rep Generative Diffusion Model with Structured Latent Geometry* ..... Article 119  
*Xiang Xu, Joseph G. Lambourne, Pradeep Kumar Jayaraman, Zhengqing Wang, Karl D.D. Willis, Yasutaka Furukawa*

*CLAY: A Controllable Large-scale Generative Model for Creating High-quality 3D Assets*..... Article 120  
*Longwen Zhang\*, Ziyu Wang\*, Qixuan Zhang, Qiwei Qiu, Anqi Pang, Haoran Jiang, Wei Yang, Lan Xu, Jingyi Yu*

*TIP-Editor: An Accurate 3D Editor Following Both Text-Prompts And Image-Prompts*..... Article 121  
*Jingyu Zhuang, Di Kang, Yan-Pei Cao, Guanbin Li, Liang Lin, Ying Shan*

*Mesh Neural Cellular Automata*..... Article 122  
*Ehsan Pajouheshgar\*, Yitao Xu\*, Alexander Mordvintsev, Eyvind Niklasson, Tong Zhang, Sabine Süsstrunk*

## Rendering, Denoising & Path Guiding

*Temporally Stable Metropolis Light Transport Denoising using Recurrent Transformer Blocks*..... Article 123  
*Chuhao Chen, Yuze He, Tzu-Mao Li*

*Target-Aware Image Denoising for Inverse Monte Carlo Rendering*..... Article 124  
*Jongmin Gu, Jonghee Back, Sung-Eui Yoon, Bochang Moon*

*Real-Time Path Guiding Using Bounding Voxel Sampling*..... Article 125  
*Haolin Lu, Wesley Chang, Trevor Hedstrom, Tzu-Mao Li*

*Specular Polynomials*..... Article 126  
*Zhimin Fan, Jie Guo, Yiming Wang, Tianyu Xiao, Hao Zhang, Chenxi Zhou, Zhenyu Chen, Pengpei Hong, Yanwen Guo, Ling-Qi Yan*

*Spin-Weighted Spherical Harmonics for Polarized Light Transport*..... Article 127  
*Shinyoung Yi, Donggun Kim, Jiwoong Na, Xin Tong, Min H. Kim*

## Perception, Image, Video

*Theory of Human Tetrachromatic Color Experience and Printing*..... Article 128  
*Jessica Lee, Nicholas Jennings, Varun Srivastava, Ren Ng*

*ColorVideoVDP: A visual difference predictor for image, video and display distortions* ..... Article 129  
*Rafał K. Mantiuk, Param Hanji, Maliha Ashraf, Yuta Asano, Alexandre Chapiro*

*Analogist: Out-of-the-box Visual In-Context Learning with Image Diffusion Model*..... Article 130  
*Zheng Gu, Shiyuan Yang, Jing Liao, Jing Huo, Yang Gao*

*Learning Images Across Scales Using Adversarial Training* ..... Article 131  
*Krzysztof Wolski, Adarsh Djeacoumar, Alireza Javanmardi, Hans-Peter Seidel, Christian Theobalt, Guillaume Cordonnier, Karol Myszkowski, George Drettakis, Xingang Pan, Thomas Leimkühler*

\* and \*\* denote equal contribution

# Table of Contents

## Simulation With Contact

Contact detection between curved fibres: high order makes a difference ..... Article 132  
*Octave Crespel\**, *Emile Hohmadel\**, *Thibaut Métivet*, *Florence Bertails-Descoubes*

## Spatial Data Structures

fVDB: A Deep-Learning Framework for Sparse, Large-Scale, and High-Performance Spatial Intelligence ..... Article 133  
*Francis Williams*, *Jiahui Huang*, *Jonathan Swartz*, *Gergely Klár*, *Vijay Thakkar*, *Matthew Cong*, *Xuanchi Ren*,  
*Ruilong Li*, *Clement Fuji-Tsang*, *Sanja Fidler*, *Eftychios Sifakis*, *Ken Museth*

Neural Gaussian Scale-Space Fields..... Article 134  
*Felix Mujkanovic*, *Ntumba Elie Nsambi*, *Christian Theobalt*, *Hans-Peter Seidel*, *Thomas Leimkühler*

## Controllable Image Generation and Completion

RealFill: Reference-Driven Generation for Authentic Image Completion..... Article 135  
*Luming Tang*, *Nataniel Ruiz*, *Qinghao Chu*, *Yuanzhen Li*, *Aleksander Hołyński*, *David E. Jacobs*,  
*Bharath Hariharan*, *Yael Pritch*, *Neal Wadhwa*, *Kfir Aberman*, *Michael Rubinstein*

## Character Animation: 2D, 3D, Robot

Semantic Gesticulator: Semantics-Aware Co-Speech Gesture Synthesis ..... Article 136  
*Zeyi Zhang\**, *Tenglong Ao\**, *Yuyao Zhang*, *Qingzhe Gao*, *Chuan Lin*, *Baoquan Chen*, *Libin Liu*

Interactive Design of Stylized Walking Gaits for Robotic Characters..... Article 137  
*Michael A. Hopkins\**, *Georg Wiedebach\**, *Kyle Cesare*, *Jared Bishop*, *Espen Knoop*, *Moritz Bächer*

## Geometry: Editing and Deformation

Biharmonic Coordinates and their Derivatives for Triangular 3D Cages ..... Article 138  
*Jean-Marc Thiery*, *Élie Michel*, *Jiong Chen*

Curvature-Driven Conformal Deformations ..... Article 139  
*Etienne Corman*

Repulsive Shells..... Article 140  
*Josua Sassen*, *Henrik Schumacher*, *Martin Rumpf*, *Keenan Crane*

## Computational Cameras and Displays

Split-Aperture 2-in-1 Computational Cameras ..... Article 141  
*Zheng Shi\**, *Ilya Chugunov\**, *Mario Bijelic*, *Geoffroi Côté*, *Jiwoon Yeom*, *Qiang Fu*, *Hadi Amata*,  
*Wolfgang Heidrich*, *Felix Heide*

\* and \*\* denote equal contribution

# Table of Contents

## Character Control

Categorical Codebook Matching for Embodied Character Controllers..... Article 142  
*Sebastian Starke, Paul Starke, Nicky He, Taku Komura, Yuting Ye*

Interactive Character Control with Auto-Regressive Motion Diffusion Models..... Article 143  
*Yi Shi, Jingbo Wang, Xuekun Jiang, Bingkun Lin, Bo Dai, Xue Bin Peng*

MoConVQ: Unified Physics-Based Motion Control via Scalable Discrete Representations ..... Article 144  
*Heyuan Yao, Zhenhua Song, Yuyang Zhou, Tenglong Ao, Baoquan Chen, Libin Liu*

## Procedural Geometry

Terrain Amplification using Multi-scale Erosion..... Article 145  
*Hugo Schott, Eric Galin, Eric Guérin, Adrien Peytavie, Axel Paris*

Interactive Invigoration: Volumetric Modeling of Trees with Strands..... Article 146  
*Bosheng Li, Nikolas A. Schwarz, Wojtek Patubicki, Sören Pirk, Bedrich Benes*

Surface-Filling Curve Flows via Implicit Medial Axes ..... Article 147  
*Yuta Noma, Silvia Sellán, Nicholas Sharp, Karan Singh, Alec Jacobson*

## Radiance Field Processing

Bilateral Guided Radiance Field Processing ..... Article 148  
*Yuehao Wang, Chaoyi Wang, Bingchen Gong, Tianfan Xue*

Binary Opacity Grids: Capturing Fine Geometric Detail for Mesh-Based View Synthesis ..... Article 149  
*Christian Reiser, Stephan Garbin, Pratul P. Srinivasan, Dor Verbin, Richard Szeliski, Ben Mildenhall, Jonathan T. Barron, Peter Hedman, Andreas Geiger*

TensoSDF: Roughness-aware Tensorial Representation for Robust Geometry and Material Reconstruction..... Article 150  
*Jia Li, Lu Wang, Lei Zhang, Beibei Wang*

## Sound, Light, Radiofrequency

Cricket: A Self-Powered Chirping Pixel..... Article 151  
*Shree K. Nayar, Jeremy Klotz, Nikhil Nanda, Mikhail Fridberg*

## Art, Illusion, Fabrication

Computational Illusion Knitting..... Article 152  
*Amy Zhu, Yuxuan Mei, Benjamin Jones, Zachary Tatlock, Adriana Schulz*

Creating LEGO™ Figurines from Single Images..... Article 153  
*Jiahao Ge, Mingjun Zhou, Wenrui Bao, Hao Xu, Chi-Wing Fu*

Technical Papers Committee ..... xv

Technical Papers Reviewers ..... xvii

Cover Image Credits ..... xvii

Author Index..... xviii

\* and \*\* denote 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 (Ohay) 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