Advances in Monte Carlo Rendering: The Legacy of Jaroslav Křivánek

Alexander Keller
NVIDIA

Pascal Gautron
NVIDIA

Jiří Vorba
Weta Digital

Iliyan Georgiev
Autodesk

Martin Šik
Chaos Czech

Eugene d'Eon
NVIDIA

Pascal Grittmann
Saarland University

Petr Vévoda
Charles University Prague

Ivo Kondapaneni
Charles University Prague

ABSTRACT

Jaroslav Křivánek's research aimed at finding that one robust light transport simulation algorithm that would efficient render any scene with arbitrary illumination complexity. He had a clear and unique vision how to reach this ambitious goal. On his way, he created an impressive track of research contributions. In this course, his collaborators will tell the story of Jaroslav's quest for that "one" algorithm and discuss his impact and legacy.

ACM Reference Format:
Alexander Keller, Pascal Gautron, Jiří Vorba, Iliyan Georgiev, Martin Šik, Eugene d'Eon, Pascal Grittmann, Petr Vévoda, and Ivo Kondapaneni. 2020. Advances in Monte Carlo Rendering: The Legacy of Jaroslav Křivánek. In *Proceedings of SIGGRAPH'20 Courses.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 IN MEMORIAM

Jaroslav Křivánek has been an outstanding and highly respected rendering researcher who passed away far ahead of time. Through his numerous contributions to light transport simulation, he profoundly influenced an entire domain of academia and industry.

In this course, we will recap many important contributions of Jaroslav's career, underlining their practicality and piecing them together as subsequent steps to finding the "one" robust light transport simulation algorithm that would efficiently render any given scene. Rarely has a single person had such an impact, and we believe it is worth remembering and continuing his legacy.

2 IRRADIANCE AND RADIANCE CACHING

Irradiance caching has been a popular method to amortize the computation of diffuse inter-reflections over entire scene regions. This idea marked the beginning of Jaroslav's search for a generalized and efficient light transport solution, and he extended the principle to global illumination on glossy surfaces [Křivánek et al. 2005]. While effective in principle, (ir)radiance caching has numerous caveats, such as restricted surface roughness range, interpolation artifacts, and corner oversampling. We will elaborate on solutions towards practical, robust (ir)radiance caching and its applications in production rendering.

SIGGRAPH'20 Courses, July 19–23, 2020, Washington, DC, USA

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

3 SAMPLING PATHS

Monte Carlo rendering methods are all based on stochastically sampling light transport paths that connect emitters and sensors. The key to achieving efficiency is to importance sample the paths' contribution, i.e. to find those paths that bring significant amount of light to the camera. Jaroslav recognized that devising a single robust path sampling technique that can handle all types of scenes is an elusive challenge. Instead, he focused on multiple, simple techniques, each specialized for a different illumination effect, and on efficiently combining these techniques. A recurring theme in Jaroslav's research, this effort has pushed the state of the art in both surface [Georgiev et al. 2012] and volumetric [Křivánek et al. 2014; Georgiev et al. 2013, 2019] rendering. We will review such techniques and will discuss the valuable insights they have provided.

4 ZERO-VARIANCE WALKS

The theory of zero-variance estimators from neutron transport can inform the design of low-variance estimators by making globally-informed (as opposed to purely local) importance sampling decisions at every scattering event in a medium to guide paths towards light sources in a way that balances their final contributions back at the camera. We demonstrate the theory using a novel perfectly zero-variance estimator due to Jaroslav, and also review a practical variance reduction scheme for subsurface scattering [Křivánek and d'Eon 2014].

5 PATH GUIDING

Traditional path sampling techniques are inefficient in scenes with complex geometric occlusion. This can be addressed by designing an estimator inspired by the zero-variance theory, which guide paths towards relevant regions of the scene. The work of Vorba et al. [2014], under Jaroslav's supervision, has resumed the interest in such path guiding methods, showing their practical potential on scenes with complex visibility and as a complement to methods like VCM [Georgiev et al. 2012]. More importantly, this was the first work to point out that path guiding can be viewed as learning uncertainty, and as such an abundant toolbox of machine learning techniques can be explored within the rendering context. We cover path guiding techniques explored by the team around Jaroslav [Vorba and Křivánek 2016; Herholz et al. 2016, 2019], show their connection to zero-variance theory and neutron transport, and discuss the impact of these works in research and industry.

6 DIRECT LIGHTING

Direct and indirect illumination calculations are two important components of any physically-based renderer. While the indirect component has been traditionally considered a more complex problem and has been studied in many research works, Jaroslav acknowledged that improving the efficiency of direct illumination could have a substantial impact on the overall rendering performance, especially with complex visibility and in the presence of large numbers of light sources. In this part, we will cover direct illumination sampling based on online learning of light selection probability distributions. We will show how to formulate the learning process as Bayesian regression to prevent over-fitting and ensure robustness even in the early stages of computation [Vévoda et al. 2018].

7 MULTIPLE IMPORTANCE SAMPLING

Efficiently combining various sampling techniques is vital in modern realistic rendering. For over a decade, the balance and power multiple importance sampling (MIS) heuristics have been universally accepted, and the problem was largely deemed solved by the community. Jaroslav's search for the "one" algorithm led him to challenge these widespread beliefs. We will discuss how the optimal weights can be far better than the balance heuristic [Kondapaneni et al. 2019], and we show that injecting variance information can improve robustness [Grittmann et al. 2019]. We further discuss how the theoretical insights around MIS have been used to make algorithms more lightweight, robust, and efficient [Karlík et al. 2019].

8 MARKOV CHAIN METHODS

Jaroslav saw the Markov chain methods as an alternative way towards the "one" rendering algorithm. These methods generate sequences of correlated samples, which yield faster convergence than independent Monte Carlo sampling. However, that convergence can be irregular and unpredictable, which had been overlooked [Křivánek et al. 2014; Šik and Křivánek 2019b]. We will discuss how to achieve more uniform convergence in Markov chain Monte Carlo [Šik and Křivánek 2016; Gruson et al. 2016; Šik et al. 2016] to improve its viability in practice [Šik and Křivánek 2019a].

9 THE LEGACY

Despite all the excellent contributions, ranging from the amortization of computations, over the efficient construction and combination of light transport paths, to machine learning, there is still unfinished business. Jaroslav Křivánek with his students and collaborators laid out a clear vision of research directions. With this course, we invite you to continue contributing to the search of the "one" robust and efficient light transport simulation algorithm.

REFERENCES

Iliyan Georgiev, Jaroslav Křivánek, Tomáš Davidovič, and Philipp Slusallek. 2012. Light Transport Simulation with Vertex Connection and Merging. *ACM Trans. Graph.* 31, 6, Article 192 (Nov. 2012).

Iliyan Georgiev, Jaroslav Křivánek, Toshiya Hachisuka, Derek Nowrouzezahrai, and Wojciech Jarosz. 2013. Joint Importance Sampling of Low-Order Volumetric Scattering. *ACM Trans. Graph.* 32, 6, Article 164 (Nov. 2013).

Iliyan Georgiev, Zackary Misso, Toshiya Hachisuka, Derek Nowrouzezahrai, Jaroslav Křivánek, and Wojciech Jarosz. 2019. Integral Formulations of Volumetric Transmittance. *ACM Trans. Graph.* 38, 6, Article 154 (Nov. 2019).

Pascal Grittmann, Iliyan Georgiev, Philipp Slusallek, and Jaroslav Křivánek. 2019. Variance-Aware Multiple Importance Sampling. *ACM Trans. Graph.* 38, 6, Article 152 (Nov. 2019).

Adrien Gruson, Mickaël Ribardière, Martin Šik, Jiří Vorba, Rémi Cozot, Kadi Bouatouch, and Jaroslav Křivánek. 2016. A Spatial Target Function for Metropolis Photon Tracing. *ACM Trans. Graph.* 36, 1, Article 4 (Nov. 2016), 4:1–4:13 pages.

Sebastian Herholz, Oskar Elek, Jiří Vorba, Hendrik Lensch, and Jaroslav Křivánek. 2016. Product Importance Sampling for Light Transport Path Guiding. *Computer Graphics Forum (Proceedings of Eurographics Symposium on Rendering)* 35, 4 (2016), 67–77.

Sebastian Herholz, Yangyang Zhao, Oskar Elek, Derek Nowrouzezahrai, Hendrik P. A. Lensch, and Jaroslav Křivánek. 2019. Volume Path Guiding Based on Zero-Variance Random Walk Theory. *ACM Trans. Graph.* 38, 3, Article 25 (June 2019).

Ondřej Karlík, Martin Šik, Petr Vévoda, Tomáš Skřivan, and Jaroslav Křivánek. 2019. MIS Compensation: Optimizing Sampling Techniques in Multiple Importance Sampling. *ACM Trans. Graph.* 38, 6, Article 151 (Nov. 2019).

Ivo Kondapaneni, Petr Vevoda, Pascal Grittmann, Tomáš Skřivan, Philipp Slusallek, and Jaroslav Křivánek. 2019. Optimal Multiple Importance Sampling. *ACM Trans. Graph.* 38, 4, Article 37 (July 2019).

Jaroslav Křivánek and Eugene d'Eon. 2014. A zero-variance-based sampling scheme for Monte Carlo subsurface scattering. In *ACM SIGGRAPH 2014 Talks*. 1–1.

Jaroslav Křivánek, Pascal Gautron, Sumanta Pattanaik, and Kadi Bouatouch. 2005. Radiance caching for efficient global illumination computation. *Visualization and Computer Graphics, IEEE Transactions on* 11, 5 (2005), 550 –561.

Jaroslav Křivánek, Alexander Keller, Iliyan Georgiev, Anton Kaplanyan, Marcos Fajardo, Mark Meyer, Jean-Daniel Nahmias, Ondřej Karlík, and Juan Canada. 2014. Recent Advances in Light Transport Simulation: Some Theory and a Lot of Practice. In *ACM SIGGRAPH 2014 Courses* (Vancouver, Canada) *(SIGGRAPH '14)*. ACM, New York, NY, USA, Article 17, 17:1–17:6 pages.

Jaroslav Křivánek, Iliyan Georgiev, Toshiya Hachisuka, Petr Vévoda, Martin Šik, Derek Nowrouzezahrai, and Wojciech Jarosz. 2014. Unifying Points, Beams, and Paths in Volumetric Light Transport Simulation. *ACM Trans. Graph.* 33, 4, Article 103 (July 2014).

Martin Šik and Jaroslav Křivánek. 2019a. Implementing One-Click Caustics in Corona Renderer. In *Eurographics Symposium on Rendering - DL-only and Industry Track*, Tamy Boubekeur and Pradeep Sen (Eds.). The Eurographics Association, 61–67.

Martin Šik and Jaroslav Křivánek. 2019b. Survey of Markov Chain Monte Carlo Methods in Light Transport Simulation. *IEEE Transactions on Visualization and Computer Graphics* (2019).

Petr Vévoda, Ivo Kondapaneni, and Jaroslav Křivánek. 2018. Bayesian Online Regression for Adaptive Direct Illumination Sampling. *ACM Trans. Graph. (SIGGRAPH 2018)* 37, 4, Article 125 (July 2018).

Jiří Vorba, Ondřej Karlík, Martin Šik, Tobias Ritschel, and Jaroslav Křivánek. 2014. On-line Learning of Parametric Mixture Models for Light Transport Simulation. *ACM Trans. on Graphics (Proc. SIGGRAPH)* 33, 4 (2014), 101:1–101:11.

Jiří Vorba and Jaroslav Křivánek. 2016. Adjoint-Driven Russian Roulette and Splitting in Light Transport Simulation. *ACM Trans. Graph.* 35, 4, Article 42 (July 2016).

Martin Šik and Jaroslav Křivánek. 2016. Improving Global Exploration of MCMC Light Transport Simulation. In *ACM SIGGRAPH 2016 Posters* (Anaheim, California) *(SIGGRAPH '16)*. ACM, New York, NY, USA, Article 50, 50:1–50:2 pages.

Martin Šik, Hisanari Otsu, Toshiya Hachisuka, and Jaroslav Křivánek. 2016. Robust Light Transport Simulation via Metropolised Bidirectional Estimators. *ACM Trans. Graph.* 35, 6, Article 245 (Nov. 2016), 245:1–245:12 pages.