Costly Light Transport Phenomena In Modern Real-Time Rendering Pipelines

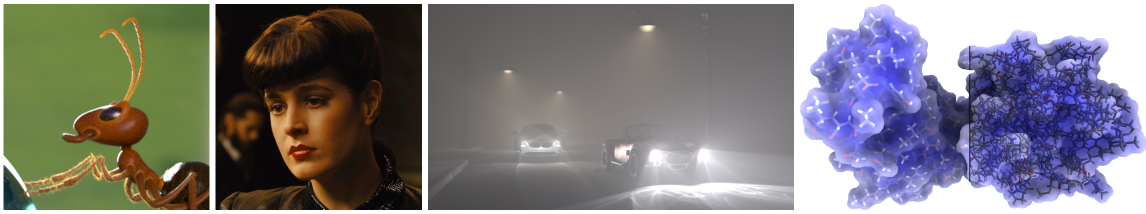

Fig. 1. Left, center left: subsurface scattering in movie productions [1]. Center right: volumetric caustics cast by car headlamps [3]. Right: interactive visualization of molecular data using a subsurface scattering approximation [2].

Photo- and sensorrealistic rendering has applications in many fields, ranging from movie production, digital prototyping, generation of AI training data to enhancing the visualization of complex structures in simulation data. In all of these examples, rendering requires some compromise between visual quality and computational budget. A key aspect of adjustment is the fidelity or accuracy of the underlying simulation of light transport which is the basis for image formation. Between acceptable rendering times in real-time (< 16ms) and offline rendering (many hours) are at least six orders of magnitude. Therefore traditionally the former has focused on visually convincing approximation and efficient usage of specialized graphics hardware (GPU), while the latter can rely more on accurate numerical simulations of the underlying physics.

Recently, advances in performance and flexibility of modern GPU hardware offer opportunities to reconsider tradeoffs of established methods and pave the way to entirely new approaches. Especially the advent of real-time raytracing raises the question of how methods that were usable only in offline rendering before can now be adapted and integrated into modern interactive rendering pipelines.

In this project, we will identify a use-case in the aforementioned range of applications and methods, and investigate the question of how to improve or advance the rendering of one or more costly light transport phenomena. Possible areas of focus include sub-surface scattering in real-time raytracing for games, real-time volumetric caustics, light diffusion for visualizations of data stored as unstructured grids, or … you name it! You will review the state of the art in the chosen area and subsequently develop your own novel rendering or visualization algorithm.

Contact: Vincent Schüßler (vincent.schuessler@kit.edu)

REFERENCES

[1] Per Christensen, Julian Fong, Jonathan Shade, Wayne Wooten, Brenden Schubert, Andrew Kensler, Stephen Friedman, Charlie Kilpatrick, Cliff Ramshaw, Marc Bannister, et al. 2018. Renderman: An advanced path-tracing architecture for movie rendering. *ACM Transactions on Graphics* 37, 3 (2018), 1–21. https://doi.org/10.1145/3182162

[2] Pedro Hermosilla, Sebastian Maisch, Pere-Pau Vázquez, and Timo Ropinski. 2018. Improving Perception of Molecular Surface Visualizations by Incorporating Translucency Effects. In *Eurographics Workshop on Visual Computing for Biology and Medicine.* https://doi.org/10.2312/vcbm.20181244

[3] Jan Novák, Derek Nowrouzezahrai, Carsten Dachsbacher, and Wojciech Jarosz. 2012. Progressive Virtual Beam Lights. *Computer Graphics Forum (Proceedings of EGSR 2012)* 31, 4 (June 2012), 1407–1413. https://doi.org/10.1111/j.1467-8659.2012.03136.x