Fast Blue-Noise Generation via Unsupervised Learning

*Daniele Giunchi, *Alejandro Sztrajman, Anthony Steed

University College London

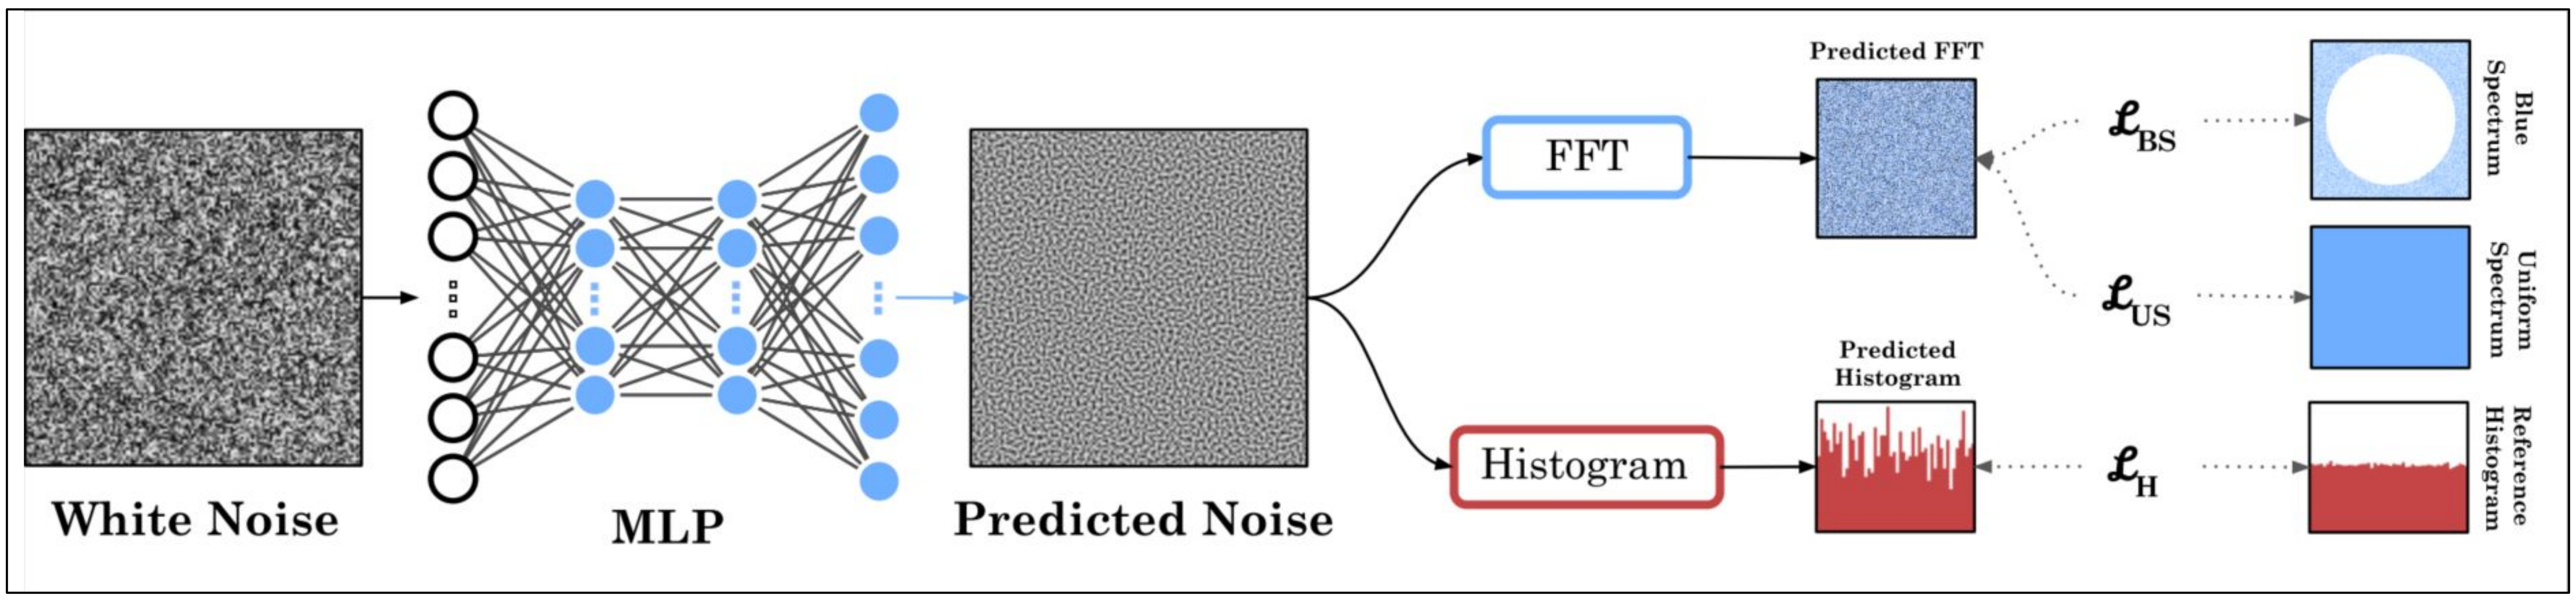

Diagram of network training, combining the three unsupervised losses: uniform spectrum loss, blue spectrum loss, and histogram loss. After training, the network produces blue noise grayscale masks based on white noise input.

Blue noise is known for its uniformity in the spatial domain, avoiding the appearance of structures such as voids and clusters, and it's characterised by a spectrum with minimal low frequency components and no high intensity spikes **[1]**. Because of this property, it has been adopted in a wide range of visual computing applications, which has motivated the development of a variety of generative methods, with different trade-offs in terms of accuracy and computational performance.

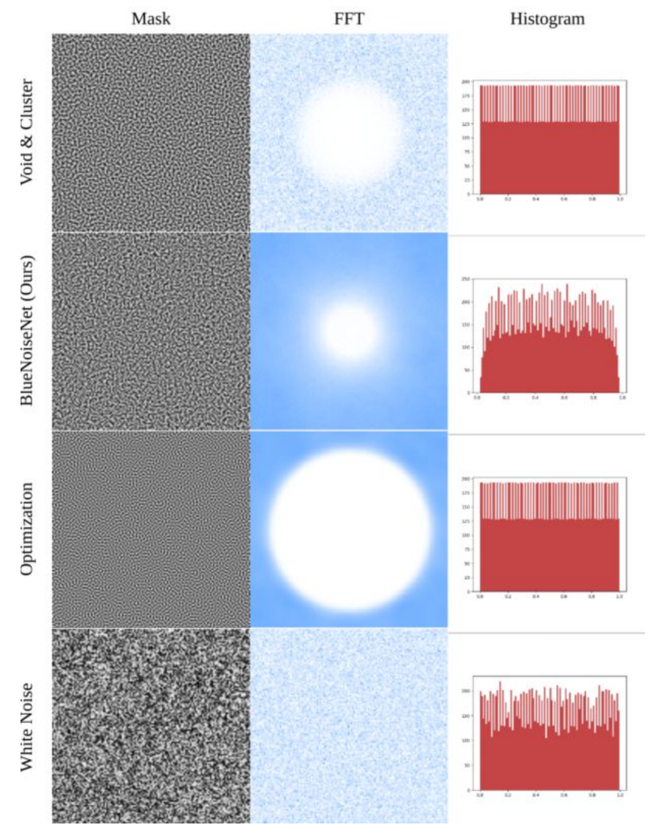

We propose a **novel unsupervised learning approach** that leverages a neural network architecture to generate blue noise masks with **high accuracy and real-time performance**, starting from a white noise input. We evaluate our method by applying the generated noise to image dithering and Monte Carlo integration. Our method is able to produce new blue noise masks much faster than other methods, at around 0.018 s per mask as showed in the figure on the right.

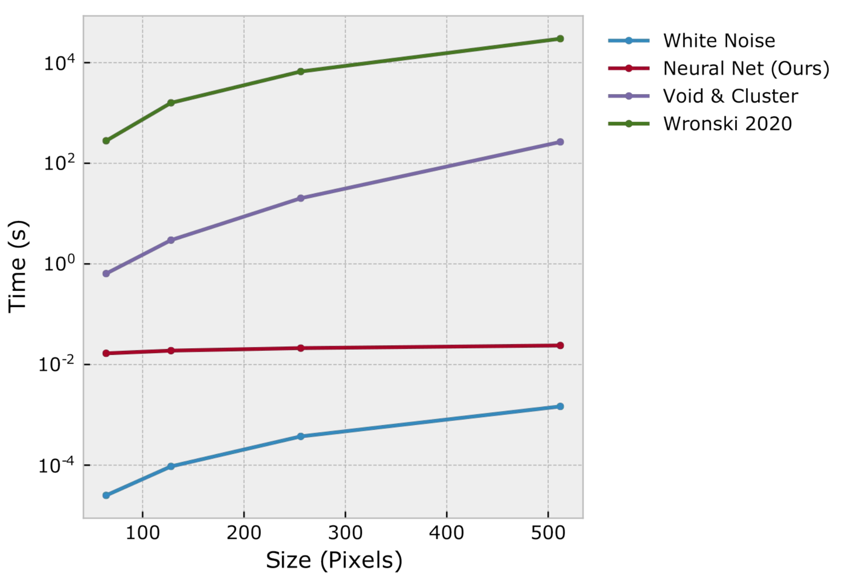

Loss Function

We leverage a linear combination of three unsupervised losses that enforce mathematical constraints on the output noise:

Blue Spectrum Loss:

$$\mathcal{L}_{\mathrm{BS}} = \sum_{ij} \hat{\phi}_{ij} \left[\max\left(0, \frac{\omega_{\mathrm{cutoff}} - r_{ij}^2}{\omega_{\mathrm{cutoff}}}\right) \right]^2$$

The blue spectrum loss $\mathcal{L}_{\mathrm{BS}}$ acts on the Fourier spectrum of the predicted noise, penalising frequencies increasingly above a cut-off threshold

Uniform Spectrum Loss:

$$\mathcal{L}_{\mathrm{US}} = \sum_{ij} \left[\nabla^2 \hat{\phi}_{ij}\right]^2 + \nabla_i \hat{\phi}_{ij} + \nabla_j \hat{\phi}_{ij}$$

The uniform spectrum loss acts on the spectrum of the predicted signal, by penalising large values of its spatial derivatives.

Histogram Loss:

$$\mathcal{L}_{\mathrm{H}} = \sum_{k} [\mathrm{Histogram}(\hat{s})_k - \mathrm{RefH}_k]^2$$

The histogram loss acts on the histogram of intensities of the noise by penalising deviations from a uniform reference histogram.

Dithering

A visual inspection of the figure on the right shows that our method produces dithered reconstructions with lower noise than white noise, and similar in quality to the optimisation-based approach by Wronski **[2]**. The Void & Cluster **[3]** method produces the highest quality reconstruction, although this is at the expense of a high computational cost, which might be prohibitive for real-time applications. The Table on the right shows average root mean square errors (RMSE) for the dithering of the David statue, using 100 different dithering masks for each noise generation method. The ground truth and dithered images are pre-processed with a Gaussian filter.

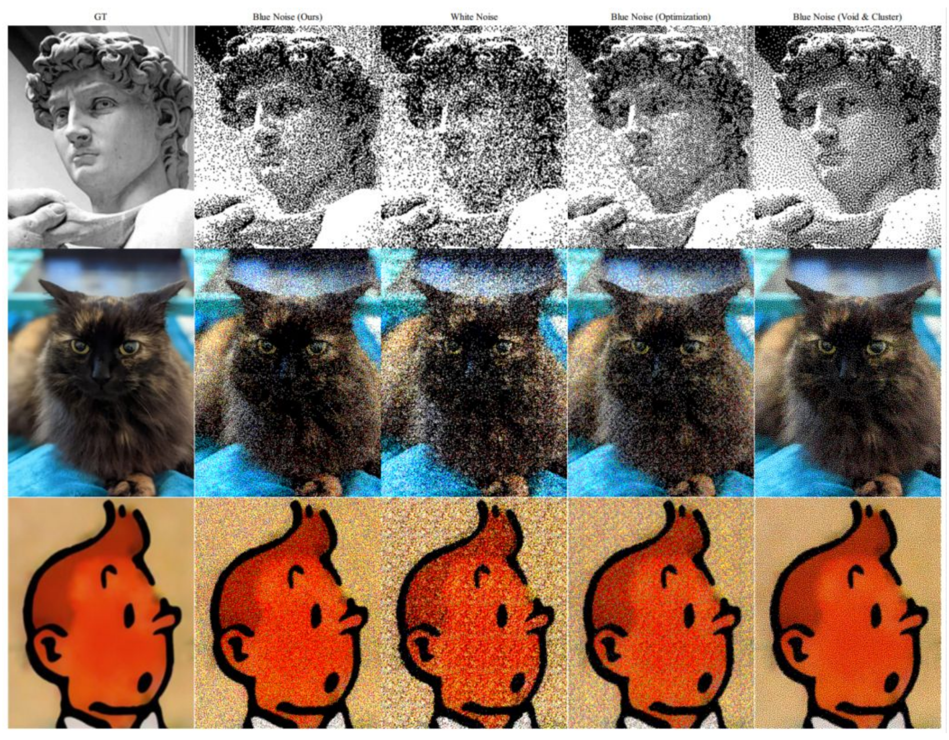

Method	RMSE
White Noise	0.069 ± 0.002
Neural Network (Ours)	0.046 ± 0.001
Wronski 2020	0.046 ± 0.001
Void & Cluster	0.029 ± 0.001

Monte Carlo Integration

We evaluate the blue noise samples generated through the application of Monte Carlo integration **[4]**. In the figure on the right we leverage the generated samples for the integration of two functions, with known analytic integrals. We show the root mean square error of the integral computation through Monte Carlo sampling, as a function of the number of samples used for the integration.

Our neural-based method shows consistently faster convergence than white noise, and in some cases lower error than void & cluster, for low number of samples. White noise and void & cluster present correspondingly the slowest and fastest convergence speeds, in terms of the number of samples, although again at the cost of a high inference time.

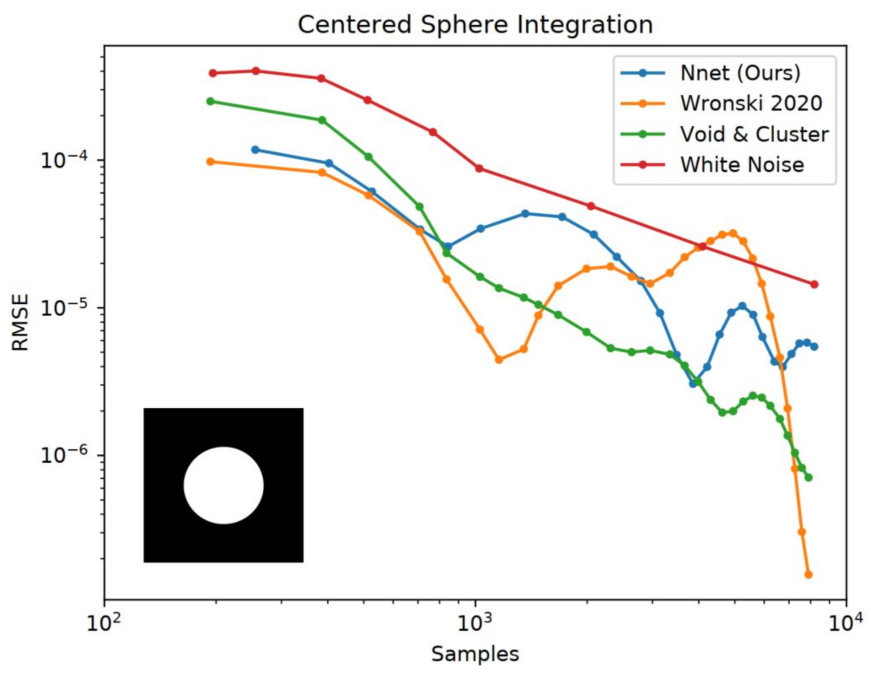

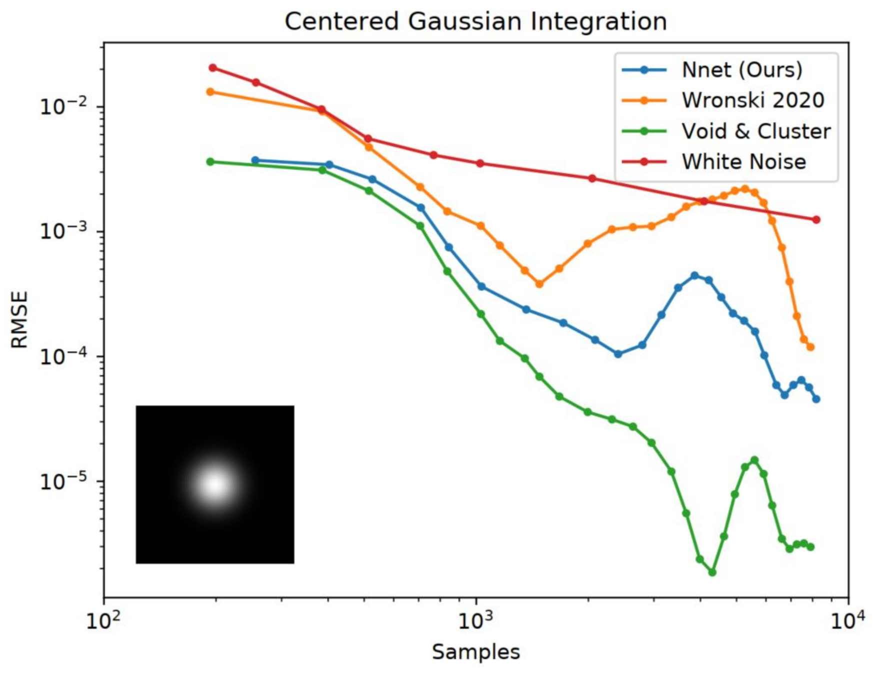

References

[1] D. M. Yan, J. W. Guo, B. Wang, X. P. Zhang and P. Wonka. A Survey of Blue-Noise Sampling and Its Applications. J. Comput. Sci. Technol. 30, pp. 439–452 (2015).

[2] B. Wronski. Optimizing blue noise dithering: backpropagation through Fourier transform and sorting. Retrieved from http://bartwronski.com. Last access 23rd January 2022.

[3] Robert A. Ulichney "Void-and-cluster method for dither array generation", Proc. SPIE 1913, Human Vision, Visual Processing, and Digital Display IV, (8 September 1993).

[4] De Freitas, Nando, and Neil James Gordon. Sequential Monte Carlo methods in practice. Ed. Arnaud Doucet. Vol. 1. No. 2. New York: Springer, 2001.

Fast Blue-Noise Generation via Unsupervised Learning

D.Giunchi, A. Sztrajman, A. Steed

University College London

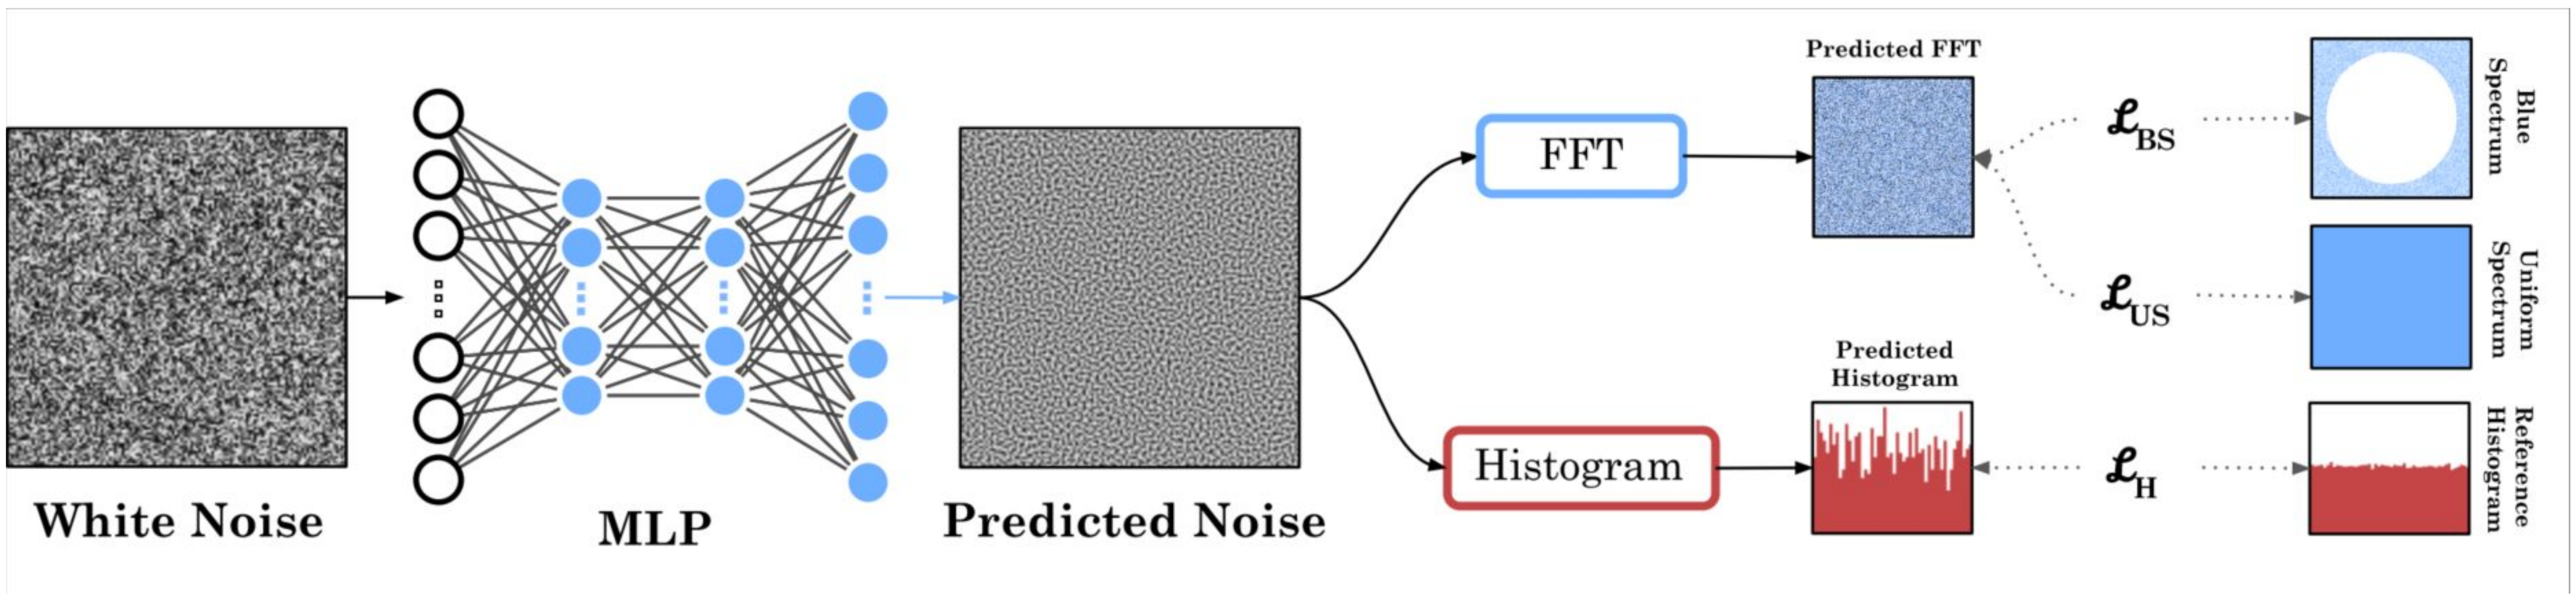

Diagram of network training, combining the three unsupervised losses: uniform spectrum loss, blue spectrum loss, and histogram loss. After training, the network produces blue noise grayscale masks based on white noise input.

Blue noise is known for its uniformity in the spatial domain, avoiding the appearance of structures such as voids and clusters and refers to a spectrum with minimallow-frequency components and no high-intensity spikes**[1]**. Because of this characteristic, it has been adopted in a wide range of visual computing applications, which has motivated the development of a variety of generative methods, with different trade-offs in terms of accuracy and computational performance.

We propose **a novel unsupervised learning approach** that leverages a neural network architecture to generate blue noise masks with **high accuracy and real-time performance**, starting from a white noise input. We evaluate our method by applying the generated noise to image dithering and Monte Carlo integration.

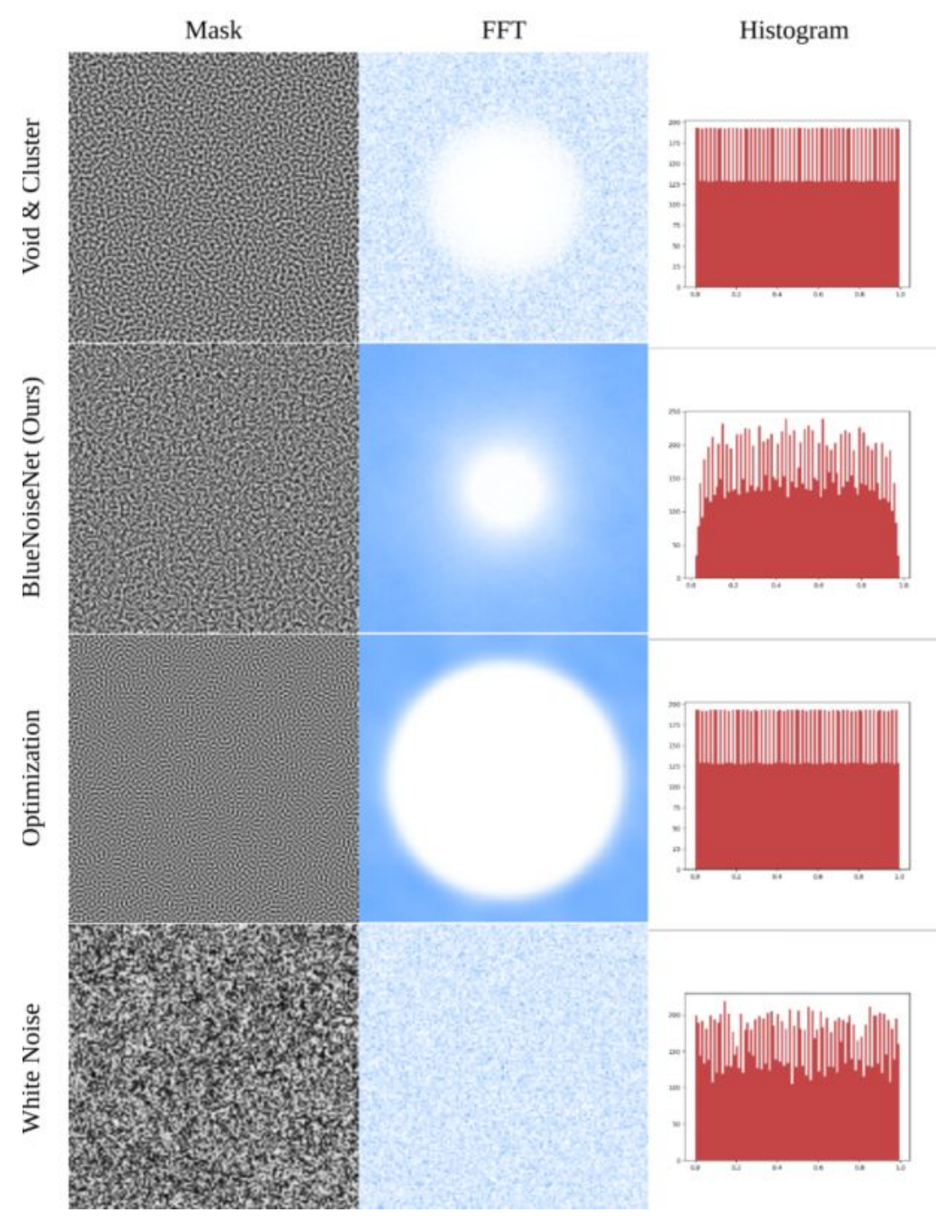

Loss Function

We leverage a linear combination of **three** unsupervised losses that enforce mathematical constraints on the output noise:

Uniform Spectrum Loss:

$$\mathcal{L}_{\mathrm{US}} = \sum_{ij} \left[\nabla^2 \hat{\phi}_{ij}\right]^2 + \nabla_i \hat{\phi}_{ij} + \nabla_j \hat{\phi}_{ij}$$

The uniform spectrum loss also acts on the spectrum of the predicted signal, by penalising large values of its spatial derivatives

Histogram Loss:

$$\mathcal{L}_{\mathrm{H}} = \sum_{k} \left[\mathrm{Histogram}(\hat{s})_k - \mathrm{RefH}_k\right]^2$$

The histogram loss acts on the histogram of intensities of the noise by penalising deviations from a uniform reference histogram

Blue Spectrum Loss:

$$\mathcal{L}_{\mathrm{BS}} = \sum_{ij} \hat{\phi}_{ij} \left[\max\left(0, \frac{\omega_{\mathrm{cutoff}} - r_{ij}^2}{\omega_{\mathrm{cutoff}}}\right)\right]^2$$

The blue spectrum loss LBS acts on the Fourier spectrum of the predicted noise, penalising frequencies increasingly above a cut-off threshold

Dithering

A visual inspection of the Figure on the right shows that our method produces dithered reconstructions with lower noise than white noise, and similar in quality to the optimisation-based approach by Wronski **[2]**. The Void & Cluster **[3]** method produces the highest quality reconstructions, although we will later show that this is at the expense of a high computational cost, which might be prohibitive for real-time applications..The Table on the right shows average root mean square errors (RMSE) for the dithering of the David statue, using 100 different dithering masks for each noise generation method. Bthe ground truth and dithered images are pre-processed with a Gaussian filter.

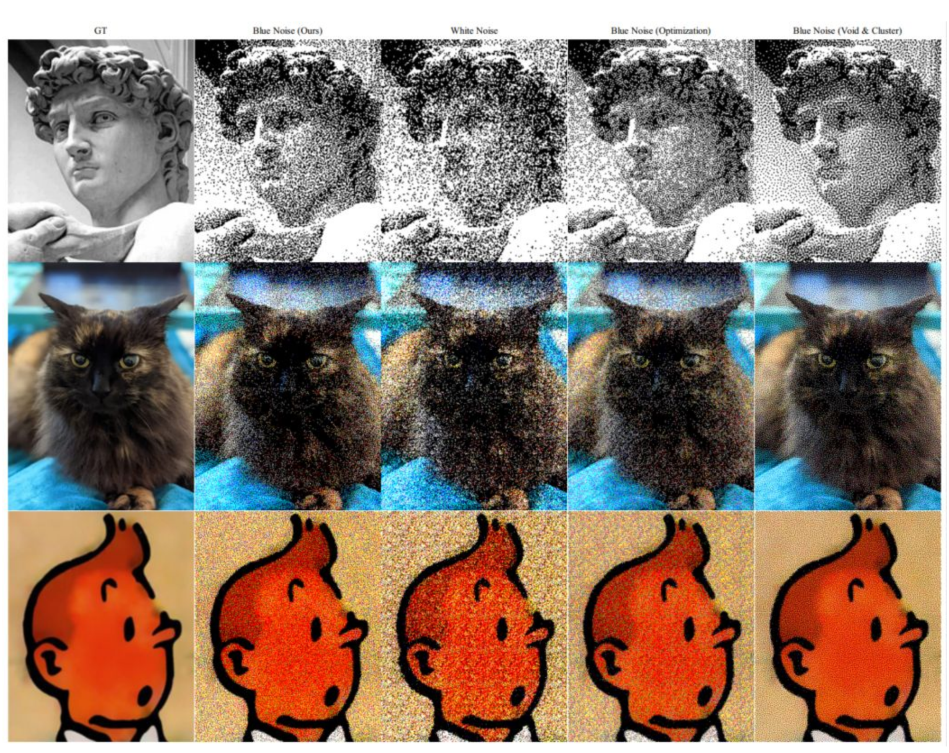

Method	RMSE
White Noise	0.069 ± 0.002
Neural Network (Ours)	0.046 ± 0.001
Wronski 2020	0.046 ± 0.001
Void & Cluster	0.029 ± 0.001

Monte Carlo Integration

We evaluate the blue noise samples generated through the application of Monte Carlo integration**[4]**. In the Figure on the right we leverage the generated samples to the integration of two functions, with known analytic integrals. We show the root mean square error of the integral computation through Monte Carlo sampling, as a function of the number of samples used for the integration. Our neural-based method shows consistently faster convergence than white noise, and in some cases lower error than void & cluster, for low number of samples. White noise and void & cluster present correspondingly the slowest and fastest convergence speeds, in terms of the number of samples. Void & cluster this comes at the cost of a high inference time of the blue noise mask.

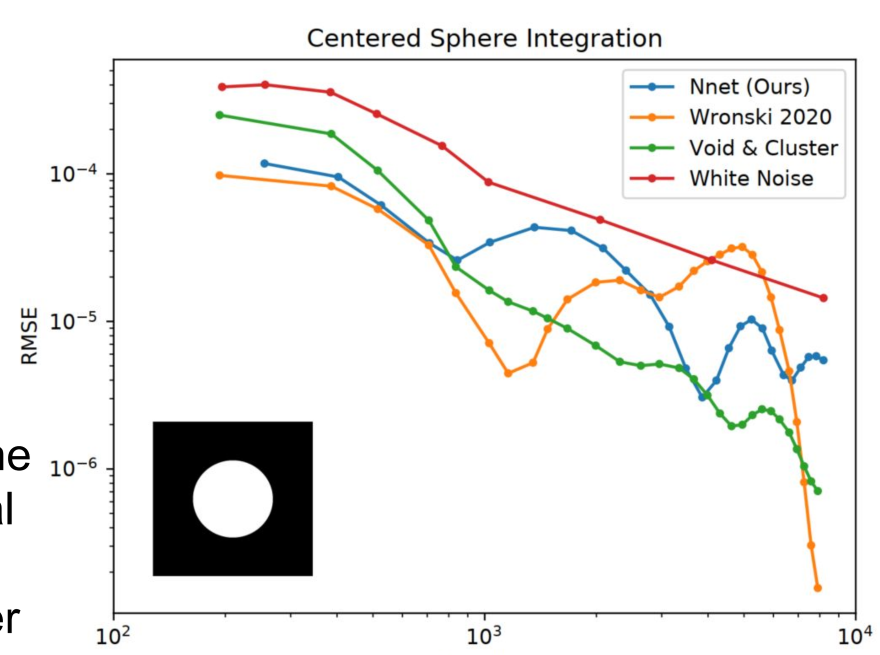

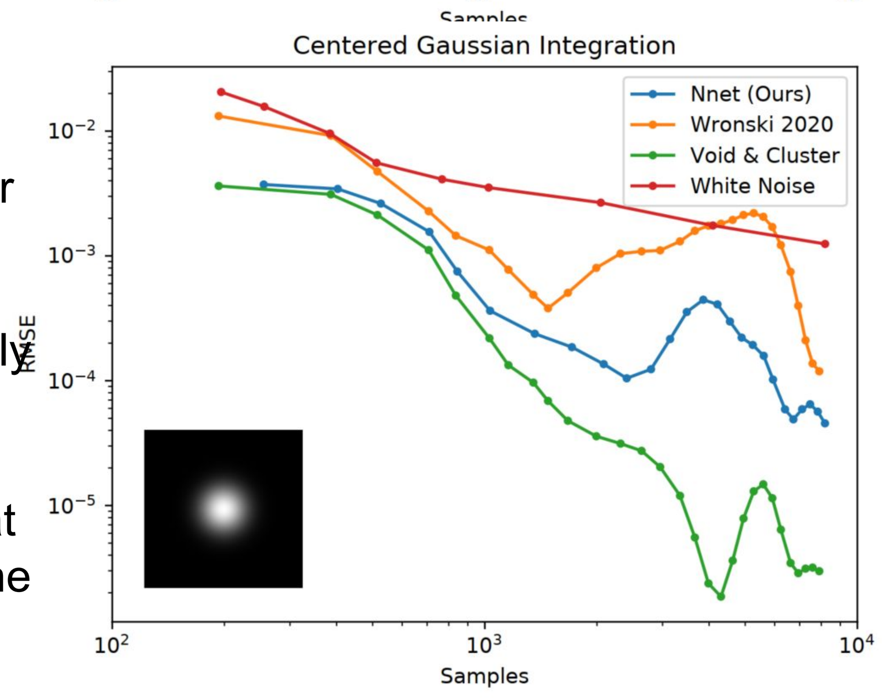

References

[1] D. M. Yan, J. W. Guo, B. Wang, X. P. Zhang and P. Wonka. A Survey of Blue-Noise Sampling and Its Applications. J. Comput. Sci. Technol. 30, pp. 439–452 (2015).

[2] B. Wronski. Optimizing blue noise dithering: backpropagation through Fourier transform and sorting. Retrieved from http://bartwronski.com. Last access 23rd January 2022

[3] Robert A. Ulichney "Void-and-cluster method for dither array generation", Proc. SPIE 1913, Human Vision, Visual Processing, and Digital Display IV, (8 September 1993)

[4] De Freitas, Nando, and Neil James Gordon. Sequential Monte Carlo methods in practice. Ed. Arnaud Doucet. Vol. 1. No. 2. New York: Springer, 2001.