



HAL
open science

Neural Denoising for Spectral Monte Carlo Rendering

Robin Rouphael, Mathieu Noizet, Stéphanie Prévost, Hervé Deleau,
Luiz-Angelo Steffanel, Laurent Lucas

► **To cite this version:**

Robin Rouphael, Mathieu Noizet, Stéphanie Prévost, Hervé Deleau, Luiz-Angelo Steffanel, et al..
Neural Denoising for Spectral Monte Carlo Rendering. Eurographics 2022, Apr 2022, Reims, France.
, Eurographics 2022 - Posters, 2022, 10.2312/egp.20221011 . hal-04042962

HAL Id: hal-04042962

<https://hal.univ-reims.fr/hal-04042962>

Submitted on 23 Mar 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.




Neural Denoising for Spectral Monte Carlo Rendering

R. ROUPHAEL^{1,2}, M. Noizet¹, S. PRÉVOST¹, H. DELEAU¹, L.A. STEFFENEL¹ and L. LUCAS¹

1. Université de Reims Champagne-Ardenne, LICIIS, LRC DIGIT, France
2. United Visual Researchers, France

PROBLEM STATEMENT

To this day, spectral Monte Carlo (MC) rendering is still to be largely adopted partially due to the specific noise, called **color noise**, induced by **wavelength-dependent phenomena**.



Hundred of thousands *Samples Per Pixels* (SPP) are needed to produce noise-free renderings, thus establishing the need for fast and **efficient post-processing denoising methods**.


RELATED WORK

In recent years, neural networks such as **denoising auto-encoders** have proven themselves to be very effective at reconstructing high quality images from noisy MC renderings [1].

Among others, RDAE [2] and KPCN [3] paved the way for robust off-the-shelf denoisers such as NVIDIA's OptiX [4] and Intel's Open Image Denoise [5].



But these denoisers are ineffective when confronted with spectral color noise and produce **chromatic artifacts**.



OVERVIEW

We propose to train a **denoising neural network on spectral rendered noisy images** in order to remove color noise at acceptable sample rates.

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.

SYSTEM OVERVIEW

Three-step pipeline

Our three-step pipeline (see figure 1) 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 (SRd) which aims at recovering color channels whereas only the noisy image is given to the secondary MC denoiser to recover the luminance channel. Both outputs are then encoded in YUV color space and combined to obtain denoised chrominance and luminance.

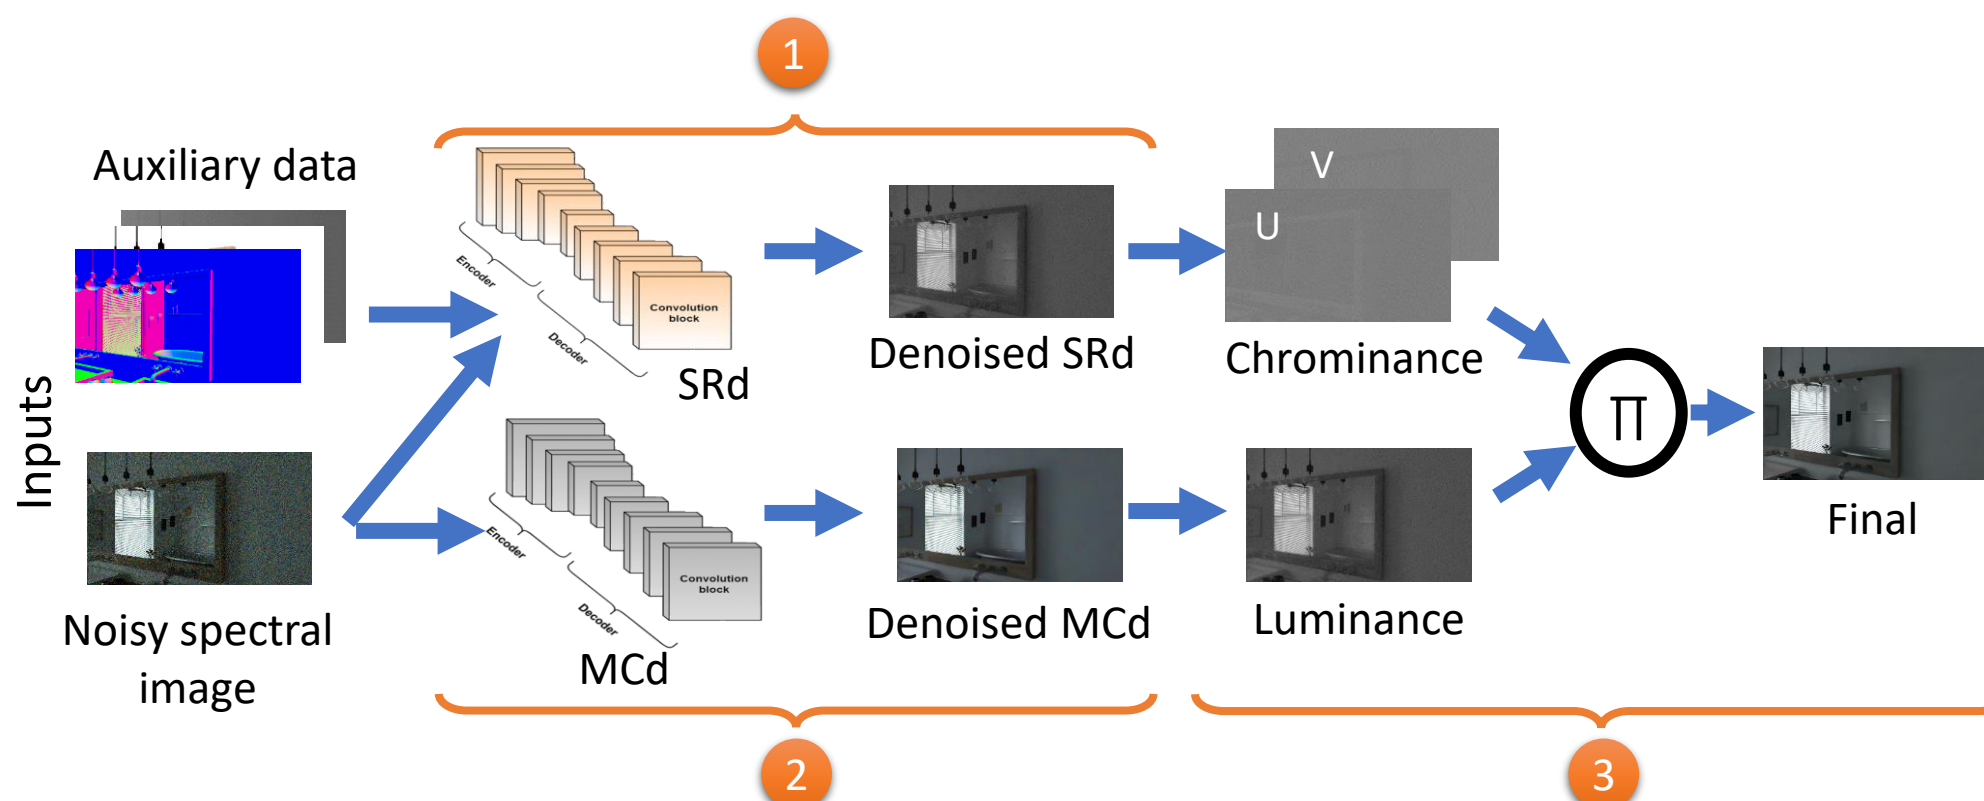


Figure 1: Three-step pipeline : (1) Spectral denoising; (2) Monte Carlo denoising; (3) Chrominance and luminance compositing.

Spectral denoiser (SRd) design (see figure 2)

Architecture: Our auto-encoder is mainly based on Chaitanya et al's design [2] using direct prediction but without the recurrent connections.

Loss function: We used the weighted combination of spatial loss and MultiScale-SSIM proposed by [7] :

$$L^{Mix} = 0.8 \times L^{MS-SSIM} + 0.2 \times L^{l1}$$

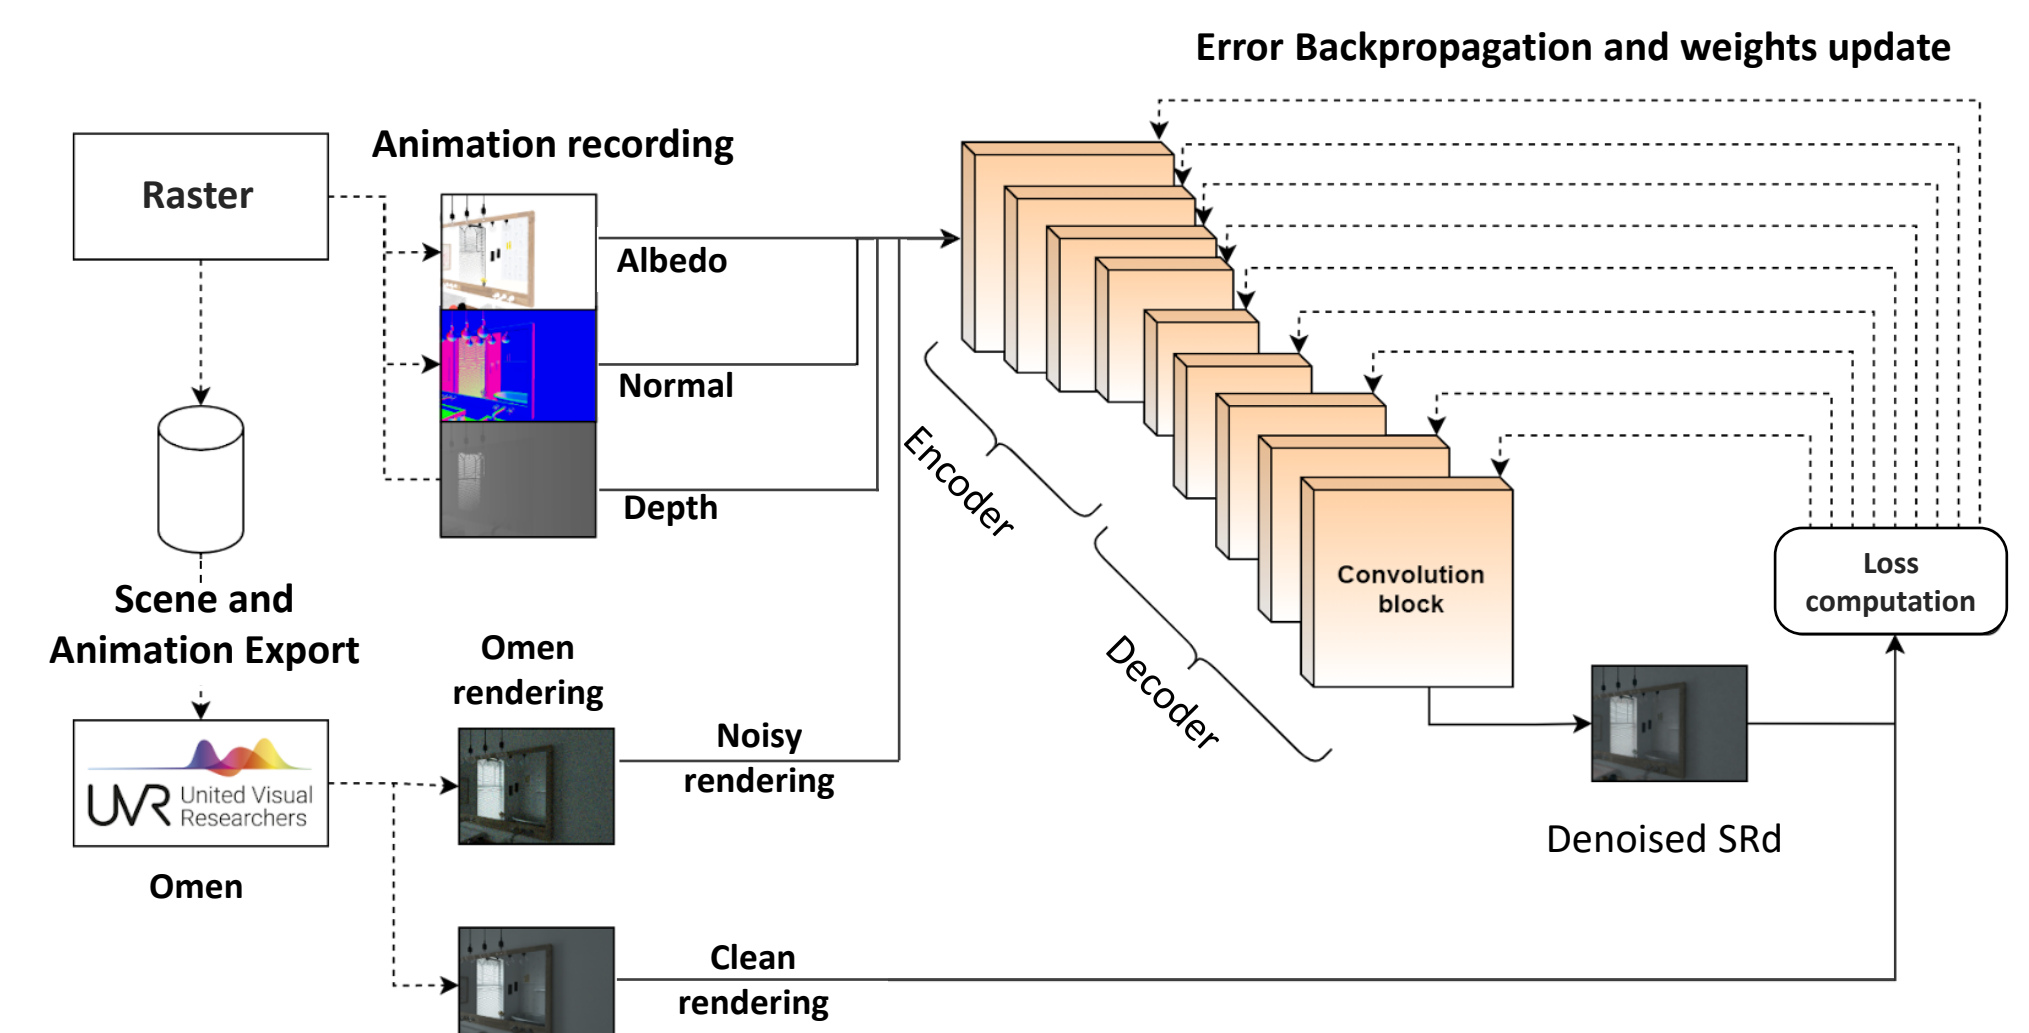


Figure 2: Spectral rendering denoiser design (zoom from step 1 of figure 1)

Training and implementation

Dataset: This dataset is composed of 700 pairs of reference and noisy input (210 for testing, 490 for training) at a resolution of 1024x768 pixels, 47 different points of view were taken from each of 4 scenes exhibiting various light phenomena (diffraction, reflections etc.).

Spectral rendering: All images were rendered with Omen [8], a spectral and polarized path tracer for predictive rendering commercialized by United Visual Researchers (UVR).



Training: The SRd 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.

RESULTS

Contrary to the OptiX and Intel denoisers, our denoiser 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.



Method	MSE	PSNR	SSIM
GT.L + Intel.C	62.454	34.946	0.999
GT.L + SRd.C	30.204	38.101	0.999
GT.C + Intel.L	867.81	23.518	0.828
GT.C + SRd.L	1002.680	22.890	0.799

Table 1: Measures of luminance/chrominance merging of results from GT and denoisers (Intel, SRd)

In table 1, when comparing denoised luminance with ground-truth colors (GT.C), we obtain higher MSE with our denoiser than with the one from Intel (1002.680 vs. 867.81).

However, our denoiser performs far better when looking at denoised colors, with an MSE divided in half (30.204 vs. 62.454). This finding led us to step 3 of our pipeline, which combines the denoised colors of our SRd with the denoised luminance of the Intel denoiser.

Method	MSE	PSNR	SSIM
SRd	1029.242	22.776	0.798
Intel	921.562	23.256	0.827
OptiX	1 0133.635	12.844	0.461
Final	900.713	23.356	0.827

Table 2: Similarity between denoiser outputs and ground truth

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 2.

REFERENCES

- [1] ZWICKER M., JAROSZ W., LEHTINEN J., MOON B., RA-MAMOORTHY 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.
- [2] CHAITANYA C. R. A., KAPLAYAN A. S., SCHIED C., SALVI M., LEFOHN A., NOWROUZSAHRAI 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.
- [3] 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.
- [4] NVIDIA : Optix denoiser. <https://developer.nvidia.com/rtx/ray-tracing/optix>, 2017.
- [5] INTEL : Intel open image denoise. <https://www.openimagedenoise.org>, 2021.
- [6] WILKIE A., NAWAZ S., DROSKE M., WEIDLICH A.HANIKAJ.: Hero Wavelength Spectral Sampling. *Computer Graphics Forum* 33, 4 (2014),123–131.
- [7] 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.
- [8] UNITED VISUAL RESEARCHERS (UVR) : Omen rendering en-gine. <https://www.united-vr.com/>, 2022.

