DEEP LEARNING FOR LIGHTING SIMULATIONS

Peter Quinn, Derek Nowrouzezahrai
HT06 - ECSE 498/499 - W2019
McGill

Abstract

The objective of this project was to improve ray tracing by using machine learning to learn and produce new light paths that contributed non-zero radiance to the final image. Using paths produced by the trained model, we successfully obtain more non-zero radiance paths when compared to random sampling.

Background

Rendering a photorealistic image uses Monte Carlo (MC) integration to solve the rendering equation [1]. MC integration has the downside of having high variance and thus slow convergence. We use a RealNVP [2] architecture to learn and subsequently sample from the distribution that describes paths with 2 bounces in the scene.

Methodology

- The RealNVP architecture:
  - FC layers (40 neurons), BN, ReLU
  - Total Parameters = 92 400
  - Training was done for 7 epochs
- Implemented using PyTorch
- Used PBRT-V3 renderer

Results

Areas with significant 2-bounce indirect illumination improve by using samples from the learned distribution. This corresponds to lower variance and should lead to faster convergence of the image. Bright pixels ('fireflies') due to sampling low probability paths in learned path images are visible.

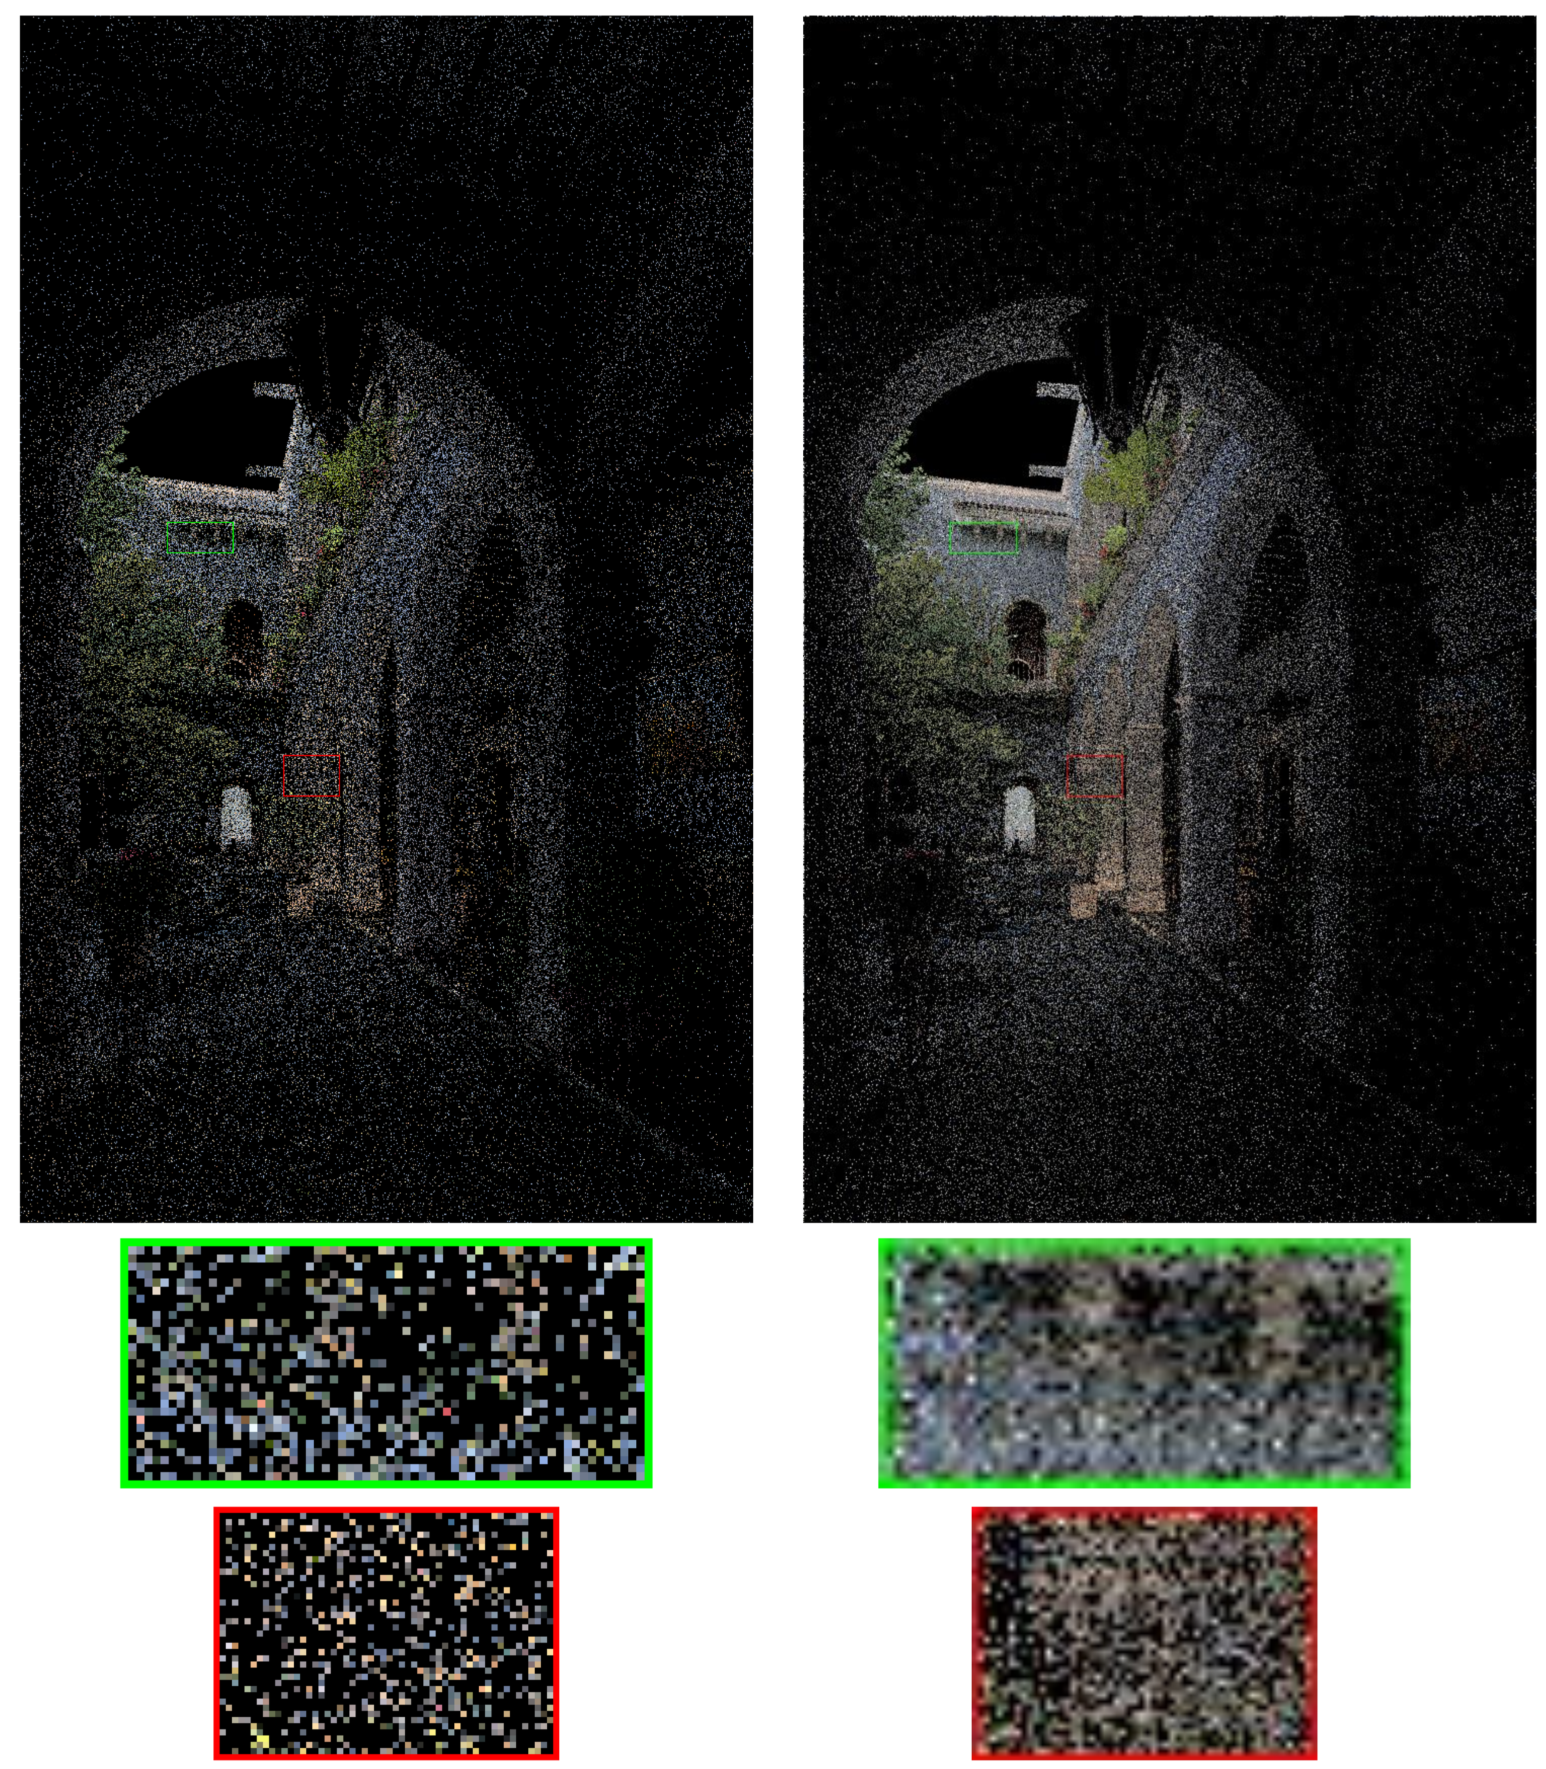

Fig. 1: 2-bounce indirect lighting in 'San Miguel' scene using randomly generated paths (left) and paths generated using machine learning model (right). Rendered with 8 SPP

Scene	Sampling Scheme	Camera Rays	Render Time	Paths Found	Non-zero Paths	%
Teapot	Random	1 081 600	11.7s	126 022	35 632	28.27%
Teapot	Learned	1 081 600	31.7s	224 922	117 329	52.16%
San Miguel	Random	6 904 064	331.9s	6 027 277	138 739	2.30%
San Miguel	Learned	6 904 064	529.1s	5 415 854	319 964	5.91%

Table 1: Comparison of random sampling vs. learned sampling

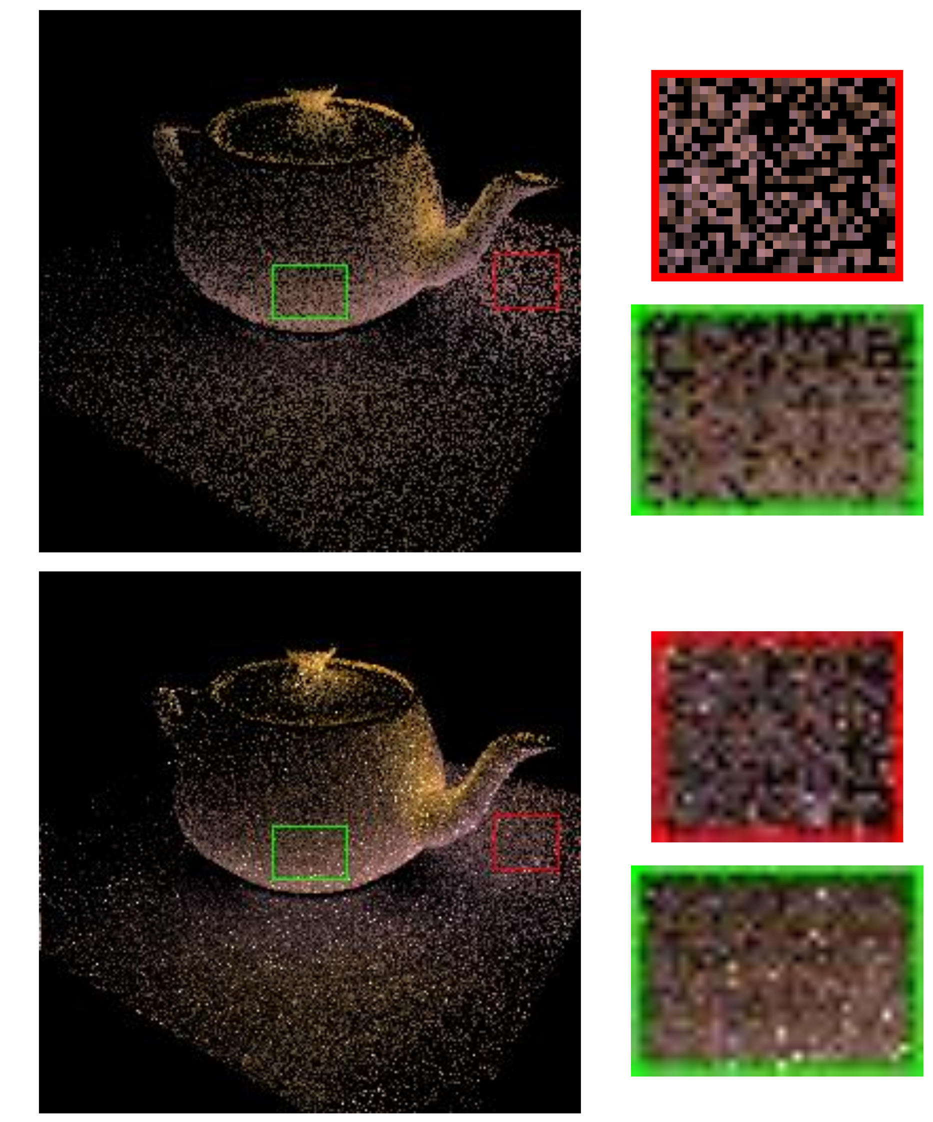

Fig. 2: 2-bounce indirect lighting in 'Teapot with area light' scene using randomly generated paths (top) and paths generated using machine learning model (bottom). Rendered with 16 SPP

Conclusion

The results of this work show that we can see some reduction in variance in the rendered image and more non zero radiance paths being traced when using learned paths.

Future Work

- Optimize neural network parameters
  - Number of layers, neurons, epochs
- Allow for paths of different/longer paths
- Use learned paths for global illumination
- Parametrize on the pixel

References

[1] J. T. Kajiya, "The rendering equation," in *ACM Siggraph Computer Graphics*, 1986, vol. 20, no. 4, pp. 143-150: ACM.
[2] L. Dinh, J. Sohl-Dickstein, and S. J. a. p. a. Bengio, "Density estimation using Real NVP," 2016.

Acknowledgements

Facebook AI for creating and distributing PyTorch.
The creators of PBRT-v3.