Towards a unified view of unsupervised non-local methods for image denoising: the NL-Ridge approach

Sébastien Herbreteau, Charles Kervrann

SERPICO **Project-Team**
E-mail: sebastien.herbreteau@inria.fr, charles.kervrann@inria.fr
Github: https://github.com/sherbret/NL-Ridge/
Inria Rennes – Bretagne Atlantique
Campus universitaire de Beaulieu, 35042 Rennes Cedex France

Towards a unified view of unsupervised non-local methods for image denoising: the NL-Ridge approach

Sébastien Herbreteau and Charles Kervrann

Inria Centre Rennes - Bretagne Atlantique, UMR144 CNRS Institut Curie, PSL Research University, Sorbonne Université, France

Overview

We propose a **unified view** to reconcile **state-of-the-art unsupervised non-local denoisers**. We derive **NL-Ridge algorithm** which leverages local linear combinations of noisy similar patches.

- ▷ We show that non-local denoisers are characterized by the **family of functions** (f_Θ) used to process groups of similar noisy image patches.
- ▷ The optimal parameters Θ^* for each group are found by minimizing an approximation of the ℓ_2 **risk** via a two-step algorithm.
- ▷ NL-Ridge linearly combines noisy similar patches. Our closed-form aggregation weights are computed, in the second step, through **multivariate Ridge regressions**.
- ▷ NL-Ridge **outperforms** well established state-of-the-art unsupervised denoisers, including BM3D [1] and NL-Bayes [2], as well as recent unsupervised deep learning methods [5, 6, 7], while being simpler conceptually.

Parameter optimization

The optimal parameters Θ^* for the local denoiser f_Θ are found by minimizing the ℓ_2 **risk**:

$$\boxed{R_\Theta(X) = \mathbb{E}\|f_\Theta(Y) - X\|_F^2}.$$

As Θ^* requires the knowledge of X which is unknown, we propose the following two-step algorithm:

- Step 1: Minimize **Stein's unbiased risk estimate** (SURE) [3] to approximate Θ^*, reposition all the denoised patches and compute the $\hat{I}_1$ image aggregating estimators at each pixel.
- Step 2: Θ is improved via "**internal adaptation**" [4] with the pilot image $\hat{I}_1$.

General principle of non-local denoisers

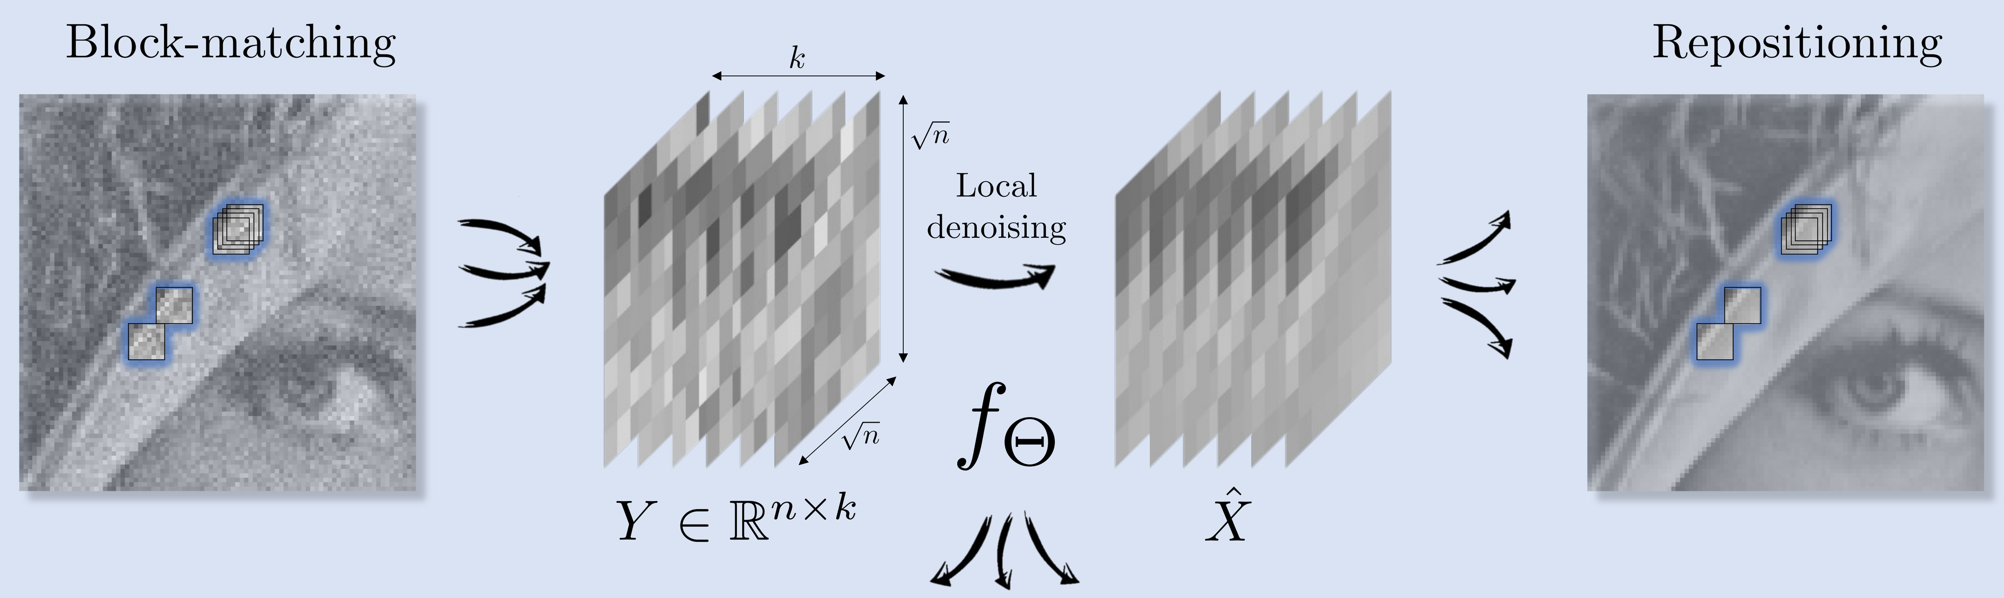

BM3D [1] assumes a locally sparse representation in a transform domain:

$$f_\Theta(Y) = P^\top(\Theta \circ (PYQ))Q^\top,$$

$P, Q \in \mathcal{O}(\mathbb{R})$, Hadamard product $\circ$.

NL-Bayes [2] was established in the Bayesian setting:

$$f_{\Theta,\beta}(Y) = \Theta Y + \beta u^\top,$$

$u \in \mathbb{R}^k$ is the all-ones vector.

NL-Ridge denoises each patch by linearly combining the most similar noisy patches:

$$\boxed{f_\Theta(Y) = Y\Theta}$$

NL-Ridge

Step 1: SURE

Proposition 1

An **unbiased estimate** *of the* **risk** $R_\Theta(X)$ *is* **Stein's unbiased risk estimate** $\mathrm{SURE}_\Theta(Y)$:

$$\begin{aligned} &- kn\sigma^2 + \|f_\Theta(Y) - Y\|_F^2 + 2\sigma^2 \operatorname{div} f_\Theta(Y) \\ &= -kn\sigma^2 + \|Y\Theta - Y\|_F^2 + 2n\sigma^2 \operatorname{tr}(\Theta) \end{aligned}$$

▷ Substituting this estimate for the risk $R_\Theta(X)$:

$$\arg\min_\Theta R_\Theta(X) \approx \arg\min_\Theta \mathrm{SURE}_\Theta(Y) = \hat{\Theta}_1$$

$$\boxed{\hat{\Theta}_1 = I_k - n\sigma^2(Y^\top Y)^{-1}}$$

Step 2: Internal adaptation

Proposition 2

The **quadratic risk** $R_\Theta(X)$ *is:*

$$R_\Theta(X) = \|X\Theta - X\|_F^2 + n\sigma^2\|\Theta\|_F^2$$

*minimized for (***multivariate Ridge regression***):*

$$\Theta^* = (X^\top X + n\sigma^2 I_k)^{-1} X^\top X$$

▷ Substituting $\hat{X}_1$ for X in the risk expression:

$$\arg\min_\Theta R_\Theta(X) \approx \arg\min_\Theta R_\Theta(\hat{X}_1) = \hat{\Theta}_2$$

$$\boxed{\hat{\Theta}_2 = (\hat{X}_1^\top \hat{X}_1 + n\sigma^2 I_k)^{-1} \hat{X}_1^\top \hat{X}_1}$$

Fig. 1: PSNR results on Barbara corrupted with additive white Gaussian noise ($\sigma = 20$).

Results

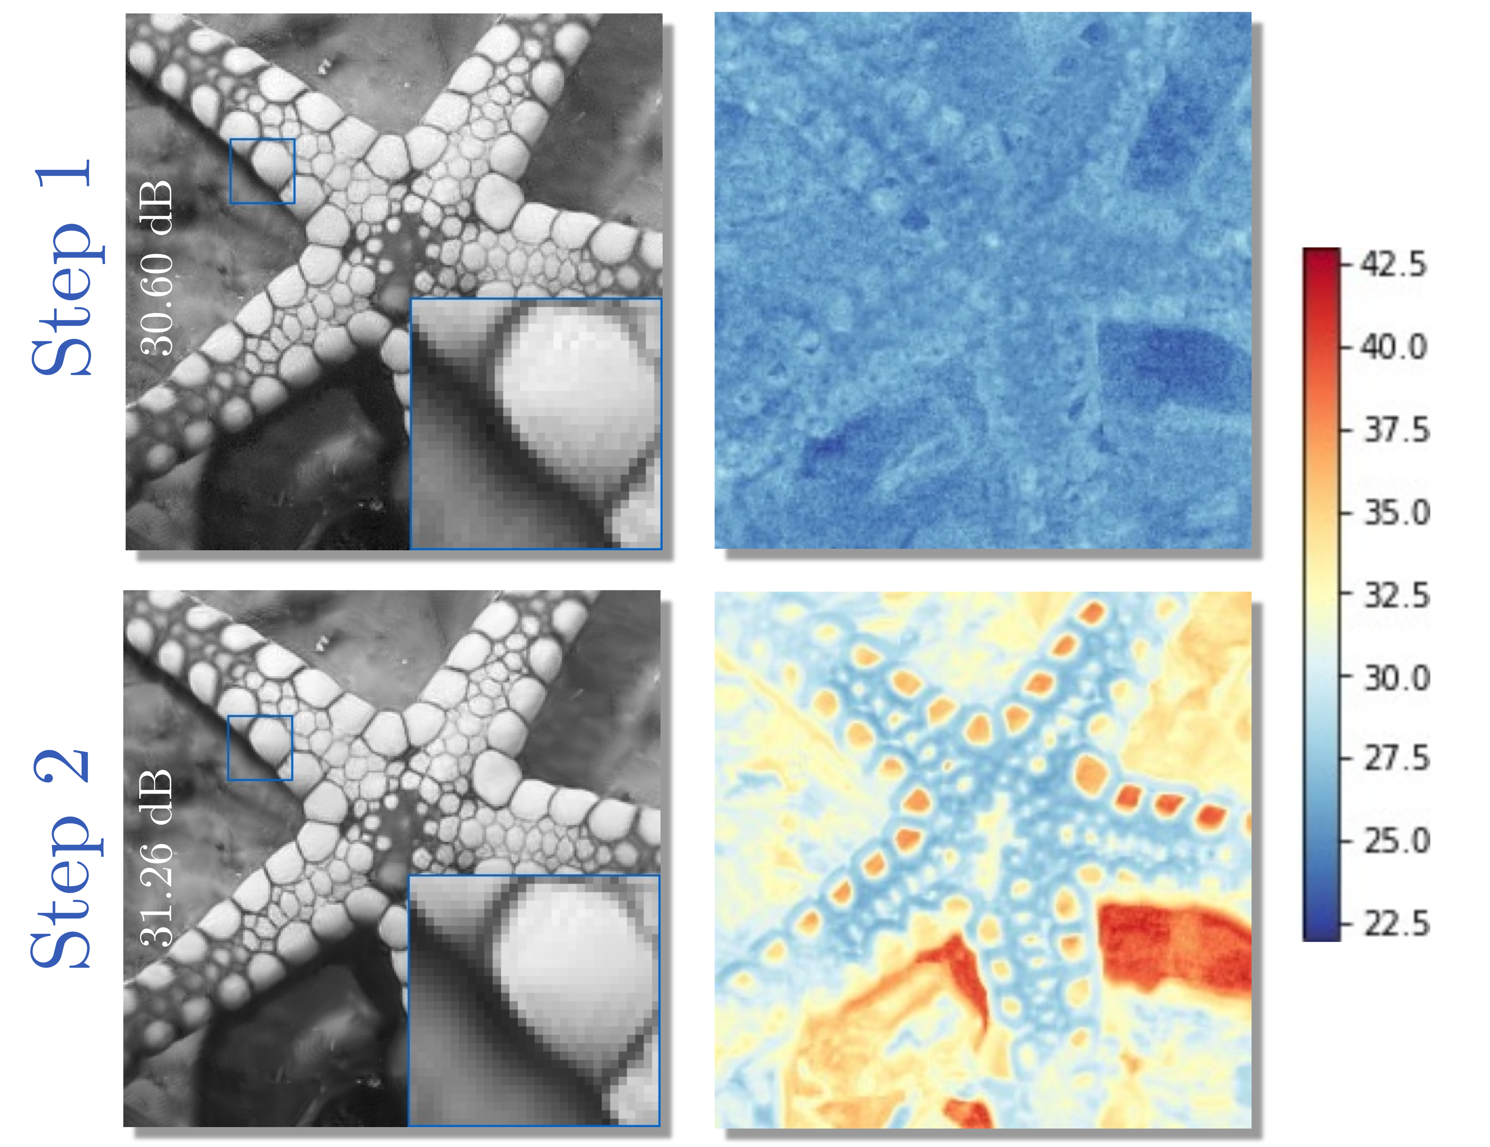

Fig. 2: Denoised images after each step ($\sigma = 15$) and PSNR colormaps (in dB) of denoised similarity matrices associated with each noisy patch.

		Methods	Set12	BSD68	Urban100
		Noisy	24.61 / 20.17 / 14.15	24.61 / 20.17 / 14.15	24.61 / 20.17 / 14.15
Unsupervised	Traditional 2-step	BM3D	32.37 / 29.97 / 26.72	31.07 / 28.57 / 25.62	32.35 / 29.70 / 25.95
		NL-Bayes	32.25 / 29.88 / 26.45	31.16 / 28.70 / 25.58	31.96 / 29.34 / 25.56
		NL-Ridge	32.46 / 30.00 / 26.73	31.20 / 28.67 / 25.67	32.53 / 29.90 / 26.29
	Traditional M-step	**GBNL-Ridge**	**32.71** / 30.24 / 26.86	**31.42** / **28.87** / 25.78	**33.00** / 30.37 / 26.49
		WNNM	32.70 / **30.26** / **27.05**	31.37 / 28.83 / **25.87**	32.97 / **30.39** / **26.83**
	Deep learning	DIP	30.12 / 27.54 / 24.67	28.83 / 26.59 / 24.13	- / - / -
		Noise2Self	31.01 / 28.64 / 25.30	29.46 / 27.72 / 24.77	- / - / -
		Self2Self	32.07 / 30.02 / 26.49	30.62 / 28.60 / 25.70	- / - / -
Supervised		DnCNN	**32.86** / **30.44** / 27.18	**31.73** / **29.23** / 26.23	32.68 / 29.97 / 26.28
		FFDnet	32.75 / 30.43 / **27.32**	31.63 / 29.19 / **26.29**	32.43 / 29.92 / **26.52**
		LIDIA	32.85 / 30.41 / 27.19	31.62 / 29.11 / 26.17	**32.80** / **30.12** / 26.51

Table. 1: PSNR (dB) results on various datasets corrupted with Gaussian noise ($\sigma = 15, 25, 50$). Best among each category is in bold. Best among each subcategory is underlined.

References

[1] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3D transform-domain collaborative filtering," in *IEEE Transactions on Image Processing*, vol. 16, no. 8, pp. 2080–2095, 2007.
[2] A. Buades, M. Lebrun, and J.-M. Morel, "A non-local bayesian image denoising algorithm," *SIAM Journal on Imaging Science*, vol. 6, no. 3, pp. 1665–1688, 2013.
[3] C. Stein, "Estimation of the mean of a multivariate normal distribution," *Annals of Statistics*, vol. 9, no. 6, pp. 1135–1151, 1981.
[4] G. Vaksman, M. Elad, and P. Milanfar, "LIDIA: Lightweight learned image denoising with instance adaptation," *Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2020.
[5] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Deep image prior," *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, p. 9446–9454, 2018.
[6] J. Batson and L. Royer, "Noise2self: Blind denoising by self-supervision," *Proceedings of the 36th International Conference on Machine Learning*, vol. 97, pp. 524–533, 2019.
[7] Y. Quan, M. Chen, T. Pang, and H. Ji, "Self2self with dropout: Learning self-supervised denoising from single image," *CVPR*, 2020.
[8] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a gaussian denoiser: residual learning of deep CNN for image denoising," in *IEEE Transactions on Image Processing*, vol. 26, no. 7, pp. 3142–3155, 2017.
[9] S. Gu, L. Zhang, W. Zuo, and X. Feng, "Weighted Nuclear Norm Minimization with Application to Image Denoising," in *IEEE Conference on Computer Vision and Pattern Recognition*, Columbus, OH, USA, 2014, pp. 2862-2869.