

Physically Based Shading in Theory and Practice

SIGGRAPH 2025 COURSE

Course Organizers

STEPHEN HILL

Lucasfilm

STEPHEN MCAULEY

Sony Santa Monica Studio

Presenters

LAURENT BELCOUR

Intel Corporation

NATY HOFFMAN

Independent

ALAIN HOSTETTLER

Industrial Light & Magic

PETER KUTZ

Adobe

KENTARO SUZUKI

Polyphony Digital Inc.

HAJIME UCHIMURA

Polyphony Digital Inc.

ANDREA WEIDLICH

NVIDIA

KENICHIRO YASUTOMI

Polyphony Digital Inc.

Additional Contributors

PASCAL BARLA

Inria

ALBAN FICHET

Intel Corporation

JAMIE PORTSMOUTH

Autodesk

Course Description

Physically based shading has transformed the way we approach production rendering and simplified the lives of artists in the process. By employing shading models that adhere to physical principles, one can readily create high quality, realistic materials that maintain their appearance under a variety of lighting environments. In contrast, traditional ad hoc models required extensive tweaking to achieve comparable results – due to less intuitive behavior or unnecessary complexity – and were liable to break under different illumination.

Consequently, physically based models have become widely adopted in film and game production, particularly as they are often no more difficult to implement or evaluate. That being said, physically based shading is not a solved problem, and thus the aim of this course is to share the latest theory as well as lessons from production.

LEVEL OF DIFFICULTY: Intermediate

Intended Audience

Practitioners from the video game, CG animation, and VFX fields, as well as researchers interested in shading models.

Prerequisites

An understanding of shading models and their use in film or game production.

Course Website

All course materials will be available [here](#).

Contact

Please address questions or comments to pbs@selfshadow.com.

Organizers


STEPHEN HILL is a Principal Rendering Engineer within Lucasfilm's Advanced Development Group, where he is engaged in physically based rendering R&D for productions such as *Carne y Arena*, and more recently *The Mandalorian*. He was previously a 3D Technical Lead at Ubisoft Montreal, where he contributed to a number of *Splinter Cell* titles as well as *Assassin's Creed Unity*.

STEPHEN MCAULEY started in video games in 2006 at Bizarre Creations before moving to Ubisoft in 2011, where he spearheaded the graphical vision on the *Far Cry* brand. In 2020, he joined Sony Santa Monica as a Lead Rendering Engineer and is now Technical Director. He focuses on physically based lighting and shading, data-driven rendering architecture and overall improvements in visual quality. He is also passionate about sharing his knowledge with the industry as a whole, running internal and external training and conferences.

Presentation Schedule

- 09:00–09:20 **Physics and Math of Shading** (*Hoffman*)
- 09:20–09:45 **OpenPBR: A Closer Look at Novel Features and Implementation Details** (*Kutz*)
- 09:45–10:05 **EON: Advancing Rough Diffuse Reflection with Energy Preservation and Clipped LTC Sampling** (*Kutz and Hill*)
- 10:05–10:30 **Spectral Rendering in a Non-Spectral Renderer:
How Can we Author and Render Fluorescence in RGB?** (*Belcour*)
- 10:30–10:45 *Break*
- 10:45–11:10 **Strand: A Production Model for Shading Hair, Fur and Feathers** (*Hostettler*)
- 11:10–11:35 **Bridging the Gap Between Offline and Real Time with Neural Materials** (*Weidlich*)
- 11:35–12:15 **Driving Toward Reality: Physically Based Tone Mapping and Perceptual Fidelity in *Gran Turismo 7*** (*Yasutomi, Suzuki and Uchimura*)

Abstracts



Physics and Math of Shading

Naty Hoffman

We will go over the fundamentals behind physically based shading models, starting with a qualitative description of the underlying physics, followed by a quantitative description of the relevant mathematical models, and finally discussing how these mathematical models can be implemented for shading.

NATY HOFFMAN has recently retired after a storied career in real-time rendering. He has worked at Meta (improving the appearance of Meta Avatars), Lucasfilm (designing and implementing advanced rendering algorithms for ILM StageCraft and VR experiences), 2K Games (driving technology development across the publishing label), Activision (doing graphics R&D for many games including the *Call of Duty* series), SCE Santa Monica Studio (coding graphics technology for *God of War III*), Naughty Dog (developing PS3 first-party libraries), Westwood Studios (leading graphics development on *Earth and Beyond*) and Intel (driving Pentium pipeline modifications and assisting with the SSE and SSE2 instruction set definitions). During his career, Naty was also responsible for many influential publications and presentations, on rendering topics including physically based shading, cinematic lighting, and color perception in graphics.




OpenPBR: A Closer Look at Novel Features and Implementation Details

Peter Kutz and Jamie Portsmouth

OpenPBR is a standard über shader designed to cover the vast majority of practical materials in VFX and animation. In this talk, we go beyond the official specification to explore some of its most challenging and noteworthy features – such as iridescence, dispersion, coat darkening/roughening heuristics, energy conservation, and subsurface scattering. We'll discuss the underlying physics, implementation pitfalls, performance considerations, and practical code-level tips to help you refine your shading pipeline. Whether you're new to OpenPBR or already using it, this deep dive will equip you with the insights needed to take full advantage of this unified surface model.

PETER KUTZ: After writing his first photorealistic path tracer as a hobby at the University of Pennsylvania and interning at Pixar, Peter spent five years at Walt Disney Animation Studios working on the Hyperion renderer – contributing to its architecture, implementing features for multiple films, and co-authoring several papers on volume rendering and beyond. He later joined the Apple Vision Pro team to develop hybrid rasterization and ray tracing for AR applications. For the past five years, Peter has been at Adobe, crafting proprietary path-tracing solutions for Adobe's in-house GPU-first renderer and helping to define standards such as the Adobe Standard Material and OpenPBR, bridging cutting-edge research with practical production needs in physically based shading.

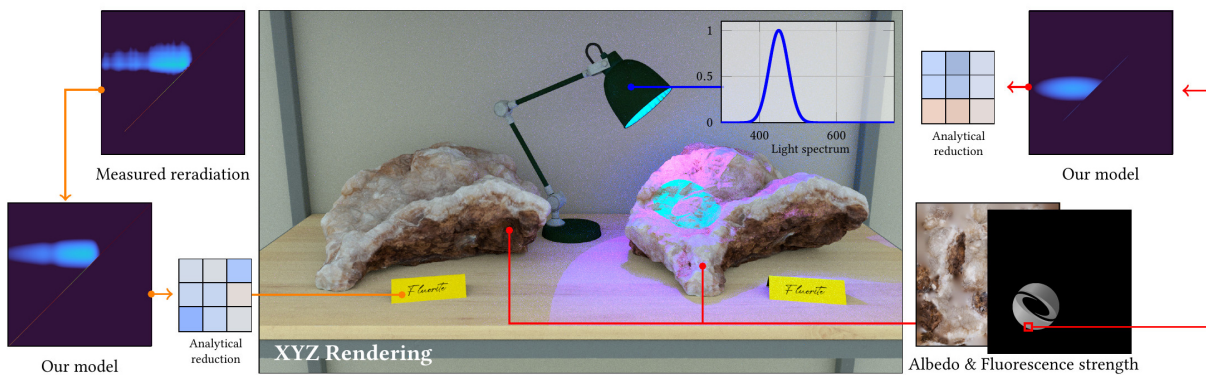
JAMIE PORTSMOUTH is a Principal Software Engineer at Autodesk, working on Autodesk's Arnold renderer. He is mostly focused on physically based rendering, most recently helping to define the OpenPBR Surface uber-shader standard.



EON: Advancing Rough Diffuse Reflection with Energy Preservation

Peter Kutz and Stephen Hill

Building on our recent publication, this talk explores how the Energy-preserving Oren–Nayar (EON) BRDF addresses key limitations of classic rough diffuse models – enforcing analytical energy conservation, ensuring reciprocity, and avoiding common artifacts. We introduce our newly developed Clipped LTC (CLTC) technique for importance sampling, reducing variance and improving performance in production rendering. Along the way, we’ll compare EON to other diffuse models, discuss recent updates since the paper’s release, and offer practical guidance for integrating it into modern pipelines.



Spectral Rendering in a Non-Spectral Renderer: How Can we Author and Render Fluorescence in RGB?

Laurent Belcour, Alban Fichet and Pascal Barla

In this talk, we will describe how fluorescence can be added to a non-spectral rendering engine. The addition of fluorescence poses some challenges, such as how to properly simulate light transport, how to model such phenomena, and how to author scenes to include this effect with fine control. We will see how adjoint vector light transport should be handled in a forward path tracer and how one can handle UV light sources there. We will detail a parametric model for fluorescence and what parameterization artists can use to edit those new appearances. Finally, we will show how this idea of adding fluorescence can lead us toward handling metameric materials, without relying on spectral rendering at all.

LAURENT BELCOUR is a computer graphics researcher working at Intel Corporation in Grenoble, France. He defended his PhD at the University of Grenoble in 2012 and then worked at Inria Bordeaux, the University of Montréal, and Unity Technologies. His research interests include material models, (quasi) Monte Carlo rendering, importance sampling, and machine learning.

ALBAN FICHET is a software research engineer at Intel in the Grenoble GraphiX Research Team. Before that, he worked at Unity Technologies. Starting from his PhD at Grenoble University and his postdoc at Institut d'Optique Graduate School in Bordeaux Alban worked on material appearance and capture, as well as accurate light transport simulation.

PASCAL BARLA received his PhD in 2006 on the topic of Expressive Rendering at INP Grenoble (France). After being recruited as a permanent researcher at Inria Bordeaux Sud Ouest in 2007, his research has expanded to the more general domain of visual appearance, with interests in both optics and perception.



Courtesy of Amazon Prime Video

Strand: A Production Model for Shading Hair, Fur and Feathers

Alain Hostettler

At ILM, we use our in-house modular material framework Lama for look development across all of our productions. As part of Lama, we provide a hair shader based on the model proposed by Chiang et al. [Chi+15]. This shader has been used extensively to render human hair as well as animal fur and feathers. However, artists have expressed difficulties in matching animal fur reference images. Furthermore, this shader is monolithic and does not follow the modular approach of Lama, which limits artistic flexibility.

This motivated us to investigate the feasibility of using more recent fur rendering research [YJR17] in production and to separate the hair shader into multiple components that can be arbitrarily layered and mixed by artists. This ultimately led us to the implementation of our modular hair shader that includes a medulla component for more accurate fur rendering and additional controls for feather rendering.

In this talk, we first give a brief introduction to hair rendering in production. This is followed by an overview of the new medulla component as proposed by Yan et al. [YJR17] and the changes we had to apply to make the model suitable for production use. Most notably, we replaced the scientific parameterisation in the original proposal with a more artist-friendly parameterisation. The monochromatic scattering and absorption coefficients of the medulla were replaced with a density and a colour parameter. The switch from a monochromatic to a coloured medulla response required additional changes to the original model.

Next, we discuss how we separated the hair model into multiple components. From a user standpoint, the separation is based on anatomical considerations; the model is split into a component for the cortex, which includes the new medulla, a component for the cuticles and a coating component (to model water or dirt layers). Internally, we take advantage of the widely used lobe separation originally proposed by Marschner et al. [Mar+03] and leverage the existing layering and mixing operators of Lama to combine the components.

Finally, we talk about the additional controls we added for more plausible feather rendering, which is based on the observation that barbs – a structural element of feathers – are typically flatter than human hair or animal fur. While not physically accurate, this additionally facilitates the use of thin films [BB17] to model iridescent feathers.

ALAIN HOSTETTLER is a Senior R&D Engineer at Industrial Light & Magic (ILM). He joined ILM in 2021 after finishing his MSc in Computer Science at ETH Zurich with a focus on computer graphics, and an internship at Disney Research Studios in Zurich, where he researched methods to improve the ML-based denoising of Monte Carlo renders. At ILM, Alain is working in the offline rendering and shading engineering team. He has implemented numerous plugins for RenderMan to support productions and is continuously working on improving and extending Lama, the modular material system that powers the look development at ILM.



Bridging the Gap Between Offline and Real Time with Neural Materials

Andrea Weidlich

Recent advances in real-time graphics have significantly narrowed the gap between offline and real-time rendering, with techniques such as path tracing making their way into modern game engines. However, while path tracing is becoming increasingly viable in real time, materials remain relatively simplistic and fail to capture the complexity of their offline counterparts, which limits the full potential of real-time photorealistic rendering.

In this talk, we will explore the current limitations of material representation in real-time graphics and discuss the challenges faced by simplified material models. We will then propose the use of neural materials as a solution to bypass expensive shader code and texture storage, by encoding appearance into simple neural representations. This approach allows us to simulate, in real time, material appearances typically reserved for offline rendering. We'll demonstrate this by sharing our experience integrating neural materials into Unreal Engine, where we've baked multi-slab Substrate materials into efficient neural representations. Finally, we'll propose neural materials as a standardized exchange format for material appearance, enabling cross-compatibility between different renderers.

ANDREA WEIDLICH is a Principal Researcher at NVIDIA in the Real-Time Rendering Research department. Before joining NVIDIA, she worked for Weta Digital, where she designed the material system attached to Weta's proprietary physically based renderer, Manuka. Her main research areas are appearance modelling and material prototyping. Andrea holds a Master of Arts in Applied Media from the University of Applied Arts Vienna and a PhD in Computer Science from Vienna University of Technology.



Driving Toward Reality: Physically Based Tone Mapping and Perceptual Fidelity in *Gran Turismo 7*

Kenichiro Yasutomi, Kentaro Suzuki and Hajime Uchimura

Gran Turismo 7 aims for photorealism by consistently respecting physics from materials to lighting, delivering physically computed light that matches real-world conditions to players. We place the highest priority on reproducing the appearance of cars as faithfully as possible. To accomplish this, car body colors are derived from spectroscopic measurements of paint samples provided by manufacturers, avoiding any arbitrary adjustments. Likewise, sunlight and skylight are precomputed based on data from atmospheric physics and atmospheric simulations, free from ad hoc modifications.

As a result, the rendered output exhibits a very wide dynamic range in both brightness and color. Nevertheless, we believe that accurately computed light can be displayed – again in a physically based manner – to achieve realism and immersion for human perception, and we continue to explore this direction. However, because display capabilities vary widely – from traditional SDR to the latest HDR – an approach to tone mapping with consistent rules is essential. Moreover, adhering to a physics-based philosophy throughout the process makes designing a physically based tone-mapping pipeline a crucial challenge.

This talk provides a comprehensive overview of the history, principles, and implementation details of tone mapping for SDR and HDR displays. We then introduce the specific techniques used in *Gran Turismo 7*, evaluate their effectiveness, and discuss future perspectives.

KENICHIRO YASUTOMI is a Lead Technical Artist at Polyphony Digital Inc. He studied urban engineering at university. However, he became fascinated by computer graphics, which he began using for presentations. Eventually, he entered the game industry as an artist.

KENTARO SUZUKI is a Lead Graphics Engineer at Polyphony Digital Inc. His computer graphics life began with the demoscene culture of the early 2000s. After studying computer graphics at university, he joined Polyphony Digital Inc., where he has contributed to the world-famous *Gran Turismo* racing game series. His primary interests are developing various real-time rendering techniques and exploring offline ray-tracing algorithms.

HAJIME UCHIMURA is an Image Processing Engineer at Polyphony Digital Inc. He started computer graphics with the MSX at the age of six. He joined Polyphony Digital Inc. after high-precision calculation research for his master's degree. His main topics are image processing and color science. He's the proud father of a girl and a boy.

References

- [BB17] L. Belcour and P. Barla. “A practical extension to microfacet theory for the modeling of varying iridescence”. In: *ACM Trans. Graph.* 36.4 (July 2017). doi: [10.1145/3072959.3073620](https://doi.org/10.1145/3072959.3073620).
- [Chi+15] M. J.-Y. Chiang, B. Bitterli, C. Tappan, and B. Burley. “A practical and controllable hair and fur model for production path tracing”. In: *ACM SIGGRAPH 2015 Talks*. SIGGRAPH ’15. Los Angeles, California: Association for Computing Machinery, 2015. doi: [10.1145/2775280.2792559](https://doi.org/10.1145/2775280.2792559).
- [Mar+03] S. R. Marschner, H. W. Jensen, M. Cammarano, S. Worley, and P. Hanrahan. “Light scattering from human hair fibers”. In: 22.3 (July 2003), pp. 780–791. doi: [10.1145/882262.882345](https://doi.org/10.1145/882262.882345).
- [YJR17] L.-Q. Yan, H. W. Jensen, and R. Ramamoorthi. “An efficient and practical near and far field fur reflectance model”. In: *ACM Trans. Graph.* 36.4 (July 2017). doi: [10.1145/3072959.3073600](https://doi.org/10.1145/3072959.3073600).