Neural Denoising for Spectral Monte Carlo Rendering

R. Rouphael[1,2], M. Noizet[1], S. Prévost[1], H. Deleau[1], L-A. Steffenel [1] and L.Lucas†

[1]Université de Reims Champagne-Ardenne, LICIIS, France
[2]United Visual Researchers, France

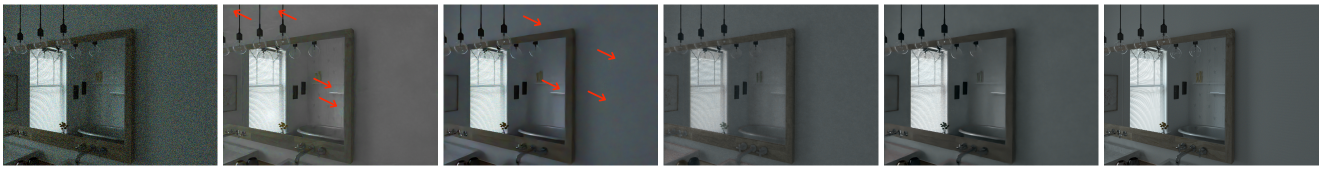

(a) Noisy Input (b) OptiX denoiser (c) Intel denoiser (d) Spectral denoiser (e) Final composition (f) Reference image

Figure 1: *Left to right : (a) Noisy 1024 SPP input generated using the Omen spectral engine [Uni22]; (b) NVIDIA's OptiX denoiser [NVi17]; (c) Intel's Open Image Denoise [Int21]; (d) Our spectral denoiser; (e) Final composition; (f) Reference generated by Omen with 262,144 SPP.*

Abstract

Spectral Monte Carlo (MC) rendering is still to be largely adopted partially due to the specific noise, called color noise, induced by wavelength-dependent phenomenons. Motivated by the recent advances in Monte Carlo noise reduction using Deep Learning, we propose to apply the same approach to color noise. Our implementation and training managed to reconstruct a noise-free output while conserving high-frequency details despite a loss of contrast. To address this issue, we designed a three-step pipeline using the contribution of a secondary denoiser to obtain high-quality results.

CCS Concepts

• ***Computing methodologies*** $\rightarrow$ ***Ray tracing; Neural networks;*** *Image processing;*

1. Introduction

Applied to industry, spectral Monte Carlo (MC) light transport simulation allows addressing predictive rendering purposes. If this physically-based issue is essential to any decision making, it is not without computational cost challenges related to convergence properties that require long rendering times for noise-free images. Recently, many approaches [HY21] have been proposed to overcome these limitations. Their core idea is to first render a noisy image with a few samples per pixel (SPP), and then use denoising algorithms to reconstruct a perceptually noise-free image. Although attractive, this principle raises numerous questions and can produce undesirable post-processing chromatic effects (see figures 1.c, 1.d), in the particular case of spectral images (e.g. stationary or slowly changing tonal noise). This is what we have observed when using off-the-shelf solutions such as NVidia [NVi17] and Intel [Int21].

To address the issue of color noise, we propose to apply these approaches in the form of a denoising neural network trained on spectral rendered noisy images. We managed to obtain noise-free outputs with high-frequency details. Although this network is our main contribution, we also designed a three-steps pipeline using a secondary MC denoiser to separately reconstruct luminance. This allows us to remove luminance artifacts in the original output by merging reconstructed chrominance and luminance and to obtain higher-quality results.

2. Related works

A comprehensive review of Monte Carlo denoising is available in [ZJL*15]. In our case, we only consider neural *a posteriori* methods, which have proven to be effective in reducing Monte Carlo noise. According to this viewpoint, RDAE [CKS*17] and KPCN [BVM*17] are two of the major references in this field. The former is designed over a recurrent auto-encoder to interactively

† This research was supported by ANR-PRCE (ANR-21-CE33-0010). The authors would like to thank the Centre Image at URCA for their material contributions in high-performance computing resources. [1]

process input image of low sample rate. The latter proposes a network able to predict pixel-centric filtering kernels to process input images with higher sample rate.
In order to handle spectral noise, which is still largely overlooked, we propose the adoption of a solution taking into account the spectral nature of the noise in a neural network architecture derived from [CKS*17] and trained on spectral rendered images.

3. System overview

Our three-steps pipeline takes as input a noisy spectral image in addition to auxiliary features such as view-space albedo, normals and depth. The full input is given to our spectral denoiser which aims at recovering color channels whereas only the noisy image is given to the MC denoiser to recover the luminance channel. Both outputs are then encoded in YUV color space and combined to obtain denoised chrominance and luminance.
Architecture. Our step one autoencoder is mainly based on Chaitanya et al's design [CKS*17] of direct prediction, but we dropped the recurrent connections. Moreover, we changed the loss function by the one proposed in [ZGFK17], which is a weighted combination of L1 Spatial Loss and MultiScale-SSIM as it has shown great results at balancing color, luminance and contrast conservation.
Training. Our training and testing dataset were generated with the spectral path tracer Omen [Uni22], which uses Heros-Wavelength [WND*14] to enable 4 wavelengths for each path. This dataset is composed of 700 pairs of reference and noisy input (210 for testing images, 490 for training images) at a resolution of 1024×768 pixels. 47 different points of view were taken from each of 4 scenes exhibiting various light phenomena (diffraction, reflections etc.). Some of these views have been duplicated using different sampling seeds to improve learning. A reference image is computed with 262,144 SPP, whereas a noisy image is computed with a range between 512 to 2048 SPP in order to include a variety of color noise levels and pattern. The denoiser in itself was implemented with Tensorflow and Keras and trained for 500 epochs using the Adam optimizer. Inference was performed on a NVIDIA GTX 1080 in 0.016 seconds on average for a resolution of 720p.

4. Preliminary results

Contrary to the OptiX and Intel denoisers (see figure 1), our SRd is able to reconstruct color noise-free outputs but visual artifacts remain (loss of contrast, change of material aspect, etc.). The nature of these artifacts pointed toward a luminance reconstruction issue. In order to validate our assumption, we performed a two-steps experiment. On one hand, merging the color channel of our denoised output with the ground truth luminance channel (GT.L), thus obtaining a denoised output free from artifacts. On the other hand, merging the ground truth color channel (GT.color) with the luminance channel of our denoised output, thus introducing the same visual artifacts into the reference. We decided to reproduce the same experiment with the off-the-shelf MC denoiser from Intel over Optix as it gave us a lower MSE and a higher SSIM (Table 1).
When comparing (Table 1) denoised luminance with ground truth colors, we obtain higher MSE with our denoiser than with the one from Intel. However, our denoiser performs far better when looking at denoised colors, with an MSE divided in half. This finding led us to step *iii* of our pipeline which combine the denoised colors of our SRd with the denoised luminance of the Intel denoiser. This approach gave us the best results and allow us to obtain results very close to ground-truth (despite some loss in the high frequencies due to Intel's over-blurring) as shown in the last line of table 1.

		MSE	PSNR	SSIM
a	SRd	1029.242	22.776	0.798
	OptiX denoiser	10133.635	12.844	0.461
	Intel Denoiser	921.562	23.256	0.827
b	YUV: GT.color + SRd.L	1002.680	22.890	0.799
	YUV: GT.color + Intel.L	867.81	23.518	0.828
	YUV: GT.L + SRd.color	30.204	38.101	0.999
	YUV: GT.L + Intel.color	62.454	34.946	0.999
c	**SRd.color + Intel.L**	900.713	23.356	0.827

Table 1: *Error statistics relative to GT: (a) Denoised output; (b) Channels merge between GT and denoised output; (c) Channels merge between our denoiser and Intel's.*

5. Conclusion and perspectives

Despite showing promising results in color reconstruction, our SRd still suffers from luminance artifacts and thus necessitates the use of a second denoiser. In our future work, we will focus on improving our neural architecture to remove this dependency and propose an all-in-one real-time spectral denoiser.

References

[BVM*17] BAKO S., VOGELS T., MCWILLIAMS B., MEYER M., NOVÁK J., HARVILL A., SEN P., DEROSE T., ROUSSELLE F.: Kernel-predicting convolutional networks for denoising Monte Carlo renderings. *ACM Transactions on Graphics 36*, 4 (July 2017), 1–14. 1

[CKS*17] CHAITANYA C. R. A., KAPLANYAN A. S., SCHIED C., SALVI M., LEFOHN A., NOWROUZEZAHRAI D., AILA T.: Interactive reconstruction of Monte Carlo image sequences using a recurrent denoising autoencoder. *ACM Transactions on Graphics 36*, 4 (July 2017), 98:1–98:12. 1, 2

[HY21] HUO Y., YOON S.-E.: A survey on deep learning-based Monte Carlo denoising. *Computational Visual Media 7* (Mar. 2021). 1

[Int21] INTEL: Intel open image denoise. https://www.openimagedenoise.org, 2021. 1

[NVi17] NVIDIA: Optix denoiser. https://developer.nvidia.com/rtx/ray-tracing/optix, 2017. 1

[Uni22] UNITED VISUAL RESEARCHERS (UVR): Omen rendering engine. https://www.united-vr.com/, 2022. 1, 2

[WND*14] WILKIE A., NAWAZ S., DROSKE M., WEIDLICH A., HANIKA J.: Hero Wavelength Spectral Sampling. *Computer Graphics Forum 33*, 4 (2014), 123–131. 2

[ZGFK17] ZHAO H., GALLO O., FROSIO I., KAUTZ J.: Loss Functions for Image Restoration With Neural Networks. *IEEE Transactions on Computational Imaging 3*, 1 (Mar. 2017), 47–57. 2

[ZJL*15] ZWICKER M., JAROSZ W., LEHTINEN J., MOON B., RAMAMOORTHI R., ROUSSELLE F., SEN P., SOLER C., YOON S.-E.: Recent Advances in Adaptive Sampling and Reconstruction for Monte Carlo Rendering. *Computer Graphics Forum 34*, 2 (May 2015), 667–681. 1