

# High-Performance Graphics 2025

Copenhagen, Denmark  
June 23 — 25, 2025

---

## General Chairs

Michael Doggett, Lund University  
Christiaan Gribble, AMD

## Papers Chairs

Aaron Knoll, AMD  
Christoph Peters, TU Delft

## Preface

We are pleased to present the papers program for High-Performance Graphics (HPG) 2025, the leading international forum for performance-oriented graphics and imaging systems research. We are excited about innovative algorithms as well as efficient, real-world implementations, and the underlying computing technology: Hardware, languages, compilers, and software systems. The HPG conference series brings together researchers, engineers, and architects to discuss the complex interactions between graphics hardware and software. Formed by the merger of the Eurographics Workshop on Graphics Hardware and the IEEE Symposium on Interactive Ray Tracing in 2009, this year marks the seventeenth in which industrial and academic researchers, practitioners, and students gather to advance the field of high-performance computer graphics through HPG's unique platform.

The HPG conference series publishes top-tier research in high-performance graphics, and this year is no exception. We received a total of 37 paper submissions, each of which was reviewed by four members of our International Program Committee (IPC). After a short but thorough review process and a discussion phase, the IPC recommended 19 innovative papers for this year's conference. After a minor revision, all accepted papers went through a second review cycle to ensure the highest quality. Ultimately, ten papers were selected for publication in a special issue of Computer Graphics Forum, and the remaining nine papers were published in the ACM/Eurographics HPG 2025 Conference Proceedings.

The 2025 papers program covers a wide range of topics including acceleration structures, compilers, graphics systems and simulators, deep learning for graphics, GPU-accelerated procedural generation, rendering, novel-view synthesis, perception, frame interpolation, neural texture compression and texture filtering. What ties all of these works together is a performance-oriented approach. Deep learning continues to become more important for HPG: Four accepted papers utilize it.

We offer our sincere thanks to the authors of all submitted papers—truly these authors are the lifeblood of HPG's successful papers program. We also offer our deep gratitude for this year's IPC members. All of them did an amazing job writing detailed and insightful reviews and discussion comments within a relatively short time frame. Together, the contributions of our authors and IPC members culminate in an outstanding selection of publications for High-Performance Graphics 2025. Stefanie Behnke has been tremendously helpful in administrating the review process on SRM and preparing the proceedings. In addition to the papers program, the conference committee has organized three keynotes, six poster presentations, three Hot 3D presentations, the student competition and multiple social events for the evenings. All of this is made possible by our sponsors, and we are most grateful for their contributions.

We look forward to 2.5 days of exciting HPG program in Copenhagen, Denmark, followed immediately by another 2.5 days of EGSR program. More than 170 attendees have registered for HPG, making this one of the largest European HPG conferences to date and an excellent forum for discussion and networking.

Sincerely,  
Aaron Knoll & Christoph Peters  
HPG 2025 Papers Chairs

## Sponsors



## Table of Contents

### Bounding Volume Hierarchies

- DOBB-BVH: Efficient Ray Traversal by Transforming Wide BVHs into Oriented Bounding Box Trees using Discrete Rotations* e70220  
Michael A. Kern, Alain Galvan, David R. Oldcorn, Daniel Skinner, Rohan Mehalwal, Leo Reyes Lozano, and Matthäus G. Chajdas
- Fused Collapsing for Wide BVH Construction* e70213  
Wilhem Barbier and Mathias Paulin
- UBVH: Unified Bounding Volume and Scene Geometry Representation for Ray Tracing* e70216  
Martin Káčerik and Jiří Bittner

### Graphics Simulators, Systems and Compilers

- Arches: A Cycle-Level Hardware Simulation Framework for Exploring Massively Parallel Ray Tracing Architectures* e70212  
Jacob Haydel, Gaurav Bhokare, Kunnong Zeng, Pengpei Hong, Sushant Kondguli, Brian Budge, Erik Brunvand, and Cem Yuksel
- RenderMan XPU: A Hybrid CPU+GPU Renderer for Interactive and Final-frame Rendering* e70218  
Per Christensen, Julian Fong, Charlie Kilpatrick, Francisco Gonzalez, Srinath Ravichandran, Akshay Shah, Ethan Jaszewski, Stephen Friedman, James Burgess, Trina M. Roy, Tom Nettleship, Meghana Seshadri, and Susan Salituro

### Procedural Generation and Sculpting

- Triangle Rejection Sampling for Density-Equipped Meshes on GPU* e70217  
Jeremie Schertzer, Theo Thonat, and Tamy Boubekeur

### Perception in Motion

- CGVQM+D: Computer Graphics Video Quality Metric and Dataset* e70221  
Akshay Jindal, Nabil Sadaka, Manu Mathew Thomas, Anton Sochenov, and Anton Kaplanyan
- Image-Based Spatio-Temporal Upsampling for Split Rendering* e70215  
Michael Steiner, Thomas Köhler, Lukas Radl, Brian Budge, and Markus Steinberger

### Splats and Points

- Splatshop: Efficiently Editing Large Gaussian Splat Models* e70214  
Markus Schütz, Christoph Peters, Florian Hahlbohm, Elmar Eisemann, Marcus Magnor, and Michael Wimmer

### Real-Time Rendering

- Spherical Harmonic Exponentials for Efficient Glossy Reflections* e70219  
Ari Silvennoinen, Peter-Pike Sloan, Michal Iwanicki, and Derek Nowrouzezahrai

## International Program Committee

Afra, Attila – Intel  
Akenine-Möller, Tomas – NVIDIA  
Apolinário, Antônio – Federal University of Bahia  
Assarsson, Ulf – Chalmers University of Technology  
Benthin, Carsten – AMD  
Bikker, Jacco – Utrecht University  
Billeter, Markus – University of Leeds  
Binder, Nikolaus – NVIDIA  
Bittner, Jiří – Czech Technical University in Prague  
Boubek, Tamy – Adobe Research  
Chajdas, Matthäus – Intel  
Clarberg, Petrik – NVIDIA Corporation  
Clua, Esteban – universidade federal fluminense  
Crassin, Cyril – NVIDIA  
Dachsbacher, Carsten – Karlsruhe Institute of Technology  
Didyk, Piotr – Università della Svizzera italiana  
Dietrich, Andreas – ESI Group  
Doggett, Michael – Lund University  
Dupuy, Jonathan – Intel Corporation  
Fenney, Simon – Imagination Technologies  
Franke, Linus – Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU)  
Garrard, Andrew – Imagination Technologies  
Grittmann, Pascal – Saarland University  
Gruen, Holger – AMD  
Hachisuka, Toshiya – University of Waterloo  
Hanika, Johannes – Karlsruhe Institute of Technology  
Harada, Takahiro – Advanced Micro Devices, Inc.  
Havran, Vlastimil – Czech Technical University in Prague  
Heitz, Eric – Intel  
Herholz, Sebastian – Intel Corporation  
Jindal, Akshay – Intel  
Kensler, Andrew – AMD  
Kerbl, Bernhard – TU Wien  
Kim, Sungye – AMD  
Knoll, Aaron – AMD  
Kuth, Bastian – Coburg University of Applied Sciences  
Lee, Sungkil – Sungkyunkwan University  
Liktó, Gabor – Intel Corporation  
Mantiuk, Rafal – University of Cambridge  
Manzke, Michael – Trinity College Dublin  
Marroquim, Ricardo – Delft University of Technology  
Meister, Daniel – Advanced Micro Devices, Inc.  
Membarth, Richard – THI, DFKI  
Meyer, Quirin – Coburg University  
Mueller, Joerg H. – Huawei Technologies

## International Program Committee

Mullia, Krishna – Adobe Research  
Neyret, Fabrice – .  
Nimier-David, Merlin – NVIDIA  
Oberberger, Max – Advanced Micro Devices  
Olano, Marc – UMBC  
Park, Woo-Chan – Sejong University  
Patney, Anjul – NVIDIA  
Paulin, Mathias – IRIT, Université de Toulouse, CNRS, UT, Toulouse, France  
Peters, Christoph – Delft University of Technology  
Pharr, Matt – NVIDIA  
Reshetov, Alexander – NVidia  
Ritschel, Tobias – University College London  
Salvi, Marco – NVIDIA  
Schied, Christoph – NVIDIA Research  
Schütz, Markus – TU Wien  
Selgrad, Kai – OTH Regensburg  
Stamminger, Marc – Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU)  
Steinberg, Shlomi – University of Waterloo  
Steinberger, Markus – Graz University of Technology, Huawei Technologies  
Usher, Will – Luminary Cloud  
Vaidyanathan, Karthik – NVIDIA  
Vardis, Konstantinos – Huawei Technologies  
Viitanen, Timo – NVIDIA  
Vining, Nicholas – NVIDIA  
Wang, Beibei – Nanjing University  
Wang, Lili – School of computer science and engineering, Beihang University  
Weidlich, Andrea – Nvidia  
Weier, Philippe – Weta Digital  
Woop, Sven – Intel Corporation  
Wyman, Chris – NVIDIA  
Zeltner, Tizian – NVIDIA  
Zirr, Tobias – Karlsruhe Institute of Technology, Intel Corporation

## Author Index

Barbier, Wilhem .....	e70213	Magnor, Marcus .....	e70214
Bhokare, Gaurav .....	e70212	Mehalwal, Rohan .....	e70220
Bittner, Jiří .....	e70216	Nettleship, Tom .....	e70218
Boubekeur, Tamy .....	e70217	Nowrouzezahrai, Derek .....	e70219
Brunvand, Erik .....	e70212	Oldcorn, David R. ....	e70220
Budge, Brian .....	e70212, e70215	Paulin, Mathias .....	e70213
Burgess, James .....	e70218	Peters, Christoph .....	e70214
Chajdas, Matthäus G. ....	e70220	Radl, Lukas .....	e70215
Christensen, Per .....	e70218	Ravichandran, Srinath .....	e70218
Eisemann, Elmar .....	e70214	Roy, Trina M. ....	e70218
Fong, Julian .....	e70218	Sadaka, Nabil .....	e70221
Friedman, Stephen .....	e70218	Salituro, Susan .....	e70218
Galvan, Alain .....	e70220	Schertzer, Jeremie .....	e70217
Gonzalez, Francisco .....	e70218	Schütz, Markus .....	e70214
Hahlbohm, Florian .....	e70214	Seshadri, Meghana .....	e70218
Haydel, Jacob .....	e70212	Shah, Akshay .....	e70218
Hong, Pengpei .....	e70212	Silvennoinen, Ari .....	e70219
Iwanicki, Michal .....	e70219	Skinner, Daniel .....	e70220
Jaszewski, Ethan .....	e70218	Sloan, Peter-Pike .....	e70219
Jindal, Akshay .....	e70221	Sochenov, Anton .....	e70221
Kaplanyan, Anton .....	e70221	Steinberger, Markus .....	e70215
Kern, Michael A. ....	e70220	Steiner, Michael .....	e70215
Kilpatrick, Charlie .....	e70218	Thomas, Manu Mathew .....	e70221
Kondguli, Sushant .....	e70212	Thonat, Theo .....	e70217
Káčarik, Martin .....	e70216	Wimmer, Michael .....	e70214
Köhler, Thomas .....	e70215	Yuksel, Cem .....	e70212
Lozano, Leo Reyes .....	e70220	Zeng, Kunnong .....	e70212

## Keynote

### Hallucinating the Future of Real-Time Rendering

*Angelo Pesce*

Roblox

#### Abstract

Nobody has a crystal ball, and long-term predictions are more likely than not of being wrong. Yet, it is interesting to reflect on an industry, chart a trajectory, and extrapolate forwards. From the vantage point of almost 30 years now writing real-time rendering code, Angelo will try to devise a method to divine the industry's future, and place some bets on what we will (or should) see next. An "open problems" talk that hopes to get at least something right, or at the very least, to provoke some interesting thoughts.

#### Short Biography

Angelo Pesce currently serves as a sr. Technical Director for the Roblox "Digital Matter" team, after having helped Roblox establish its rendering team. Previous to joining the metaverse, he was a sr. Technical Director for Activision Central Technology where he mostly helped the various Call of Duty studios with rendering R&D. His interest in Computer Graphics started in his teens, in the creative arts of the "demoscene". He has been credited on a dozen AAA game titles, having worked on rendering solutions for companies such as Milestone, Electronic Arts, Capcom and Relic Entertainment. Angelo is actively involved in the computer graphics community, having presented at multiple venues including Siggraph, GDC and Digital Dragons, served as chair for conferences and editor for publications. He is one of the co-authors of the latest edition of the book "Real-Time Rendering", and he co-organizes REAC: the Rendering Engine Architecture Conference.

## Keynote

### **Inferred reality: Embracing the new reality in real time computer graphics**

*Rama Harihara*  
AMD

#### **Abstract**

As we push towards increased photorealism in games and path tracing quality becomes a requirement at real-time frame rates, scaling silicon process alone will no longer improve perf\$. Emerging paradigms for rendering with ML assist are very promising to do more with less. This presentation gives a brief synopsis of how the industry and academia have started leveraging ML techniques to create higher levels of rendering approximation. We then outline the potential future disruption leveraging generative models for future rendering.

#### **Short Biography**

As an AMD Fellow, Rama Harihara leads the ML applied research team, pioneering advancements in real-time graphics, neural rendering, and AI-driven 3D content creation. She shapes the future of rendering technology at AMD, guiding research from initial concepts to product-ready solutions. Her role involves close collaboration with academic institutions, industry partners, and internal teams – including hardware and software architects – to ensure the successful adoption of these cutting-edge technologies. Before joining AMD, she led the compute architecture team and neural graphics research at Intel and drove graphics strategy and innovation at Samsung.

## Keynote

### The Future of Analytical Materials in a Neural World

*Andrea Weidlich*  
NVIDIA

#### Abstract

Recent advancements in real-time graphics have substantially closed the gap between offline and real-time rendering by incorporating techniques such as path tracing into modern game engines. Despite this progress, achieving realistic material appearance remains a considerable challenge. Offline rendering can produce high-quality results but is often labor-intensive and resource-intensive, making real-time “film-quality” materials difficult to achieve and limiting the potential of real-time photorealistic rendering.

In this talk, we will define what the term “film-quality materials” entails and explore the disparity between offline and real-time material representations. Recently, neural materials emerged as a promising solution to bridge that gap since they can encode appearance in small neural networks and reduce the need for complex shaders and large texture storage. However, while neural materials hold great promise, it is important not to view them as the outright successor to traditional analytical materials. We will discuss the ongoing importance of traditional analytical material research and how it can complement neural approaches. By examining the integration of analytical materials within the emerging neural-driven paradigm, we will highlight how these methods can coexist to advance the capabilities of both real-time and offline rendering.

#### Short Biography

Andrea Weidlich is a Principal Researcher at Nvidia in the Real-Time Rendering Research department working on neural appearance models. Before joining Nvidia, she worked for Weta Digital where she designed the material system attached to Weta’s proprietary physically-based renderer, Manuka, and contributed to films such as “War of the Planet of the Apes” and “Avatar 2”. Her main research areas are appearance modelling and material prototyping. Andrea holds a Master of Arts in Applied Media from University of Applied Arts, Vienna, and a Ph.D. in Computer Science from Vienna University of Technology.

## Hot 3D

### **Slang: Bridging Graphics and Machine Learning**

*Shannon Woods*  
NVIDIA

#### **Abstract**

We present Slang, an open-source shading language that brings automatic differentiation to real-time graphics programming. While neural techniques are increasingly vital for achieving photorealistic rendering, implementing them has traditionally required maintaining separate codebases in ML frameworks and graphics APIs. We demonstrate how Slang's built-in automatic differentiation enables graphics developers to implement gradient-based optimization directly in shader code, while maintaining the performance characteristics needed for real-time applications. We show how Slang reduces implementation complexity by automatically generating differential code, eliminating the need to manually maintain forward and backward passes. The presentation will walk through a simplified Neural Materials implementation, demonstrating how Slang's Python bindings enable seamless integration with ML workflows while preserving the flexibility and control required for graphics programming. By bridging the gap between graphics and machine learning ecosystems, Slang enables developers to more easily incorporate neural techniques into their rendering pipelines.

## Hot 3D

# Intel Open Image Denoise 2 and Beyond: AI-Accelerated Ray Tracing for Everyone

*Attila Afra*

### Abstract

Since its launch in 2019, Intel Open Image Denoise has become one of the most widely adopted AI-based denoising solutions for ray tracing, culminating in recently receiving an Academy Sci-Tech Award for its contribution to filmmaking. The success of the library can be largely attributed to its open-source nature, high performance and image quality, rich cross-vendor support, and simple API, making it unique among denoisers in the industry. Initially conceived as a CPU denoiser, Intel Open Image Denoise 2 has introduced optimized backends for Intel, NVIDIA, AMD, and Apple GPUs as well, enabling real-time performance on various hardware architectures. The next major version is now under development, which aims to bring both quality and performance to a new level by employing a more advanced neural network architecture, temporal denoising, and many other improvements. In this talk, we will delve into the technical details of both current and upcoming features of Intel Open Image Denoise, offering insights into the challenges of productizing AI denoising for a broad user base.

## Hot 3D

### **MaterialX: The Past, Present and Future of Portable Materials**

*Alain Hostettler*

Industrial Light & Magic

#### **Abstract**

MaterialX was launched in 2012 at Industrial Light & Magic with the aim to standardise material representations. This was driven by the need to decouple the building and rendering of assets from specific applications and to make assets available across a wide range of media types. Since then, the project has reached several notable milestones; in 2016 the first public specification of MaterialX was released, in 2017 MaterialX became open source with substantial contributions from Autodesk, Pixar, Sony Pictures Imageworks, and others. In 2021 the project was officially adopted by the Academy Software Foundation. In this talk, we will give a brief introduction to MaterialX and its history, followed by an overview of the current state of the project and some of the recent highlights. We take a look into the future and present our roadmap of upcoming developments. The talk concludes with a sneak peek of unpublished recent work.