Nonlinear Noise2Noise for Efficient Monte Carlo Denoiser Training

Andrew Tinits
University of Waterloo
Waterloo, Ontario, Canada
amtinits@uwaterloo.ca

Stephen Mann
University of Waterloo
Waterloo, Ontario, Canada
smann@uwaterloo.ca

Abstract

This work was originally presented as a technical paper at SIGGRAPH Asia 2025. The full paper can be found at the project page.

The Noise2Noise method allows for training machine learning-based denoisers with pairs of input and target images where both the input and target can be noisy. This removes the need for training with clean target images, which can be difficult to obtain. However, Noise2Noise training has a major limitation: nonlinear functions applied to the noisy targets will skew the results. This bias occurs because the nonlinearity makes the expected value of the noisy targets different from the clean target image. Since nonlinear functions are common in image processing, avoiding them limits the types of preprocessing that can be performed on the noisy targets. Our main insight is that certain nonlinear functions can be applied to the noisy targets without adding significant bias to the results. We develop a theoretical framework for analyzing the effects of these nonlinearities, and describe a class of nonlinear functions with minimal bias.

We demonstrate our method on the denoising of high dynamic range (HDR) images produced by Monte Carlo rendering, where generating high-sample count reference images can be prohibitively expensive. Noise2Noise training can have trouble with HDR images, where the training process is overwhelmed by outliers and performs poorly. We consider a commonly used method of addressing these training issues: applying a nonlinear tone mapping function to the model output and target images to reduce their dynamic range. This method was previously thought to be incompatible with Noise2Noise training because of the nonlinearities involved. We show that certain combinations of loss functions and tone mapping functions can reduce the effect of outliers while introducing minimal bias. We apply our method to an existing machine learning-based Monte Carlo denoiser, where the original implementation was trained with high-sample count reference images. Our results approach those of the original implementation, but are produced using only noisy training data.

CCS Concepts

• **Computing methodologies → Image processing**; **Ray tracing**; *Neural networks.*

Keywords

Noise2Noise, nonlinear functions, Jensen gap, Monte Carlo denoising, high dynamic range, tone mapping

ACM Reference Format:
Andrew Tinits and Stephen Mann. 2025. Nonlinear Noise2Noise for Efficient Monte Carlo Denoiser Training. In *SIGGRAPH Asia 2025 Conference Papers (SA Conference Papers '25), December 15–18, 2025, Hong Kong, Hong Kong*. ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3757377.3763931

Poster at Graphics Interface 2026 (GI '26), Waterloo, ON, Canada

ACM ISBN 979-8-4007-2137-3/2025/12
https://doi.org/10.1145/3757377.3763931

1 Introduction

There are many applications in image processing where clean images are hard to obtain, such as with long-exposure photography or Monte Carlo rendering. In these cases, the noise is gradually reduced through a lengthy process, making clean images expensive to produce. Denoising offers a solution to this problem, where the image generating process can be stopped before convergence, and post-processing techniques can be applied to reduce the remaining noise to an acceptable level. Denoising has contributed to recent successes in adapting Monte Carlo rendering for real-time applications, for example. Common denoising methods use either hand-tuned filters or machine learning models to reduce the noise.

Learning-based denoisers are trained on large datasets of images and their associated feature data. In the supervised learning setting, these datasets contain pairs of noisy images with their corresponding clean reference images. However, difficulty in obtaining clean references presents a challenge to training denoisers with supervised learning. A potential solution to this problem is to remove the need for clean reference images by changing to a different learning setting where clean references are not required. One such method is Noise2Noise, a weakly supervised learning approach where denoisers are trained with pairs of noisy images [Lehtinen et al. 2018]. For the common case where the noise distribution has zero mean, Noise2Noise is used with the L_2 loss to recover the expectation of the noisy targets $\mathbb{E}[\hat{y}]$, which is equal to the clean target y [Lehtinen et al. 2018]. The error from the noisy targets approaches zero as the number of training examples increases [Lehtinen et al. 2018].

Noise2Noise, in the zero-mean noise setting, has a major limitation: nonlinear functions applied to the noisy targets will skew the results [Lehtinen et al. 2018]. This bias occurs because the nonlinearity φ makes the expected value of the noisy targets $\mathbb{E}[\varphi(\hat{y})]$ different from the clean target $\varphi(\mathbb{E}[\hat{y}])$ [Lehtinen et al. 2018], since in general $\mathbb{E}[\varphi(\hat{y})] \neq \varphi(\mathbb{E}[\hat{y}])$. Nonlinear functions are common in image processing, so the current practice of avoiding nonlinearity limits the types of preprocessing that can be performed on the noisy targets. Our main insight is that certain nonlinear functions can be applied to the noisy targets without adding significant bias to the results. We develop a theoretical framework, based on the Jensen gap

$\mathbb{E}[\varphi(\hat{y})] - \varphi(\mathbb{E}[\hat{y}])$ [Jensen 1906], for analyzing the effects of these nonlinearities on Noise2Noise training. We then use our framework to describe a class of nonlinear functions with minimal bias.

Noise2Noise training has been widely adopted for other denoising problems, but has not seen wide adoption for the denoising of Monte Carlo rendered images. Monte Carlo rendering is a randomized algorithm that samples the space of possible light paths through a scene. Although this process converges in the limit to the correct image, achieving a low noise level can require many thousands of samples, which can be prohibitively expensive. Monte Carlo noise has zero mean, which makes it an ideal candidate for Noise2Noise training. However, Monte Carlo renderings often have a high dynamic range (HDR) spanning many orders of magnitude, especially at low samples per pixel (spp). Noise2Noise training can have trouble with HDR images, where the L_2 loss is overwhelmed by outliers and the training fails to converge [Lehtinen et al. 2018]. The Noise2Noise authors offer a solution to this problem with their L_{HDR} loss function [Lehtinen et al. 2018], but in our testing this solution is insufficient, and the training still performs poorly. We attribute this difference in our findings to our more complex model and our significantly larger and more varied training dataset.

One existing solution to the HDR training issues is to apply a nonlinear tone mapping function to the model output and target images to reduce their dynamic range [Gharbi et al. 2019]. This method was previously thought to be incompatible with Noise2Noise training because of the nonlinearities involved [Lehtinen et al. 2018]. We show that this method can in fact be used with Noise2Noise training. We apply our theoretical framework to show that certain combinations of loss functions and tone mapping functions can reduce the effect of outliers while introducing minimal bias to the results. Our experimental results confirm the effectiveness of this approach. We apply our method to an existing machine learning-based Monte Carlo denoiser [Gharbi et al. 2019], where the original implementation was trained with high-sample count reference images. Our denoising results approach those of the original implementation, but are produced using only low-sample count noisy training data. Monte Carlo denoising of varied scenes with complex models is therefore an example of an application where training from noisy data is made possible by our nonlinear Noise2Noise method.

2 Related Work

2.1 Weakly Supervised Image Denoising

Noise2Noise [Lehtinen et al. 2018] was not the first work to recognize that clean reference images are not necessary for image denoising. Outside of machine learning methods, there are classic algorithms such as NLM [Buades et al. 2005] and BM3D [Dabov et al. 2007] that denoise a single noisy image without requiring any additional information. These methods look at statistical information from the rest of the image to determine how to denoise a particular pixel. Within machine learning, Noise2Noise is the first denoising method that can learn from noisy images without requiring prior knowledge of either the noise distribution or the distribution of the clean images [Lehtinen et al. 2018]. Other methods that can train on noisy images, such as AmbientGAN [Bora et al. 2018] and GradNet [Guo et al. 2019], require either an explicit statistical model of the noise (AmbientGAN), or additional regularization that makes assumptions about the clean results (GradNet). In the case of Monte Carlo denoising, the noise distribution cannot be characterized analytically, so methods with the former requirement cannot be used [Lehtinen et al. 2018].

Since the publication of Noise2Noise, several works have extended the idea so that only individual noisy images are required for training, rather than the noisy image pairs required by Noise2Noise [Lehtinen et al. 2018]. These works include Noise2Void [Krull et al. 2019], Noise2Self [Batson and Royer 2019], and the work of Laine et al. [2019]. Our idea of applying nonlinearity to the noisy targets can equally be used with these methods and others that make use of the Noise2Noise property. Single-image denoising is useful in settings where noisy image pairs are difficult to obtain, such as in biomedical imaging. However, these methods generally come with reduced performance when compared to Noise2Noise because they have less information available during training [Krull et al. 2019]. It therefore makes sense to use Noise2Noise in settings such as Monte Carlo denoising, where noisy image pairs are easy to obtain.

This work is a continuation of previous experimental work by the first author [Tinits 2022]. To our knowledge, these two works are the first to successfully apply nonlinear functions to the noisy targets in Noise2Noise training.

2.2 Monte Carlo Denoising

Monte Carlo denoising methods can be separated into two main categories: traditional algorithms, and machine learning-based methods. We briefly review machine-learning based methods below. For a more comprehensive review, see the survey papers by Zwicker et al. [2015] and Huo and Yoon [2021].

The use of machine learning for Monte Carlo denoising was initiated by Kalantari et al. [2015], who used a multi-layer perceptron (MLP) to predict the parameters of a cross-bilateral filter. Bako et al. improved this idea by using a deep convolutional neural network (CNN) to predict filter kernels separately for each pixel [2017]. This kernel prediction idea is used by many later works [Back et al. 2020; Gharbi et al. 2019; Lin et al. 2020, 2021; Meng et al. 2020; Munkberg and Hasselgren 2020; Vogels et al. 2018]. Other recent works take the simpler approach of predicting the pixel radiances directly [Chaitanya et al. 2017; Guo et al. 2019; Lu et al. 2020; Wong and Wong 2019; Xu et al. 2019; Yang et al. 2018, 2019]. Gharbi et al. propose another alternative where splatting kernels are used to allow individual pixels to determine their own contributions [2019]. This concept is extended by Munkberg and Hasselgren [2020], who divide the samples into layers and filter them separately to improve efficiency, and by Lin et al. [2021] and Cho et al. [2021], who add information about individual Monte Carlo paths. Other recent works use different model architectures, such as generative adversarial networks (GANs) [Lu et al. 2021, 2020; Xu et al. 2019], self-attention [Chen et al. 2023, 2024; Oh and Moon 2024; Yu et al. 2021], and diffusion models [Vavilala et al. 2024]. Several works in recent years have explored unsupervised learning [Back et al. 2022; Choi et al. 2024; Guo et al. 2019; Xu et al. 2020]. Weakly supervised learning for Monte Carlo denoising was studied by Cho et al. [2021], but their method uses high-sample count reference images for training.

References

Jonghee Back, Binh-Son Hua, Toshiya Hachisuka, and Bochang Moon. 2020. Deep combiner for independent and correlated pixel estimates. *ACM Transactions on Graphics* 39, 6 (Nov. 2020), 242:1–242:12. doi:10.1145/3414685.3417847

Jonghee Back, Binh-Son Hua, Toshiya Hachisuka, and Bochang Moon. 2022. Self-Supervised Post-Correction for Monte Carlo Denoising. In *ACM SIGGRAPH 2022 Conference Proceedings (SIGGRAPH '22)*. Association for Computing Machinery, New York, NY, USA, 1–8. doi:10.1145/3528233.3530730

Steve Bako, Thijs Vogels, Brian McWilliams, Mark Meyer, Jan Novák, Alex Harvill, Pradeep Sen, Tony DeRose, and Fabrice Rousselle. 2017. Kernel-Predicting Convolutional Networks for Denoising Monte Carlo Renderings. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2017)* 36, 4 (July 2017). doi:10.1145/3072959.3073708

Joshua D. Batson and Loic A. Royer. 2019. Noise2Self: Blind Denoising by Self-Supervision. *ICML* (2019). http://proceedings.mlr.press/v97/batson19a.html

Ashish Bora, Eric Price, and Alexandros G. Dimakis. 2018. AmbientGAN: Generative models from lossy measurements. In *ICLR*. https://openreview.net/forum?id=Hy7fDog0b

Antoni Buades, Bartomeu Coll, and Jean-Michel Morel. 2005. A non-local algorithm for image denoising. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, Vol. 2. IEEE, 60–65. doi:10.1109/CVPR.2005.38

Chakravarty R. Alla Chaitanya, Anton S. Kaplanyan, Christoph Schied, Marco Salvi, Aaron Lefohn, Derek Nowrouzezahrai, and Timo Aila. 2017. Interactive reconstruction of Monte Carlo image sequences using a recurrent denoising autoencoder. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 98. doi:10.1145/3072959.3073601

Bingyi Chen, Zengyu Liu, Li Yuan, Zhitao Liu, Yi Li, Guan Wang, and Ning Xie. 2023. Monte Carlo Denoising via Multi-scale Auxiliary Feature Fusion Guided Transformer. In *SIGGRAPH Asia 2023 Technical Communications (SA '23)*. Association for Computing Machinery, New York, NY, USA, 1–4. doi:10.1145/3610543.3626179

Chuhao Chen, Yuze He, and Tzu-Mao Li. 2024. Temporally Stable Metropolis Light Transport Denoising using Recurrent Transformer Blocks. *ACM Trans. Graph.* 43, 4 (July 2024), 123:1–123:14. doi:10.1145/3658218

In-Young Cho, Yuchi Huo, and Sung-Eui Yoon. 2021. Weakly-supervised contrastive learning in path manifold for Monte Carlo image reconstruction. *ACM Transactions on Graphics* 40, 4 (July 2021), 38:1–38:14. doi:10.1145/3450626.3459876

Hajin Choi, Seokpyo Hong, Inwoo Ha, Nahyup Kang, and Bochang Moon. 2024. Online Neural Denoising with Cross-Regression for Interactive Rendering. *ACM Trans. Graph.* 43, 6 (Nov. 2024), 221:1–221:12. doi:10.1145/3687938

Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. 2007. Image Denoising by Sparse 3-D Transform-Domain Collaborative Filtering. *IEEE Transactions on Image Processing* 16, 8 (Aug. 2007), 2080–2095. doi:10.1109/TIP.2007.901238

Michaël Gharbi, Tzu-Mao Li, Miika Aittala, Jaakko Lehtinen, and Frédo Durand. 2019. Sample-based Monte Carlo Denoising Using a Kernel-splatting Network. *ACM Trans. Graph.* 38, 4 (July 2019), 125:1–125:12. doi:10.1145/3306346.3322954

Jie Guo, Mengtian Li, Quewei Li, Yuting Qiang, Bingyang Hu, Yanwen Guo, and Ling-Qi Yan. 2019. GradNet: unsupervised deep screened poisson reconstruction for gradient-domain rendering. *ACM Transactions on Graphics* 38, 6 (Nov. 2019), 223:1–223:13. doi:10.1145/3355089.3356538

Yuchi Huo and Sung-eui Yoon. 2021. A survey on deep learning-based Monte Carlo denoising. *Computational Visual Media* 7, 2 (June 2021), 169–185. doi:10.1007/s41095-021-0209-9

J. L. W. V. Jensen. 1906. Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta Mathematica* 30 (Jan. 1906), 175–193. doi:10.1007/BF02418571 Publisher: Institut Mittag-Leffler.

Nima Khademi Kalantari, Steve Bako, and Pradeep Sen. 2015. A Machine Learning Approach for Filtering Monte Carlo Noise. *ACM Trans. Graph.* 34, 4 (July 2015), 122:1–122:12. doi:10.1145/2766977

Alexander Krull, Tim-Oliver Buchholz, and Florian Jug. 2019. Noise2Void - Learning Denoising From Single Noisy Images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2129–2137. doi:10.1109/CVPR.2019.00223

Samuli Laine, Tero Karras, Jaakko Lehtinen, and Timo Aila. 2019. High-Quality Self-Supervised Deep Image Denoising. In *Advances in Neural Information Processing Systems 32 (NeurIPS 2019)*, Vol. 32. 6970–6980. https://proceedings.neurips.cc/paper/2019/hash/2119b8d43eafcf353e07d7cb5554170b-Abstract.html

Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila. 2018. Noise2Noise: Learning Image Restoration without Clean Data. In *International Conference on Machine Learning*. PMLR, 2965–2974. http://proceedings.mlr.press/v80/lehtinen18a.html

Weiheng Lin, Beibei Wang, Lu Wang, and Nicolas Holzschuch. 2020. A Detail Preserving Neural Network Model for Monte Carlo Denoising. *Computational Visual Media* 6 (April 2020), 157–168. doi:10.1007/s41095-020-0167-7

Weiheng Lin, Beibei Wang, Jian Yang, Lu Wang, and Ling-Qi Yan. 2021. Path-based Monte Carlo Denoising Using a Three-Scale Neural Network. *Computer Graphics Forum* 40, 1 (Feb. 2021), 369–381. doi:10.1111/cgf.14194

Yifan Lu, Siyuan Fu, Xiao Hua Zhang, and Ning Xie. 2021. Denoising Monte Carlo renderings via a multi-scale featured dual-residual GAN. *The Visual Computer* 37, 9 (Sept. 2021), 2513–2525. doi:10.1007/s00371-021-02204-4

Yifan Lu, Ning Xie, and Heng Tao Shen. 2020. DMCR-GAN: Adversarial Denoising for Monte Carlo Renderings with Residual Attention Networks and Hierarchical Features Modulation of Auxiliary Buffers. In *SIGGRAPH Asia 2020 Technical Communications*. Association for Computing Machinery, 1–4. doi:10.1145/3410700.3425426

Xiaoxu Meng, Quan Zheng, Amitabh Varshney, Gurprit Singh, and Matthias Zwicker. 2020. Real-time Monte Carlo Denoising with the Neural Bilateral Grid. In *Eurographics Symposium on Rendering (EGSR)*. The Eurographics Association. doi:10.2312/sr.20201133

Jacob Munkberg and Jon Hasselgren. 2020. Neural Denoising with Layer Embeddings. In *Computer Graphics Forum*, Vol. 39. 1–12. doi:10.1111/cgf.14049

Geunwoo Oh and Bochang Moon. 2024. Joint self-attention for denoising Monte Carlo rendering. *The Visual Computer* 40, 7 (July 2024), 4623–4634. doi:10.1007/s00371-024-03446-8

Andrew Tinits. 2022. *Learning Sample-Based Monte Carlo Denoising from Noisy Training Data*. Master's thesis. University of Waterloo, Waterloo, Ontario, Canada. http://hdl.handle.net/10012/18071

Vaibhav Vavilala, Rahul Vasanth, and David Forsyth. 2024. Denoising Monte Carlo Renders With Diffusion Models. doi:10.48550/arXiv.2404.00491

Thijs Vogels, Fabrice Rousselle, Brian McWilliams, Gerhard Röthlin, Alex Harvill, David Adler, Mark Meyer, and Jan Novák. 2018. Denoising with Kernel Prediction and Asymmetric Loss Functions. *ACM Transactions on Graphics* 37, 4 (July 2018), 124:1–124:15. doi:10.1145/3197517.3201388

Kin-Ming Wong and Tien-Tsin Wong. 2019. Deep residual learning for denoising Monte Carlo renderings. *Computational Visual Media* 5, 3 (Sept. 2019), 239–255. doi:10.1007/s41095-019-0142-3

Bing Xu, Junfei Zhang, Rui Wang, Kun Xu, Yong-Liang Yang, Chuan Li, and Rui Tang. 2019. Adversarial Monte Carlo denoising with conditioned auxiliary feature modulation. *ACM Transactions on Graphics* 38, 6 (Nov. 2019), 224:1–224:12. doi:10.1145/3355089.3356547

Zilin Xu, Qiang Sun, Lu Wang, Yanning Xu, and Beibei Wang. 2020. Unsupervised Image Reconstruction for Gradient-Domain Volumetric Rendering. *Computer Graphics Forum* 39, 7 (2020), 193–203. doi:10.1111/cgf.14137

Xin Yang, Dawei Wang, Wenbo Hu, Lijing Zhao, Xinglin Piao, Dongsheng Zhou, Qiang Zhang, Baocai Yin, Qiang Cai, and Xiaopeng Wei. 2018. Fast Reconstruction for Monte Carlo Rendering Using Deep Convolutional Networks. *IEEE Access* 7 (Dec. 2018), 21177–21187. doi:10.1109/ACCESS.2018.2886005

Xin Yang, Dawei Wang, Wenbo Hu, Li-Jing Zhao, Bao-Cai Yin, Qiang Zhang, Xiao-Peng Wei, and Hongbo Fu. 2019. DEMC: A Deep Dual-Encoder Network for Denoising Monte Carlo Rendering. *Journal of Computer Science and Technology* 34 (Sept. 2019), 1123–1135. doi:10.1007/s11390-019-1964-2

Jiaqi Yu, Yongwei Nie, Chengjiang Long, Wenju Xu, Qing Zhang, and Guiqing Li. 2021. Monte Carlo denoising via auxiliary feature guided self-attention. *ACM Transactions on Graphics (TOG)* 40, 6 (Dec. 2021), 1–13. doi:10.1145/3478513.3480565

Matthias Zwicker, Wojciech Jarosz, Jaakko Lehtinen, Bochang Moon, Ravi Ramamoorthi, Fabrice Rousselle, Pradeep Sen, Cyril Soler, and Sung-Eui Yoon. 2015. Recent Advances in Adaptive Sampling and Reconstruction for Monte Carlo Rendering. *Computer Graphics Forum (Proceedings of Eurographics - State of the Art Reports)* 34, 2 (May 2015), 667–681. doi:10.1111/cgf.12592