2 Syllabus

2.1 14:00 — Opening statement and introduction, Jiri Vorba (15 min)

We introduce path guiding as a family of methods for variance reduction in Monte Carlo based path tracers and state objectives of this course (see above). We provide motivation for adopting path guiding in production, give an overview of some typical scenarios where path guiding methods can greatly reduce numbers of traced paths (e.g. indirect illumination, caustics, selection between many lights) and we provide taxonomy of these methods based on the random decisions that they strive to improve.

2.2 14:15 — Theoretical Background, Jaroslav Krivanek (15 min)

View briefly review basic Monte Carlo rendering algorithms such as path tracing, light tracing, photon mapping or bidirectional path tracing. We define the meaning of "path guiding" as an attempt to perform globally optimal decisions in path construction. We discuss various data-driven approaches for learning the importance and relate it to basic problems in statistical learning such as density estimation and regression. We briefly discuss the relation to the theory of zero-variance random walk and we list the theoretical and practical requirements on good guiding methods such as computational efficiency, low memory overhead, minimal overhead in simple scenes, robustness etc.

2.3 14:30 — Bayesian Inference in Many-light Sampling, Jaroslav Krivanek (15 min)

Production scenes often feature tens, hundreds or even thousands of lights while only varying subsets are visible across the scene. This makes sampling of direct illumination challenging and unpredictable. If path guiding learns optimal sampling decisions only for indirect illumination, the variance coming from choosing a light source connection can still be inhibitive. Jaroslav will discuss a recent approach by Vévoda et al. [2018] to sampling many lights based on Bayesian inference. Bayesian learning is based on taking prior assumptions that are formed by subsequent observations (i.e. sampled paths) that provide evidence about real statistical properties of a given scene including occlusion.

2.4 14:45 — Guiding and Shadow Rays, Alexander Keller (30 min)

Alex will introduce an alternative data-driven approach to many-light sampling which computes and stores light source visibility first, and makes good use of this information during next-event estimation. This results in a guiding technique for shadow rays, and has the potential to greatly reduce noise in path tracing.

2.5 15:15 — "Practical Path Guiding" in Production, Thomas Müller (30 min)

We implemented the "Practical Path Guiding" algorithm [Müller et al., 2017] in Disney's Hyperion renderer [Burley et al., 2018] for use in movie production. Thomas will introduce the algorithm at a high level and describe three extensions based on recently published material Müller et al. [2018] that we developed to further improve the algorithm's effectiveness in a production environment. These extensions are (i) inverse-variance-weighted sample combination to avoid wasted samples, (ii) spatio-directional filtering to increase robustness against high-frequency illumination, and (iii) on-line learning of the BSDF : guiding ratio to improve handling of highly glossy materials. Mitsuba source code containing the extensions is available publicly at https://github.com/tom94/practical-path-guiding.

2.6 15:45 — Break (15min)

2.7 16:00 — Volumetric path guiding, Sebastian Herholz (30 min)

In participating media, path construction is influenced by scattering direction and distance sampling, Russian roulette, and splitting strategies. Sebastian will present a volumetric path construction technique Herholz et al. [2019] where all these sampling decisions are guided by a cached estimate of the adjoint light transport solution. This sampling strategy is based on the theory of zero-variance transport estimators, and it accounts for the spatial and directional variation in volumetric transport. Specifically, paths are constructed incrementally by sampling collision distances proportionally to the product of transmittance and the adjoint transport solution (i.e., in-scattered radiance). Scattering directions are likewise sampled according to the product of the phase function and the incident radiance estimate. Combined with an adaptive Russian roulette and splitting strategy tailored to volumes, variance is greatly reduced as compared to uni-directional methods. Sebastian will also discuss his experience with implementing this method in a production renderer such as Weta Digital's Manuka.

2.8 16:30 — Guiding in path space, Johannes Hanika (30 min)

Johannes will take the ideas presented in the last talk, especially about volume sampling and Russian roulette, and relate them to guiding of full paths in path space Simon et al. [2018]. Guiding new samples along full guide paths instead of marginalised distributions which only guide low dimensional parts at a time transparently includes all aspects. It also allows us to use simpler, uni-modal functions to represent a continuous PDF around guide samples. However both marginalised caches and full paths come with advantages and drawbacks. These properties will be systematically analysed and categorised, leading over seamlessly into the next talk.

2.9 17:00 — Open problems and future work, Jiri Vorba (15min)

We identify the most pressing problems and share them with the research community to enable further exploration of path guiding methods. We also discuss the possibility of combination between some of these methods so that the best of them would form more efficient algorithms.

Some of the aspects we cover are

- the problem of so called global exploration which is related to finding a useful path first before it can be locally explored,
- efficiency driven Russian roulette,
- problems specific to marginalised/cached guiding methods,
- problems related to “path space path guiding”,
- problems specific to direct illumination sampling and light selection.

References

Brent Burley, David Adler, Matt Jen-Yuan Chiang, Hank Driskill, Ralf Habel, Patrick Kelly, Peter Kutz, Yining Karl Li, and Daniel Teece. 2018. The Design and Evolution of Disney's Hyperion Renderer. *ACM Trans. on Graphics* 37, 3, Article 33 (Aug. 2018). https://doi.org/10.1145/3182159

Sebastian Herholz, Yangyang Zhao, Oskar Elek, Derek Nowrouzezahrai, Hendrik Lensch, and Jaroslav Křivánek. 2019. Zero-Variance Based Sampling for Volume Path Guiding. *ACM Trans. on Graphics (conditionally accepted)* (2019).

Thomas Müller, Markus Gross, and Jan Novák. 2017. Practical Path Guiding for Efficient Light-Transport Simulation. *Computer Graphics Forum (Proc. Eurographics Symposium on Rendering)* 36, 4 (June 2017), 91–100.

Thomas Müller, Brian McWilliams, Fabrice Rousselle, Markus Gross, and Jan Novák. 2018. Neural Importance Sampling. *CoRR* abs/1808.03856 (2018). arXiv:1808.03856 http://arxiv.org/abs/1808.03856

Florian Simon, Alisa Jung, Johannes Hanika, and Carsten Dachsbacher. 2018. Selective guided sampling with complete light transport paths. *Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 37, 6 (Dec. 2018). https://doi.org/10.1145/3272127.3275030

Petr Vévoda, Ivo Kondapaneni, and Jaroslav Křivánek. 2018. Bayesian online regression for adaptive direct illumination sampling. *ACM Trans. on Graphics (Proc. SIGGRAPH)* 37, 4 (Aug. 2018), 125:1–125:12.