Machine Learning and Rendering

Alexander Keller
NVIDIA
akeller@nvidia.com

Jaroslav Křivánek
Charles University, Prague
Jaroslav.Krivanek@mff.cuni.cz

Jan Novák
Disney Research, Zürich
jan.novak@disneyresearch.com

Anton Kaplanyan
Oculus Research
kaplanyan@gmail.com

Marco Salvi
NVIDIA
msalvi@nvidia.com

ABSTRACT

Machine learning techniques just recently enabled dramatic improvements in both realtime and offline rendering. In this course, we introduce the basic principles of machine learning and review their relations to rendering. Besides fundamental facts like the mathematical identity of reinforcement learning and the rendering equation, we cover efficient and surprisingly elegant solutions to light transport simulation, participating media, noise removal, and anti-aliasing.

CCS CONCEPTS

• **Computing methodologies → Rendering**; **Reinforcement learning**; *Neural networks*;

KEYWORDS

Rendering, path tracing, anti-aliasing, parametric mixture models, neural networks, reinforcement learning, integral equations.

ACM Reference Format:
Alexander Keller, Jaroslav Křivánek, Jan Novák, Anton Kaplanyan, and Marco Salvi. 2018. Machine Learning and Rendering. In *Proceedings of SIGGRAPH '18 Courses.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3214834.3214841

1 INTRODUCTION

Machine learning can be roughly partitioned into supervised, semi-supervised, and unsupervised learning. While the predominant example for supervised learning are deep neural networks, reinforcement learning stands for reward based learning, and techniques like clustering allow for learning without feedback.

In fact, machine learning techniques have deep mathematical relations to rendering, for example, reinforcement learning may be described by the same integral equation that is known as the rendering equation [2], see Fig. 1. Deep neural networks and parametric mixture models may be used to learn function approximations and thus complement the classical representations used in computer graphics. The capability of generalization provided by autoencoder neural networks may be used for function reconstruction from a sparse set of samples.

SIGGRAPH '18 Courses, August 12-16, 2018, Vancouver, BC, Canada

ACM ISBN 978-1-4503-5809-5/18/08.
https://doi.org/10.1145/3214834.3214841

In this course, we will cover the emerging applications of machine learning in both realtime and offline rendering.

1.1 Path Guiding by learning Mixture Models

Monte Carlo techniques for light transport simulation may be improved by importance sampling during constructing light transport paths. Representing such distributions for importance sampling by a parametric mixture model trained in a progressive manner from a potentially infinite stream of particles enables recovering good sampling distributions in scenes with complex lighting within finite memory. Using these distributions for guiding light transport paths significantly improves the performance of light transport simulation algorithms [4], see Fig. 2.

1.2 Learning Distributions for Deep Scattering

Combining Monte Carlo integration and neural networks allows for efficiently synthesizing images of atmospheric clouds [3]. Instead of simulating all light transport during rendering, the spatial and directional distribution of radiant flux is learned from tens of cloud exemplars. To render a new scene, visible points of the cloud are sampled and a hierarchical 3D descriptor of the cloud geometry with respect to the shading location and the light source is extracted. The descriptor is input to a deep neural network that predicts the radiance function for each shading configuration. A GPU implementation synthesizes images of clouds that are nearly indistinguishable from the reference solution within seconds to minutes, see Fig. 3. The method thus represents a viable solution for applications such as cloud design and, thanks to its temporal stability, for high-quality production of animated content.

1.3 Rendering using a Recurrent Autoencoder

Reconstructing global illumination at interactive rates given extremely low sampling budgets is possible using a recurrent denoising autoencoder neural network [1]. The addition of recurrent connections to the network drastically improves temporal stability for sequences of sparsely sampled input images. The method also has the desirable property of automatically modeling relationships based on auxiliary per-pixel input channels, such as depth and normals. Compared to existing methods that run at comparable speeds, the approach yields significantly higher quality, see Fig. 4.

1.4 Deep Learning Anti-Aliasing

Supersampling is the de-facto standard technique for generating high quality antialiased images for off-line rendering, but its high computational cost is limiting adoption in interactive applications.

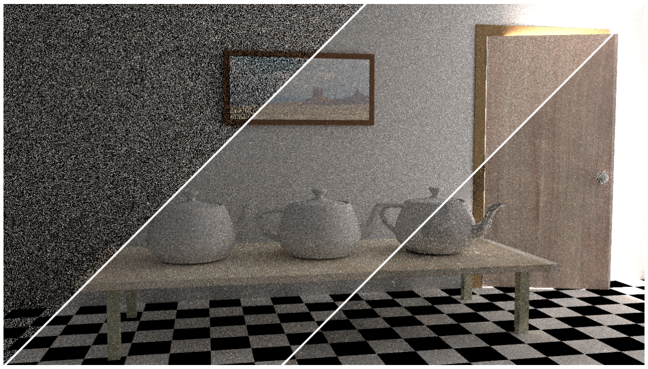

Figure 1: Simple path tracing (left) combined with simple reinforcement learning (middle) outperforms even the much more complicated Metropolis light transport algorithm (right) at the same computational budget.

Figure 2: Online learning of how to guide light transport paths dramatically increases the efficiency of light transport simulation.

To overcome this issue we amortize the cost of supersampling by training a convolutional neural network to integrate information from current and past frames in a temporally coherent fashion. For sequences rendered with 1 sample per pixel, our network generates an image quality nearly equivalent to renderings at 16 samples per pixel at a fraction of the original cost (see Fig. 5).

(a) PT reference (b) RPNN (ours) (c) PT (d) RPNN

Figure 3: Learning distributions allows for the much more efficient simulation of complex phenomena.

Figure 4: Using recurrent autoencoders enables close to realtime light transport simulation (left: input at one sample per pixel, middle: autoencoder reconstruction, right: reference).

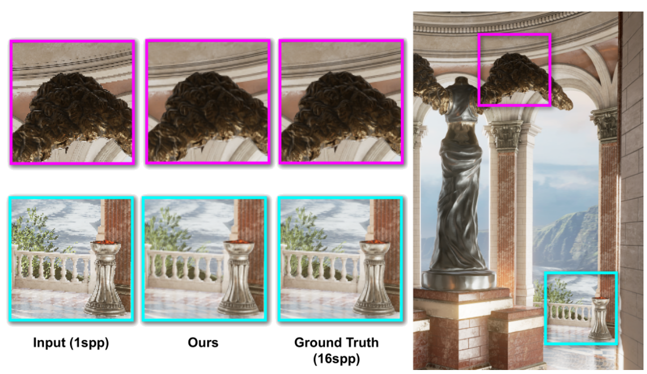

Figure 5: Realtime high quality temporally stable anti-aliasing can be learned.

REFERENCES

[1] C. Chaitanya, A. Kaplanyan, C. Schied, M. Salvi, A. Lefohn, D. Nowrouzezahrai, and T. Aila. 2017. Interactive Reconstruction of Monte Carlo Image Sequences Using a Recurrent Denoising Autoencoder. *ACM Trans. Graph.* 36, 4, Article 98 (July 2017), 12 pages. https://doi.org/10.1145/3072959.3073601

[2] K. Dahm and A. Keller. 2017. Learning Light Transport the Reinforced Way, to appear in Monte Carlo and Quasi-Monte Carlo Methods 2016. *CoRR* abs/1701.07403 (2017). http://arxiv.org/abs/1701.07403

[3] S. Kallweit, T. Müller, B. McWilliams, M. Gross, and J. Novák. 2017. Deep Scattering: Rendering Atmospheric Clouds with Radiance-Predicting Neural Networks. *ACM Trans. Graph. (Proc. of Siggraph Asia)* 36, 6, Article 231 (Nov. 2017), 11 pages. https://doi.org/10.1145/3130800.3130880

[4] J. Vorba, O. Karlík, M. Šik, T. Ritschel, and J. Křivánek. 2014. On-line Learning of Parametric Mixture Models for Light Transport Simulation. *ACM Trans. Graph.* 33, 4, Article 101 (July 2014), 11 pages. https://doi.org/10.1145/2601097.2601203