UNIVERSITÄT BERN

SEATTLE, WA

CVPR

JUNE 17-21, 2024

Masked and Shuffled Blind Spot Denoising for Real-World Images

Hamadi Chihaoui Paolo Favaro

Computer Vision Group, Institute of Informatics, University of Bern, Switzerland

Context: Blind Spot Denoising (BSD)

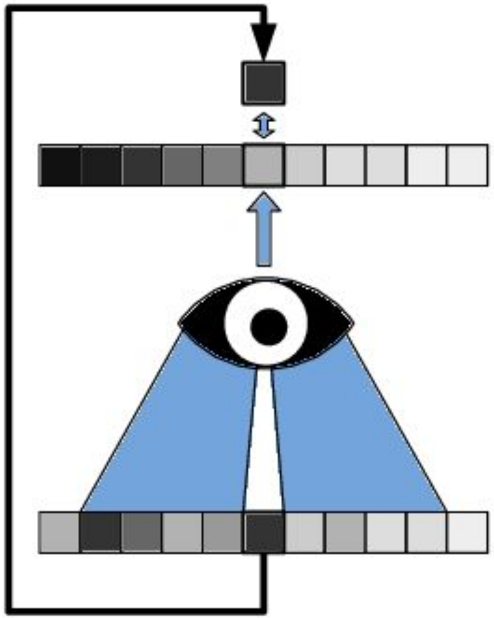

Blind spot denoising [9] concept (source :https://arxiv.org/pdf/1811.10980)

Goal and Motivation

Goal: To restore images corrupted by real noise **without supervision**

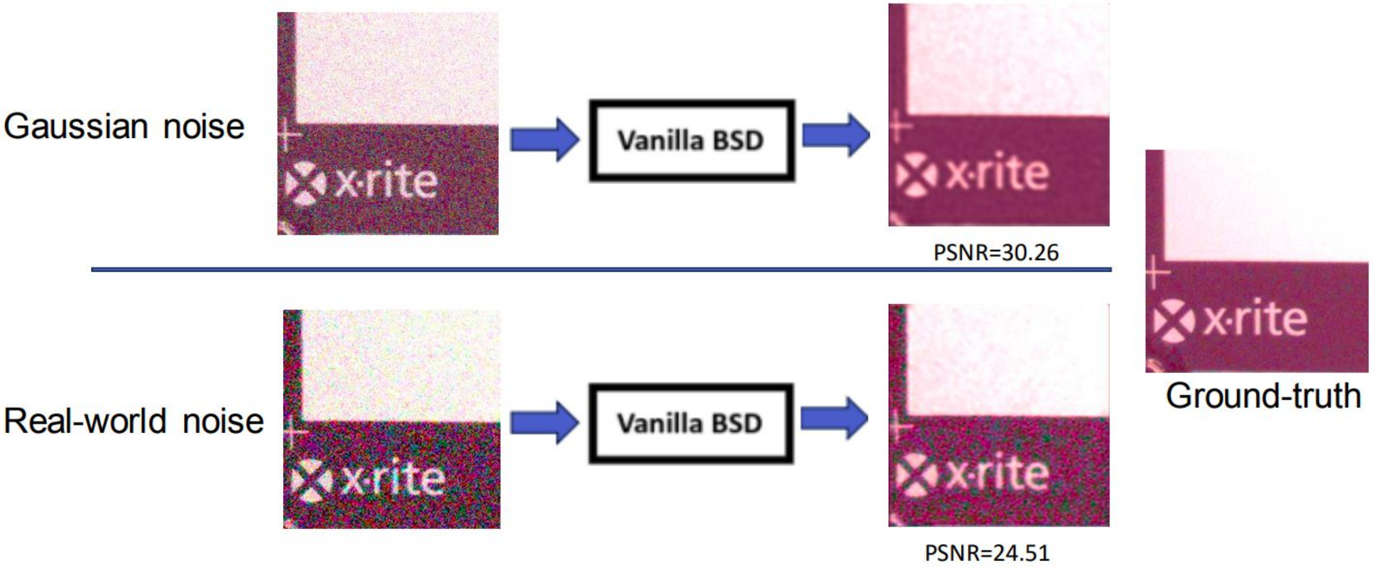

Challenges

- Because real-world noise is spatially correlated, it is more difficult to discriminate from texture patterns than iid noise

- State-of-the-art methods for iid Gaussian noise, such as Blind Spot Denoising (BSD), struggle with real-data denoising

- BSD uses context pixels to predict center pixels; if the noise at the context pixels is correlated with the noise at the center pixels, then the predictions can be biased

Idea

Extend BSD to real-world noise by:

- **adaptive masking ratio** to decrease the prediction bias by reducing the correlation between context pixels and center pixels (larger center masks allow BSD to handle more non local correlations)

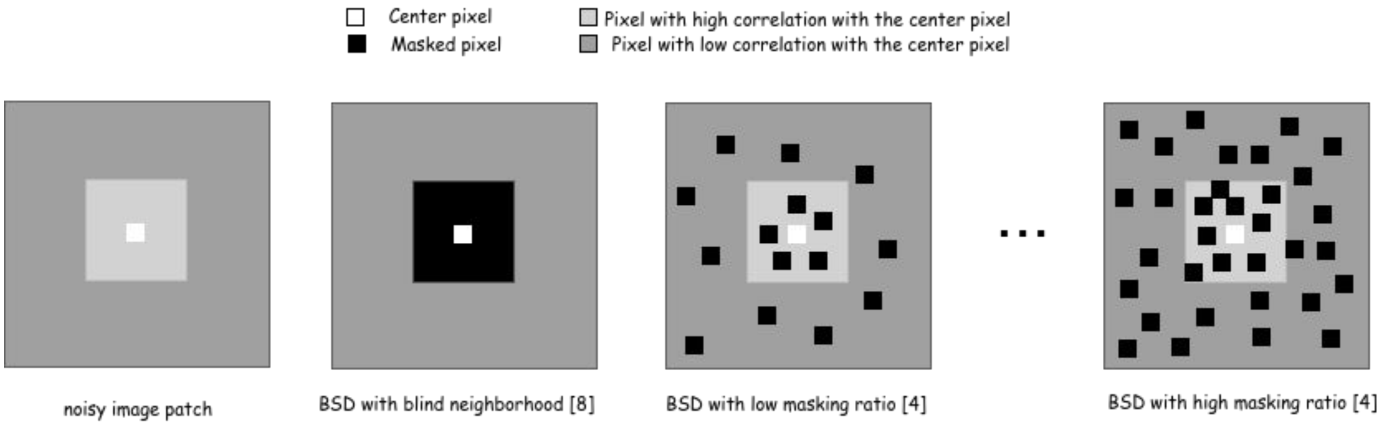

Idea

Extend BSD to real-world noise by:

- using **local pixel shuffling** of pixels with similar colors to break the correlation in the case of a highly spatially correlated regime

Patch corrupted with correlated noise

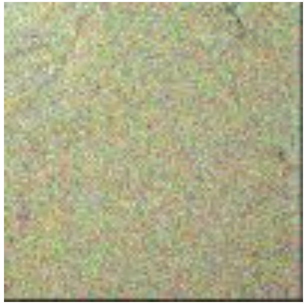

Same patch when locally shuffled

Contributions

- An analysis of BSD, showcasing the impact of various masking ratios on correlated noise, and presenting a method for estimating the noise correlation level
- Introduction of the local pixel shuffling technique to address noise correlation at its source
- Marked enhancements over the baseline BSD and sota results in self-supervised real-world denoising across multiple datasets

Impact of the Masking Ratio on BSD Performance

- Noise covariance $$\Sigma[\mathbf{i},\mathbf{j}] = \begin{cases} \sigma^2 & \mathbf{i}=\mathbf{j} \\ \beta^{\frac{k-\|\mathbf{i}-\mathbf{j}\|}{k}}\sigma^2 & 0<\|\mathbf{i}-\mathbf{j}\|\le k \\ 0 & \text{otherwise} \end{cases}$$
- The effectiveness of BSD is greatly influenced by:

i) the spatial correlation of the noise

ii) the masking ratio

- A lower masking ratio is better for the iid noise while

a higher masking ratio is better suited for highly correlated noise.

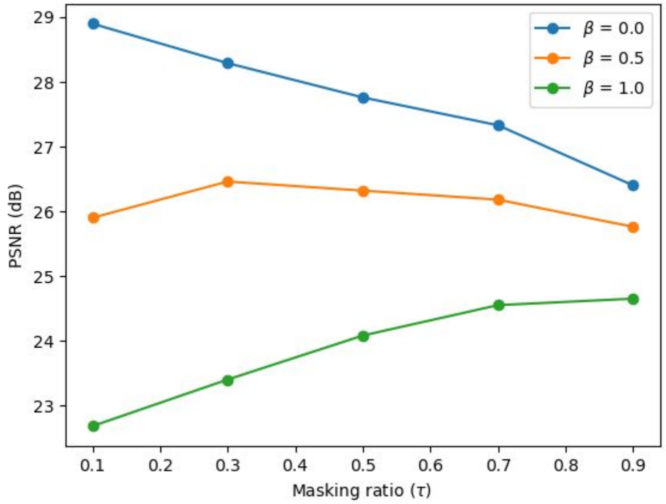

Handling High Noise Correlation with Local Pixel Shuffling

- The noisy image is partitioned into:

  1. flat regions

  2. textured regions

- A random pixel permutation is applied to each

sxs (e.g., s=4) tile within the flat regions.

- The local pixel shuffling improves BSD performance

in the case of highly spatially correlated noise.

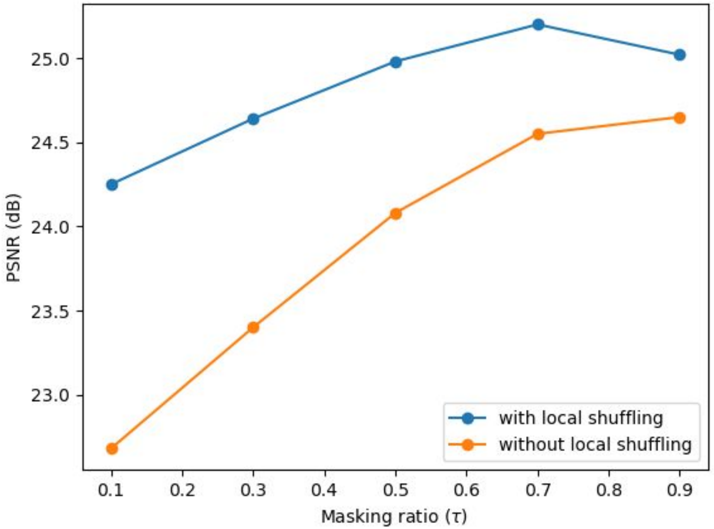

MASH

- Infer the noise correlation magnitude
- Define the optimal masking ratio based on (1.)
- Apply the local pixel shuffling or not based on (1.)
- Run the BSD with the optimal configuration

A Proxy for Noise Correlation Magnitude

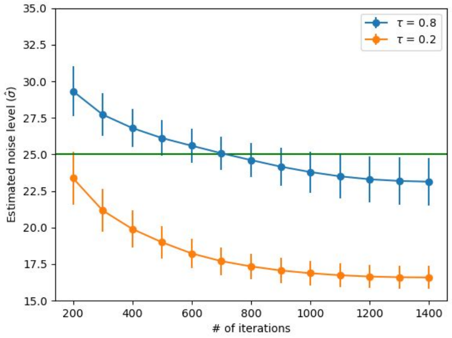

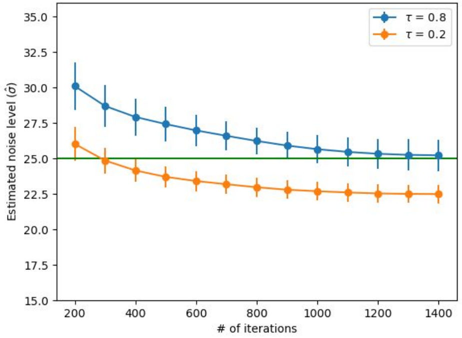

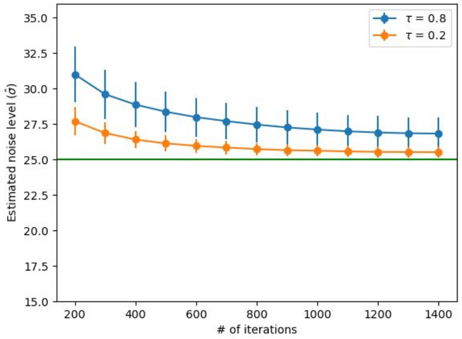

- Apply BSD with low and high masking ratios τ^{low} and τ^{high} respectively
- Compute the noise level gap $\varepsilon = |\hat{\sigma}_{\tau^{\text{high}}} - \hat{\sigma}_{\tau^{\text{low}}}|$
- $$\tau^{\text{optimal}} = \begin{cases} \tau^{\text{low}} & \text{if } \varepsilon \leq \varepsilon^{\text{low}}; \\ \tau^{\text{medium}} & \text{if } \varepsilon^{\text{low}} < \varepsilon < \varepsilon^{\text{high}}; \\ \tau^{\text{high}} & \text{if } \varepsilon^{\text{high}} < \varepsilon \end{cases}$$
- If $\varepsilon^{\text{high}} < \varepsilon$, apply the local pixel shuffling

Results

Method	SIDD Validation	SIDD Benchmarck	FMDD	PolyU
BM3D [2]	25.65/0.475	25.65/0.685	30.06/0.771	37.40/0.953
DIP[5]	32.11/0.740	-	32.90/0.854	37.17/0.912
Self2Self[4]	29.46/0.595	29.51/0.651	30.76/0.695	37.52/0.926
PD-denoising[7]	33.97/0.820	33.61/0.894	33.01/0.856	37.04/0.940
NN+denoiser[6]	-	33.18/0.895	32.21/0.831	37.66/0.956
APBSN-single[3]	30.90/0.818	30.71/0.869	28.43/0.804	29.61/0.897
ScoreDVI[1]	34.75/**0.856**	34.60/0.920	33.10/0.865	**37.77/0.959**
Baseline	33.12/0.805	32.67/0.850	32.25/0.824	37.12/0.911
Ours	**35.06**/0.851	**34.78**/0.900	**33.71/0.882**	37.62/0.932

(a) Quantitative comparison

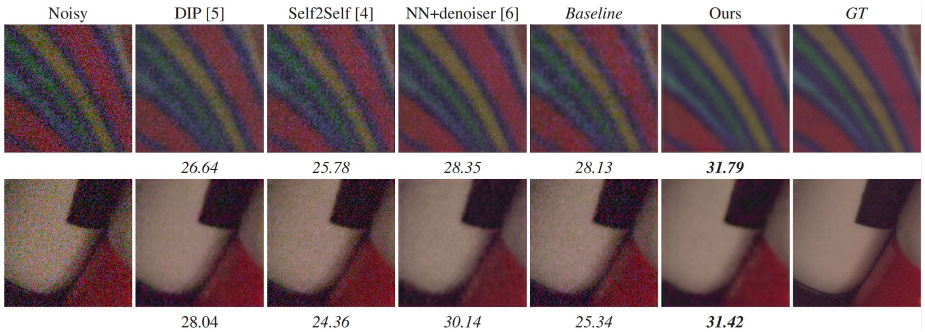

(b) Qualitative comparison

Ablations and Computational Cost

Adaptive masking ratio	Local pixel shuffling	SIDD	FMDD
No	No	33.12	32.25
No	Yes	33.86	32.92
Yes	No	34.45	33.56
Yes	Yes	35.06	33.71

(a) Ablations of MASH components

Method	Infer. time (s)	Params (M)	FLOPs (G)
DIP [5]	146.2	13.4	31.06
Self2Self [4]	3546.5	1.0	9.55
NN+denoiser [6]	897.6	13.4	31.06
APBSN-single [3]	121.4	3.66	234.63
ScoreDVI [1]	81.2	13.5	37.87
Baseline	24.6	0.99	11.44
Ours	75.3	0.99	11.44

(b) Efficiency comparison

References

[1] J. Cheng *et al*. Score priors guided deep variational inference for unsupervised real-world single image denoising. ICCV 2023

[2] K. Dabov *et al*. Bm3d image denoising with shape adaptive principal component analysis. SPARS 2009

[3] W. Lee *et al*. Apbsn: Self-supervised denoising for real-world images via asymmetric pd and blind-spot network. CVPR 2022

[4] Y. Quan *et al*. Self2self with dropout: Learning self-supervised denoising from single image. CVPR 2020

[5] D. Ulyanov *et al*. Deep image prior. CVPR 2018

[6] D. Zheng *et al*. An unsupervised deep learning approach for real-world image denoising. ICLR 2020

[7] Y. Zhou *et al*. When awgn-based denoiser meets real noises. AAAI 2020

[8] J. Li *et al*. Spatially Adaptive Self-Supervised Learning for Real-World Image Denoising. CVPR 2023

[9] A. Krull *et al*. Noise2Void - Learning Denoising from Single Noisy Images

Thank you!