

Extended Abstract: Path Guiding in Production

Jiří Vorba
Weta Digital Ltd.

Johannes Hanika
Weta Digital Ltd.

Sebastian Herholz
University of Tübingen

Thomas Müller*
NVIDIA

Jaroslav Krivánek
Charles University, Prague
Render Legion

Alexander Keller
NVIDIA



Figure 1: Path guiding allows efficient rendering of notoriously difficult light transport conditions: caustics behind smooth dielectric surfaces. These production shots from *Alita: Battle Angel* show its versatility. We applied guiding on close-up shots of Alita’s eyes featuring SDS caustics as well as on vast underwater scenes with god-rays and caustics at the lake bottom and on the main character. ©2018 Twentieth Century Fox Film Corporation. All rights reserved.

ABSTRACT

Path guiding is a family of adaptive techniques for variance reduction in physically based rendering which includes techniques for sampling both direct and indirect illumination, surfaces and volumes but also for sampling optimal path lengths. Since the adoption of path tracing as a de facto standard in VFX industry several years ago there has been increased interest in producing high-quality images at low sample counts. Path guiding has proven to be useful for this task and therefore has been adopted by Weta Digital. Recently, it has been implemented in the Walt Disney Animation Studios’ Hyperion and Pixar’s Renderman. It has received attention in the research community in the past few years which resulted in several publications introducing both orthogonal and competing guiding methods. In this course, we share our practical experience with path guiding that we gained from production rendering systems at Weta Digital and the Walt Disney Animation Studios as well as from our research in this area. One goal of this course is to sort and classify recent work on various guiding methods. Some are competing, some are orthogonal, and some could be partially

*The presented work was conducted while the author was employed at ETH Zürich and Disney Research.

SIGGRAPH ’19 Courses, July 28 - August 01, 2019, Los Angeles, CA, USA

© 2019 Copyright held by the owner/author(s).

This is the author’s version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Proceedings of SIGGRAPH ’19 Courses*, <https://doi.org/10.1145/3305366.3328091>.

combined to form even better sampling schemes. We identify these relationships and provide suggestions for important avenues of future research. We also cover introduction for people not familiar with path guiding, covering theory and some underlying concepts from machine learning and neutron transport.

KEYWORDS

path guiding, path tracing, production, machine learning

ACM Reference Format:

Jiří Vorba, Johannes Hanika, Sebastian Herholz, Thomas Müller, Jaroslav Krivánek, and Alexander Keller. 2019. Extended Abstract: Path Guiding in Production. In *Proceedings of SIGGRAPH ’19 Courses*. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3305366.3328091>

1 MOTIVATION

Since the advent of path tracing in movie production several years ago, the complexity of rendered scenes has grown substantially. They often include many hard-to-sample light transport effects like strong indirect illumination, occlusions, many lights, a broad range of materials with various scattering profiles, caustics, sub-surface scattering, hair and fur, and volumetric effects. This challenges path tracing based on Monte Carlo (MC) integration which is infamous for its poor convergence rate. In practice, we need to trace many paths and wait tens of hours to obtain a clean image.

To decrease the number of sampled paths, production systems implement various (multiple-)importance sampling schemes, denoising, and even clamping of difficult light transport. Only recently, attention has been paid to *guiding*, data-driven path sampling schemes. Key to their success is that, unlike other schemes, they incorporate knowledge from all previously sampled paths to improve sampling efficiency. There is, however, a substantial amount of trickery involved to make this work in practice.

2 INTRODUCTION

Path tracing is based on sampling of many paths between virtual camera and light sources and collecting transported light along them. Sampling of one path involves multiple decisions like choosing a direction after each scattering event, free path sampling in a volume, absorption or choosing a light source for connection. These decisions have a crucial impact on the efficiency of path tracing. Traditional importance sampling schemes, like for example importance sampling BSDFs, are general and do not adapt to various scenes. In contrast, *path guiding* strives to adapt decisions to a given scene in order to reduce the number of traced paths required for a converged image. We give an overview of typical scenarios where path guiding can greatly reduce variance and we provide some taxonomy of guiding methods.

3 THEORETICAL BACKGROUND

We briefly review MC rendering algorithms and define *path guiding* as global importance sampling of decisions during path sampling. We expose the related issues, namely that the *importance* necessary for the optimal sampling is not available upfront. Subsequently, 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 and briefly discuss the relation to the theory of zero-variance random walk [?]. We list requirements on good guiding methods such as computational efficiency, low memory overhead, minimal overhead in simple scenes, robustness etc., and relate the basic concepts of the existing methods to these requirements.

4 GUIDING AND SHADOW RAYS

When connecting endpoints of segments of paths by shadow rays, we need to test their mutual visibility. However, when visibility information is not included in importance sampling, images may be noisier because paths are not sampled proportionally to their contribution. Realizing that a hitpoint of a photon is the end of a ray that started where light came from allows one to use photons as an informed guess about visibility and thus to notably reduce noise. We discuss data structures and algorithms to efficiently store, learn, and use visibility information in path tracing.

5 GUIDING METHODS IN PRODUCTION

In Disney's Hyperion renderer, we implemented the "Practical Path Guiding" algorithm [?] for use in movie production. We will introduce the algorithm at a high level and describe three extensions based on recently published material [?] that 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.

6 VOLUMETRIC ZERO-VARIANCE BASED PATH GUIDING

Recent work on path guiding in volumes applies strict rules given by zero-variance theory. We demonstrate the importance of guiding all individual sampling decisions (directional, distance, Russian roulette or splitting) when aiming to reduce the variance of the MC estimator and explain how parametric mixture models can be used to represent incident and in-scattered radiance needed to guide these decisions. Furthermore, we discuss the details one needs to consider when implementing path guiding in a production renderer such as Weta Digital's *Manuka*.

7 GUIDING IN PATH SPACE

Most path guiding solutions for sampling indirect illumination [?????] are based on spatial caching of 2D directional distributions. We relate these solutions to guiding full paths in path space [Simon et al. 2018]. However, both marginalized caches and full paths come with advantages and drawbacks. These properties are systematically analysed and categorised, leading seamlessly into open problems and future work.

8 OPEN PROBLEMS AND FUTURE WORK

We believe that there is still an unexplored potential of guiding methods to further reduce number of samples required for rendering clean images. One of the main goals of this course is to identify and share the most pressing problems with the research community and thus to enable further exploration of path guiding. We also discuss the possibility of combining some of these methods so that the best of them would form more efficient algorithms.

ACKNOWLEDGEMENTS

This work was partially supported by Czech Science Foundation grant no 19-07626S.