Reviewer:
Dr. rer. nat. Johannes Schudeiske
Karlsruhe Institute of Technology
Am Fasanengarten 5
Gebäude 50.34, 1. OG,
76131 Karlsruhe, Germany

08.09.2023

Evaluation Report of the Ph.D. Thesis of Petr Vévoda: Advanced Monte Carlo methods in Image Synthesis

1 Content of the Thesis

This thesis advances the state of the art in Monte Carlo rendering in multiple regards. Overall it is concerned with variance reduction, the single most important topic in Monte Carlo integration. The work under review moves this goal closer in three ways.

Chapter 1 introduces a specialised learning-based sampling technique for direct illumination. This is low-overhead and shown to be highly effective. It utilises a light tree similar to LightCuts for scalability and selects the resulting light clusters for each shading point using learned probabilities, based on Bayesian online regression. An actively used production implementation exists and proves the utility of the technique.

Chapter 2 is a theoretical contribution which improves the mixing weights for combined estimators. In particular it introduces the notion of *negative weights* for multiple importance sampling. This leads to new variance bounds for the combined estimators and a proof-of-concept method for limited use cases which is based on the solution of a linear system. The work presented here is not directly applicable in practice since it is limited to a low number of transport path vertices (direct illumination) and sampling techniques. However, it did spur follow-up research and opened a door for a family of research papers improving the weighting schemes.

Chapter 3 presents a precomputed sky model. This model can be used to replace extremely costly atmospheric scattering simulations and thus completely removes any variance from this part of the computation. The writing of this chapter includes a lot of previously published content to make it self-contained and to provide context. The technique considers short wavelength infrared portions of light and thus considers several effects that are normally ignored in computer graphics. The chapter is executed extremely diligently and comes with extensive validation. The spectral sky radiance model is surprisingly limited to solar elevation angles of -4.2°, where astronomical dawn and nautical dawn already start at -18° and -12°, respectively. This is a limitation inherited from earlier works though.

2 Novelty of the proposed Methods

Each chapter is based on a peer reviewed publication in a high-quality journal. Thus the scientific value and novelty of the presented approaches is undisputed.

The learning-based method for direct illumination, as presented in Chapter 1, was a timely contribution, solving a problem that most rendering experts were not aware they had at the time. Often CG scenes would be setup in a way that there are only a few and carefully placed important light sources influencing the scene. The problem with complex direct illumination from very many light sources only arises in production settings, not in academic test scenes. The collaboration that took place here was very fruitful in that regard. The solution based on Bayesian learning is theoretically well grounded and is a result at the intersection of machine learning and Monte Carlo rendering. This was also not widely explored at the time of publication, and the work was a major step towards the acceptance of such methods, showing the potential and the robustness to the research community.

Chapter 2 really broke through a wall: for many years, researchers had accepted that the MIS weights as proposed by Veach are good enough, and even proven to be optimal in some sense. The mere thought that this should be improved upon shows great scientific value, making the paper an early proponent of a family of MIS papers, such as *Variance-Aware Multiple Importance Sampling*, *Correlation-aware multiple importance sampling for bidirectional rendering algorithms* and *Efficiency-aware multiple importance sampling for bidirectional rendering algorithms*.

Finally, the sky model proposed in chapter 3 provides several new features to computer graphics, such as the full spherical definition and the evaluation at different heights. The contribution in this thesis is mostly the extension to SWIR, which makes it suitable for other applications such as solar cell simulation.

3 Importance for the Monte Carlo rendering community

Chapter 1 has direct practical applicability for rendering systems. Above that, it is a step to deepen the understanding of the rendering research community how learning based methods can help create robust adaptive solutions.

Chapter 2 spurred scientific follow up work and showed the research community that there are large gains to be found in an area that has been thought to be well explored in the past.

Chapter 3 provides a finished solution for sky rendering and I am sure will find wide adoption.

All chapters represent ubiquitous problems in Monte Carlo rendering: everybody has direct illumination, everybody combines multiple estimators, many outdoor shots require a sky model, and in particular the considerably improved realism with in-scattering for finite distances.

4 Generality and Applicability in other Areas

The first two chapters are directly applicable to Monte Carlo light transport simulation in many areas such as architectural or product visualisation and visual effects. The results would also carry over to neutron transport simulation. Multiple importance sampling is one of the rare techniques that have made their way from rendering to computational statistics. In that sense any improvement made here will have direct relevance to this field, too.

The third chapter, by virtue of the SWIR extension, widens the field to other natural sciences with applications in solar cell simulation and sensor simulation for automotive assistant systems or autonomous driving.

5 Creative scientific Contribution

The thesis clearly shows strong scientific contribution. While some works are shared with other authors or even appeared in previous theses the individual contributions are marked in extraordinary clarity and are absolutely appropriate in magnitude for a dissertation. The work clearly shows the imprint of the respective supervisors.

Introducing machine learning to direct illumination estimation is a great example of creative scientific work: at a time when almost no graphics professional believed in path guiding for rendering, a new method was devised and showed how this can be robust and practical. The focus on direct illumination is also remarkable, while most researchers were focussing on indirect. This shows strong research intuition and proves intellectual leadership in the field.

Chapter 2 proved out of the box thinking by introducing the counter-intuitive notion of negative MIS weights.

Good science requires two aspects: identifying important problems and ideas how to approach them, and the fine grained work to deliver it. The third chapter is particularly strong in this second aspect, due to extensive data analysis and validation of the model. Also the first chapter demonstrated the extra work necessary to deliver a technique in a practical context.

6 Quality of the Presentation

The thesis is generally well written. The quality of the papers constituting the scientific content is very high, which is reflected in the text here. The content has been augmented by helpful illustrations as well as introductory text to put the works into context. The text is well structured and forms a cohesive story showing the path to variance reduction. This fosters understanding of the individual papers in the context of the greater research area.

The overall introduction is very short. Basics of stochastic theory and Monte Carlo integration are summarised in a single paragraph on page 4. The introduction of necessary equations of radiative transport are deferred until the individual chapters require

them. This makes the document hard to access for non-expert readers.

7 Summary

This thesis shows great strength in implementation and practical, applicable improvements. Especially chapter 2 also provides solid theoretical basic research.

While the author was not the sole primary investigator on most projects included here, he shows the ability to advance the theory in the field, especially by discovering that the balance heuristic variance bounds are loose, and providing an approach to devise optimal MIS weights.

Despite minor weaknesses in writing and volume of the author's contributions I consider the submitted thesis a good piece of academic work which meets the requirements of a Ph.D. thesis. It clearly demonstrates the engagement and ability of the author to advance the field.

For the reasons and under consideration of the reservations named above, I recommend the Ph.D. thesis of Mgr. Petr Vévoda entitled "Advanced Monte Carlo methods in Image Synthesis" for acceptance.

Signature

Dr. rer. nat. Johannes Schudeiske
Karlsruhe, 08.09.2023