Supplementary Material: Neural Importance Sampling of Many Lights

PEDRO FIGUEIREDO, Texas A&M University, USA
QIHAO HE, Texas A&M University, USA
STEVE BAKO, Aurora Innovation, USA
NIMA KHADEMI KALANTARI, Texas A&M University, USA

CCS Concepts: • **Computing methodologies** → **Ray tracing**; **Neural networks**; • **Mathematics of computing** → *Sequential Monte Carlo methods*.

Additional Key Words and Phrases: many lights rendering, ray tracing, importance sampling, neural networks, next event estimation

ACM Reference Format:
Pedro Figueiredo, Qihao He, Steve Bako, and Nima Khademi Kalantari. 2025. Supplementary Material: Neural Importance Sampling of Many Lights. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '25), August 10–14, 2025, Vancouver, BC, Canada.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3721238.3730754

1 VARL GPU Implementation

We compare our GPU implementation of VARL [Wang et al. 2021] in PBRT v4 [Pharr et al. 2023] versus the original CPU implementation in PBRT v3 [Pharr et al. 2016] provided by the authors. Although numerical comparisons are infeasible due to differences in PBRT versions, we show an equal-sample qualitative comparison on two scenes (Living Room, Staircase2) in Fig. 1. As seen, our implementation generates results with less apparent noise and fewer discretization artifacts while being an order of magnitude faster. Additionally, the runtime speed-up increases with higher light counts: in Living Room (64 lights), the gain is approximately 10x, while in Staircase2 (2.2k lights) the gain increases to approximately 38x.

2 Global Illumination

Following existing many-light methods [Conty Estevez and Kulla 2018; Lin and Yuksel 2020; Wang et al. 2021; Yuksel 2019], we show only direct illumination (and indirect illumination via VPLs) in the main paper. However, here we present comparisons for global illumination in the form of path tracing with multiple importance sampling (BSDF sampling for indirect rays and a many-light method for next-event estimation) using six bounces. Figure 2 shows an equal-time comparison of our method against ATS [Conty Estevez and Kulla 2018], SLCRT [Lin and Yuksel 2020], and VARL [Wang et al. 2021] on two scenes. Note that we do not compare against ReSTIR [Bitterli et al. 2020] since it is designed for 1-bounce direct illumination defined in screen space. As seen, our approach produces significantly better results, both visually and numerically, for Bistro. For Bedroom, the differences between various approaches is generally small, mainly because the noise from indirect illumination dominates the results. Nonetheless, our approach still produces the best results both numerically and visually on this scene.

3 Additional Convergence Analysis

We complement the convergence analysis from the main paper by evaluating with an additional metric: mean absolute percentage error (MAPE) in Figs. 3 (equal-time) and 4 (equal-sample). As seen, our method ranks first in all scenes for both equal-sample and equal-time scenarios.

Acknowledgments

This project was funded in part by the NSF CAREER Award #2238193. We thank Ryusuke Villemin for the valuable discussions, Vincent Serritella for generating the experimental test data, and Magnus Wrenninge and the Aurora simulation team for support. We are grateful to Wang et al. [2021] for releasing the source code of their work. We would like to thank the following artists for sharing their scenes and models that appear in our figures: Mareck (Bathroom), SlykDrako (Bedroom), Amazon Lumberyard [2017] (Bistro), JayArtist (Living Room), Guillermo M. Leal Llaguno (San Miguel), Wig42 (Staircase), NewSee2l035 (Staircase2), and Mike Winkelmann [2019] (Zero Day).

References

Benedikt Bitterli, Chris Wyman, Matt Pharr, Peter Shirley, Aaron Lefohn, and Wojciech Jarosz. 2020. Spatiotemporal reservoir resampling for real-time ray tracing with dynamic direct lighting. *ACM Trans. Graph.* 39, 4, Article 148 (Aug. 2020), 17 pages. doi:10.1145/3386569.3392481

Alejandro Conty Estevez and Christopher Kulla. 2018. Importance Sampling of Many Lights with Adaptive Tree Splitting. *Proc. ACM Comput. Graph. Interact. Tech.* 1, 2, Article 25 (Aug. 2018), 17 pages. doi:10.1145/3233305

Daqi Lin and Cem Yuksel. 2020. Real-time stochastic lightcuts. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 3, 1 (2020), 1–18.

Amazon Lumberyard. 2017. Amazon Lumberyard Bistro, Open Research Content Archive (ORCA). http://developer.nvidia.com/orca/amazon-lumberyard-bistro `http://developer.nvidia.com/orca/amazon-lumberyard-bistro`.

Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2016. *Physically Based Rendering: From Theory to Implementation (3rd ed.)* (3rd ed.). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA. 1266 pages.

Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2023. *Physically Based Rendering: From Theory to Implementation (4th ed.)* (4th ed.). The MIT Press, Cambridge, MA, USA. 1280 pages.

Yu-Chen Wang, Yu-Ting Wu, Tzu-Mao Li, and Yung-Yu Chuang. 2021. Learning to cluster for rendering with many lights. *ACM Trans. Graph.* 40, 6, Article 277 (Dec. 2021), 10 pages. doi:10.1145/3478513.3480561

Mike Winkelmann. 2019. Zero-Day, Open Research Content Archive (ORCA). https://developer.nvidia.com/orca/beeple-zero-day `https://developer.nvidia.com/orca/beeple-zero-day`.

Authors' Contact Information: Pedro Figueiredo, Texas A&M University, USA, pedrofigueiredo@tamu.edu; Qihao He, Texas A&M University, USA, phyqh@tamu.edu; Steve Bako, Aurora Innovation, USA, sbako@aurora.tech; Nima Khademi Kalantari, Texas A&M University, USA, nimak@tamu.edu.

SIGGRAPH Conference Papers '25, Vancouver, BC, Canada

ACM ISBN 979-8-4007-1540-2/2025/08
https://doi.org/10.1145/3721238.3730754

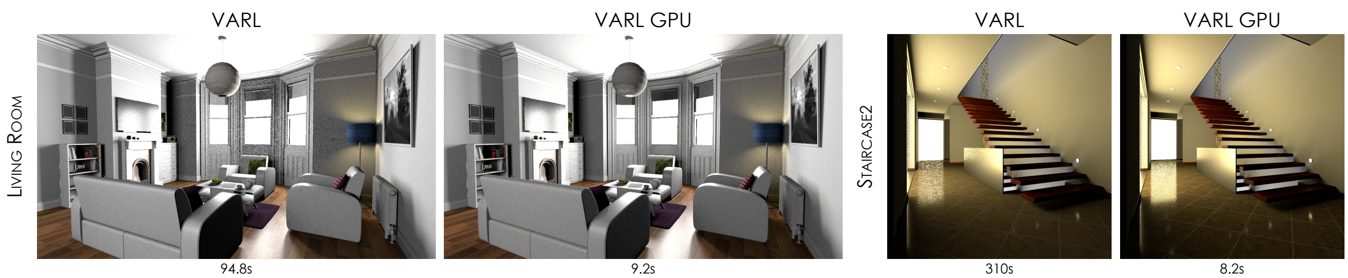

Fig. 1. Equal-sample comparison (256spp) between the original VARL [Wang et al. 2021] CPU implementation in PBRT v3 [Pharr et al. 2016] and our GPU implementation in PBRT v4 [Pharr et al. 2023] (VARL GPU). As seen, VARL GPU shows fewer discretization artifacts while being an order of magnitude faster.

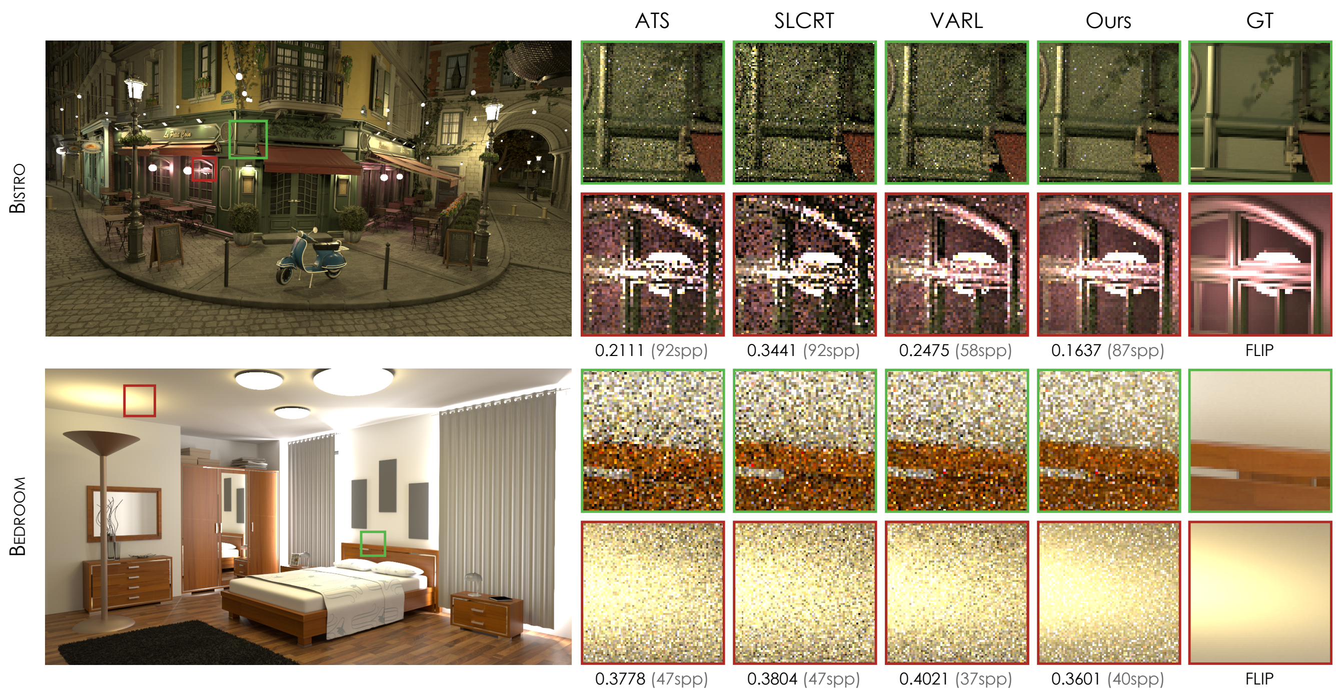

Fig. 2. Equal-time global illumination comparison of ATS [Conty Estevez and Kulla 2018], SLCRT [Lin and Yuksel 2020], VARL [Wang et al. 2021], and our method. The time budget increases with scene complexity and resolution. For Bedroom, the budget is 5 seconds; Bistro 20 seconds.

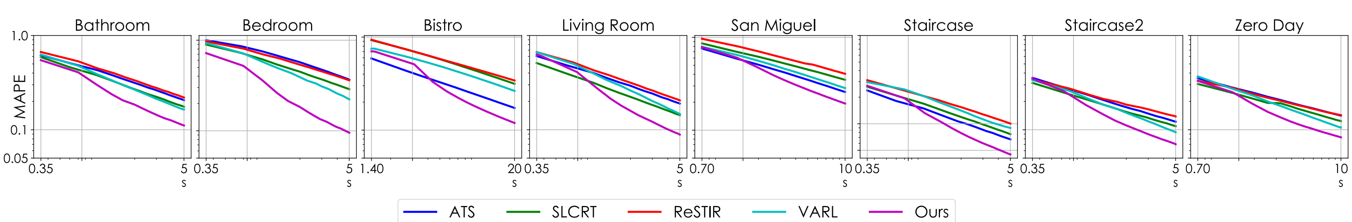

Fig. 3. Equal-time convergence plots of all the approaches on the eight scenes measured in mean absolute percentage error (MAPE). The time budget increases with scene complexity. For all scenes, the second vertical line is approximately where our method stops learning and uses the learned distributions to sample the remaining paths.

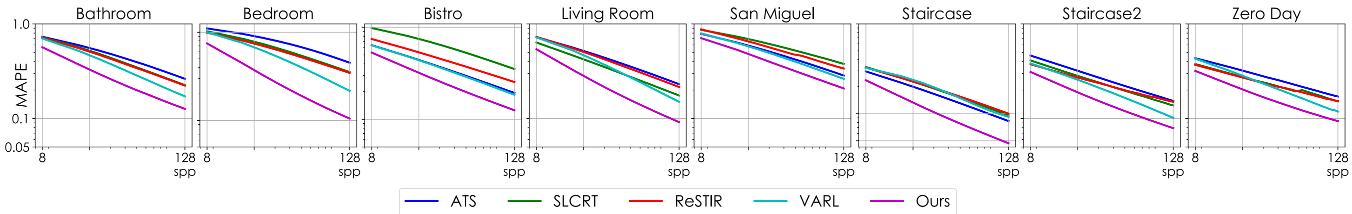

Fig. 4. Equal-sample convergence plots of all the approaches on the eight scenes from 8 to 128 spp measured in mean absolute percentage error (MAPE). The second vertical line is approximately where our method stops learning and uses the learned distributions to sample the remaining paths.

Cem Yuksel. 2019. Stochastic Lightcuts. In *High-Performance Graphics (HPG 2019)* (Strasbourg, France). The Eurographics Association, 27–32. doi:10.2312/hpg.20191192