DTU Compute
Department of Applied Mathematics and Computer Science

Investigating Machine Learning for Monte-Carlo noise removal in rendered images

Andrea Luongo, Alessandro Dal Corso

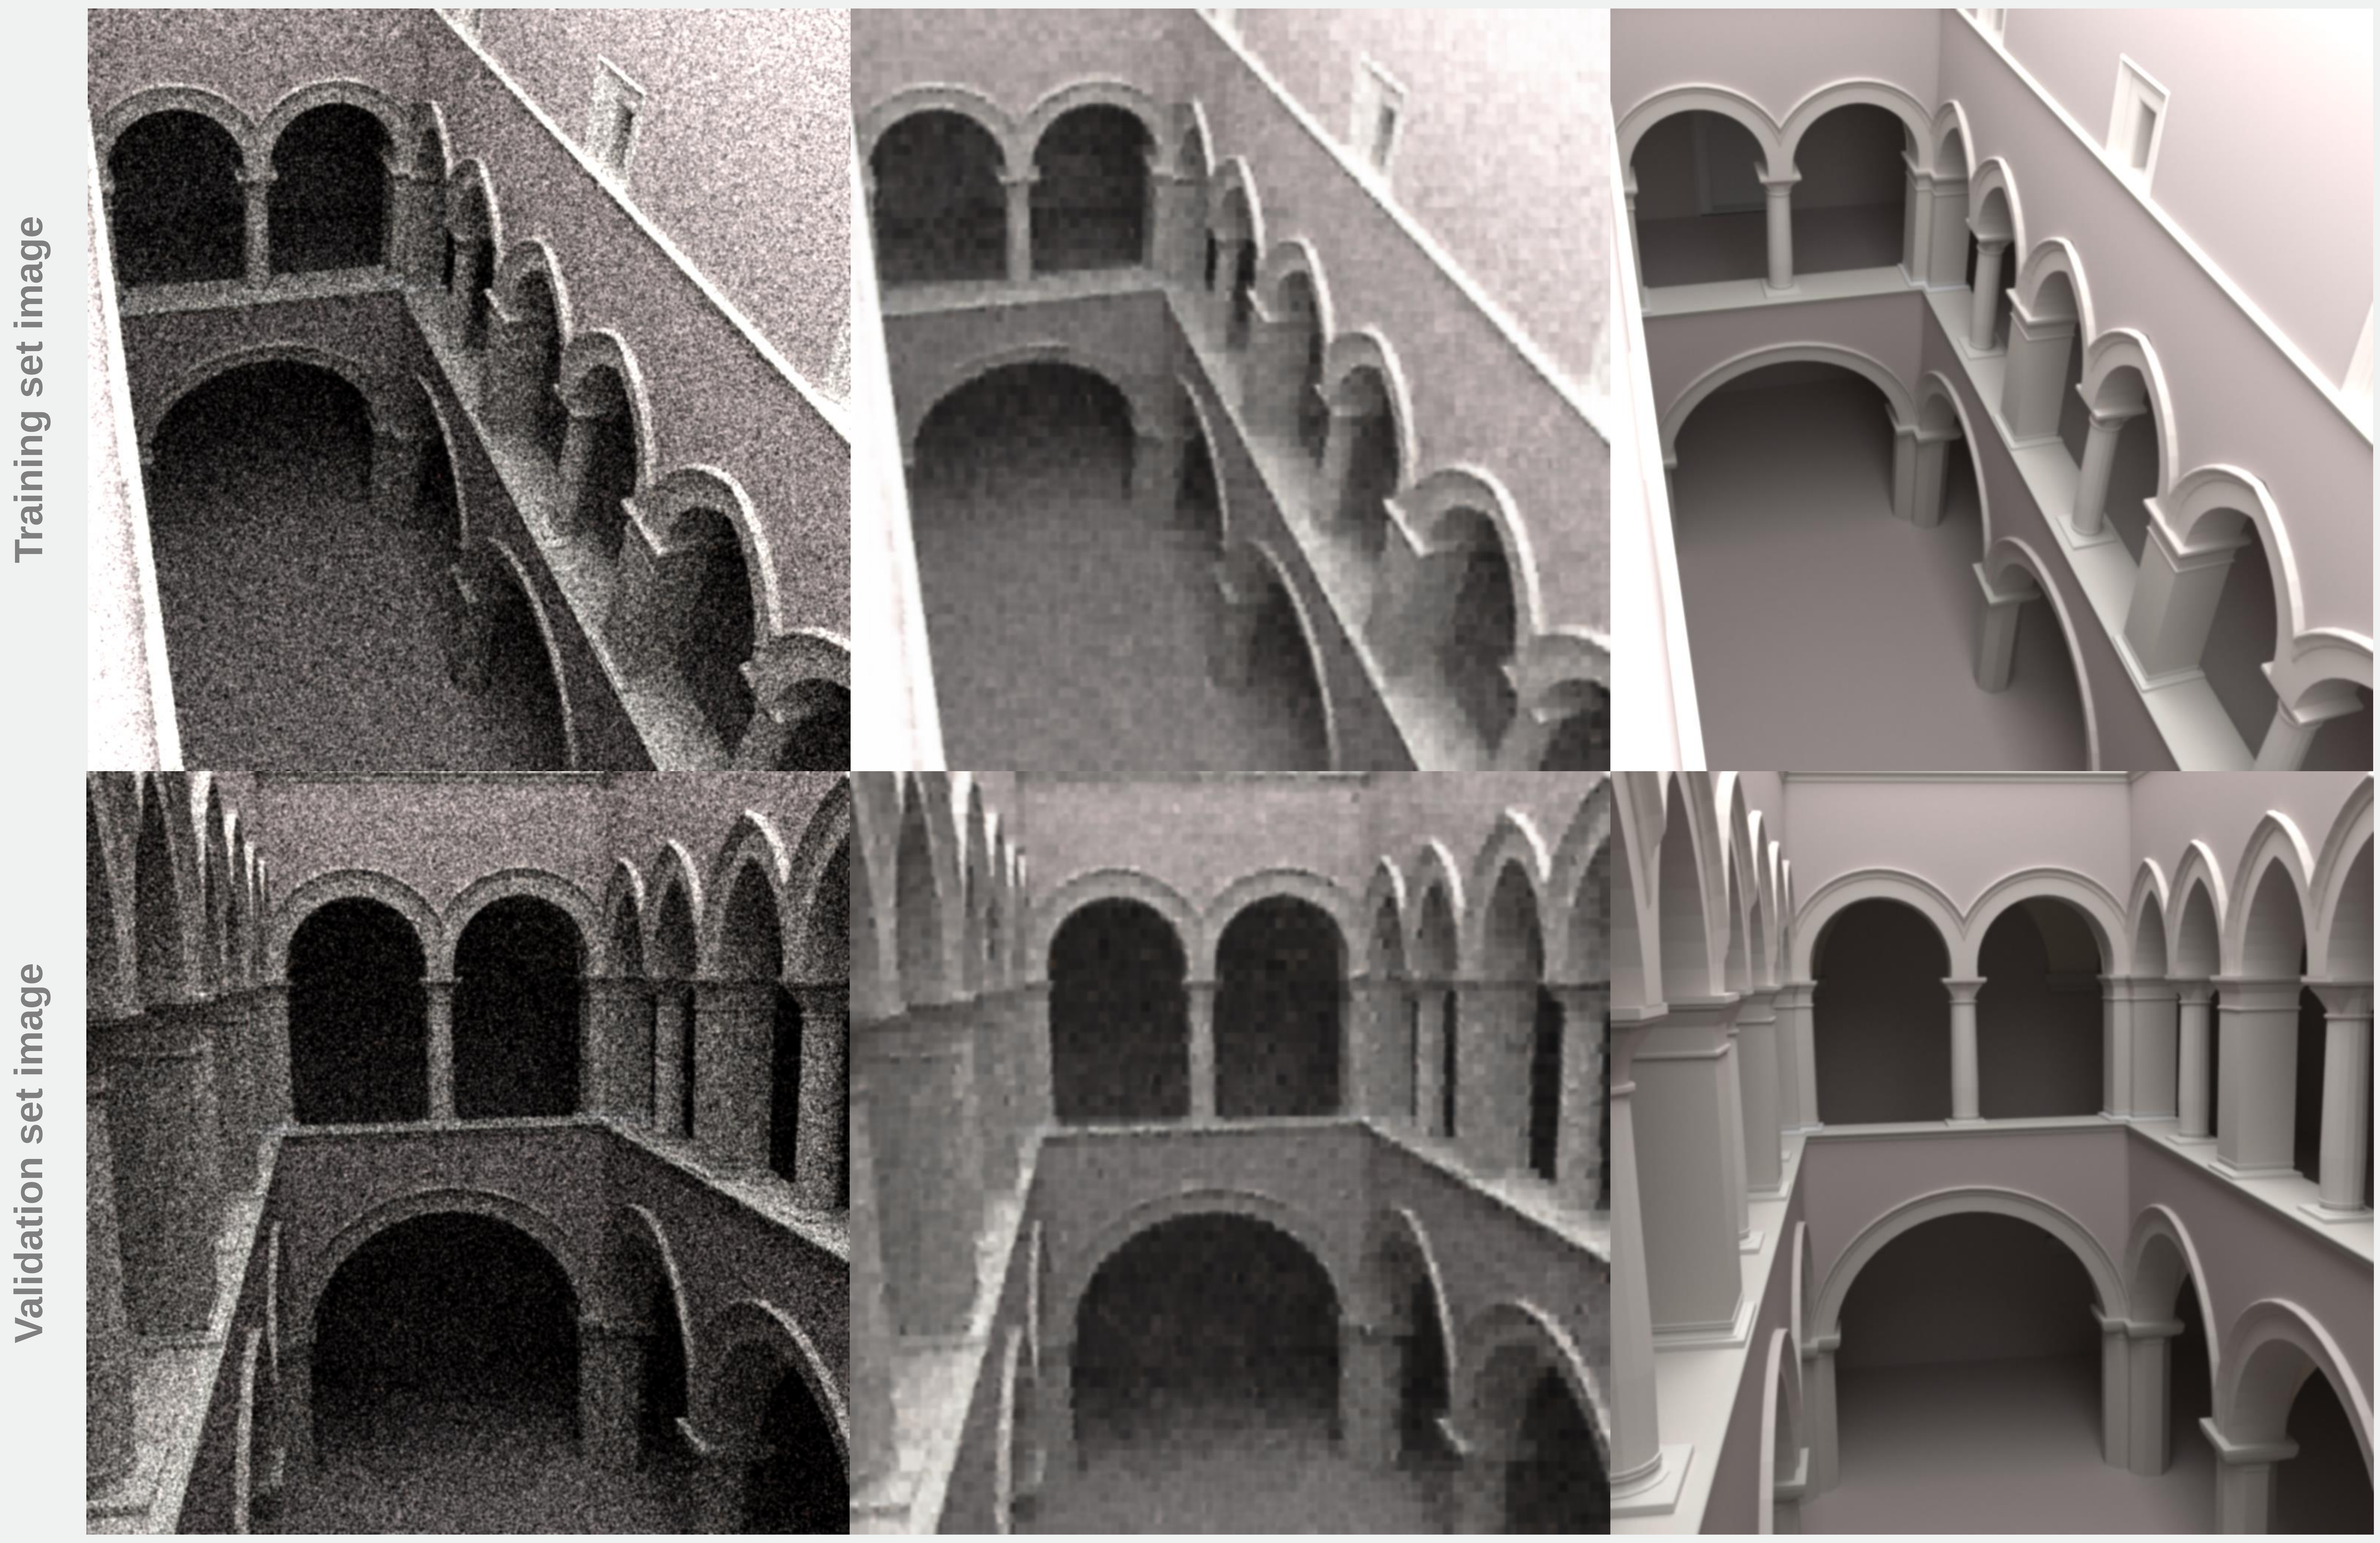

Noisy (5 samples) | Filtered | Reference (10000 samples)

Introduction

Recent advances in ray tracing performance have seen a resurgence of path tracing techniques for production rendering, such as architectural rendering or animated movies. Path tracing techniques generate a noisy unbiased image, that can be improved by tracing an elevated number of paths. However, when we increase the number of samples, path tracing becomes progressively less efficient at removing noise. In production rendering applications, noise in the final image is unacceptable, so advanced denoising techniques are applied onto the final image [1].
In our experiments, we tried to apply a machine learning approach to denoising Monte Carlo images, in the same spirit as in [2]. We employ clean and noisy Monte Carlo images to train a neural network to denoise images not part of the training set.

Method

We based our neural network implementation on the work by Vincent et al. [3,4] on denoising autoencoders. We performed supervised learning on a dataset containing patches from rendered images, with different levels of Monte Carlo noise. Our network consists of two parts: first, we pre-train a series of denoising autoencoders to initialize the parameters of our network. Then, we add an extra MLP layer and to update the parameters by performing fine tuning.

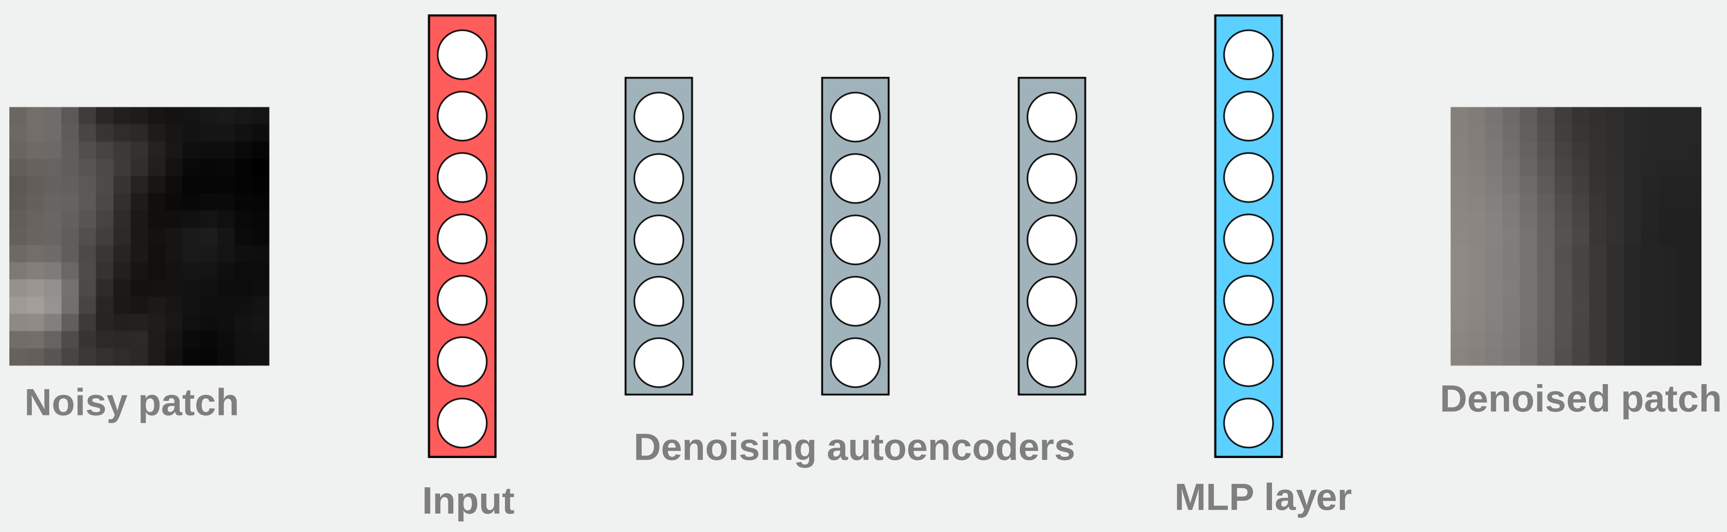

Results

Results are shown in the figure above. We implemented our neural network using the Theano framework running on a NVIDIA GTX 780 Ti GPU. The training dataset were patches of 16x16 pixels from three rendered 1024x1024 views of the Sponza scene. Renderings were generated using a diffuse path tracer implemented on the NVIDIA OptiX ray tracing engine. All surfaces are diffuse with $\rho_d = (1{,}1{,}1)$ reflectance. A constant environment sky light serves as illumination. Validation was performed on another view of the same scene.

Conclusions and Future work

We believe that denoising using machine learning is a viable path for production rendering to reduce rendering times. As future steps we would like to expand our framework to reproduce the work by [2], that originally included reconstruction from scenes with different materials and illumination conditions. Moreover, we would like to investigate an approach using deep convolutional neural networks (CNNs), as they are usually more suitable to an image processing environment.

References

[1] M. Zwicker, W. Jarosz, J. Lehtinen, B. Moon, R. Ramamoorthi, F. Rousselle, P. Sen, C. Soler, S. Yoon. Recent Advances in Adaptive Sampling and Reconstruction for Monte Carlo Rendering. *Computer Graphics Forum (Proceedings of Eurographics),* May 2015.
[2] N. K. Kalantari, S. Bako, and P. Sen. A machine learning approach for filtering Monte Carlo noise. *ACM Trans. Graph.* 34, 2015.
[3] P. Vincent, H. Larochelle, Y. Bengio, P.A. Manzagol. Extracting and Composing Robust Features with Denoising Autoencoders*, ICML 2008*.
[4] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, P.A. Manzagol. Stacked Denoising Autoencoders: Learning Useful Representations in a Deep Network with a Local Denoising Criterion. ICML 2010.

DTU Compute – Department of Applied Mathematics and Computer Science