
On Nearest Neighbors in Non Local Means Denoising

I. Frosio
NVIDIA
USA
ifrosio@nvidia.com

J. Kautz
NVIDIA
USA
jkautz@nvidia.com

Abstract

To denoise a reference patch, the Non-Local-Means denoising filter processes a set of neighbor patches. Few Nearest Neighbors (NN) are used to limit the computational burden of the algorithm. Here we show analytically that the NN approach introduces a bias in the denoised patch, and we propose a different neighbors' collection criterion, named Statistical NN (SNN), to alleviate this issue. Our approach outperforms the traditional one in case of both white and colored noise: fewer SNNs generate images of higher quality, at a lower computational cost.

1 Introduction

Non-Local-Means (NLM) denoising has been widely investigated by researchers [2, 10, 3]. Denoising of a given patch is obtained as a weighted average of the surrounding patches, with weights proportional to the patch similarity. The filtering parameters such as the number and size of the patches affect both the quality of the output images [7], and the computational burden of the filter [15]. A widely used practice is to reduce the number of neighbors collected for each reference patch: the 3D Block-Matching (BM3D) filter achieves state-of-the-art results in this way [4]. The neighbors' set is collected through a Nearest-Neighbors (NN) approach. Reducing the set size leads to images with sharp edges [7], but it also introduces low-frequency artifacts (Fig. 1).

Our contribution is to show that this artifact occurs because the estimate of the noise-free patch from the set of NNs is biased. To the best of our knowledge, this is the first time this problem is explicitly investigated, although other authors (e.g. [10, 7, 18]) analyzed other sources of bias in NLM. We propose a strategy to collect neighbors, named Statistical NN (SNN), which reduces the prediction error of the estimate of the noise-free patch. Using fewer neighbors, SNN leads to an improvement in perceived image quality, both in case of white and colored Gaussian noise. In the latter case, visual inspection reveals that NLM with SNN achieves an image quality comparable to the state-of-the-art, at a much lower computational cost.

2 Related work

NLM denoising [2, 3] averages similar patches and then aggregates the averages. Given a noisy reference patch with P elements, $\boldsymbol{\mu}_r = [\mu_r^0 \mu_r^1 \dots \mu_r^{P-1}]$, its squared distance from a noisy neighbor patch $\boldsymbol{\gamma}_k$, is $\delta^2(\boldsymbol{\mu}_r, \boldsymbol{\gamma}_k) = \frac{1}{P} \cdot \sum_{i=0}^{P-1} (\mu_r^i - \gamma_k^i)^2$. Following [3], the weight of $\boldsymbol{\gamma}_k$ in the average is:

$$w_{\boldsymbol{\mu}_r, \boldsymbol{\gamma}_k} = \exp \{ \max[0, \delta^2(\boldsymbol{\mu}_r, \boldsymbol{\gamma}_k) - 2\sigma^2] \} / h^2, \quad (1)$$

where h is the *filtering parameter* and σ^2 is the variance of zero-mean, white Gaussian noise [2, 3]. The estimate of the noise-free patch, $\hat{\boldsymbol{\mu}}(\boldsymbol{\mu}_r)$, is then:

$$\hat{\boldsymbol{\mu}}(\boldsymbol{\mu}_r) = \sum_k w_{\boldsymbol{\mu}_r, \boldsymbol{\gamma}_k} \cdot \boldsymbol{\gamma}_k / \sum_k w_{\boldsymbol{\mu}_r, \boldsymbol{\gamma}_k}. \quad (2)$$

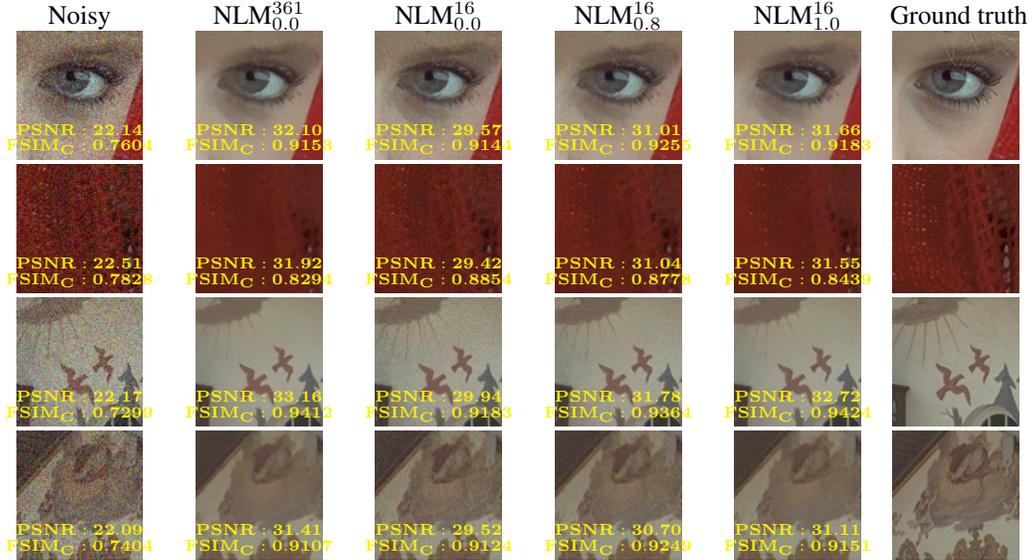


Figure 1: Noisy, 100×100 patches from the Kodak dataset, corrupted by zero-mean Gaussian noise, $\sigma = 20$. Traditional NLM ($NLM_{0,0}^{361}$) uses 361, 3×3 patches in a 21×21 search window; it denoises effectively the flat areas (the skin, the wall), but it blurs the small details (texture of the textile, fresco details). Using 16 NNs for each patch ($NLM_{0,0}^{16}$) improves the small details, but introduces colored noise in the flat areas. The proposed SNN technique ($NLM_{1,0}^{16}$) uses 16 neighbors and mimics the results of traditional NLM in flat areas (high PSNR), while keeping the visibility of small details (high FSIM_C). Best results are achieved when patches are collected through SNN, with $o = 0.8$ ($NLM_{0,8}^{16}$). Better seen at 400% zoom.

Using multiple scales in the distance function can improve patch matching [12, 11], leading to higher quality of the filtered images at the price of a higher computational cost. Wu et al. [18] identify a source of bias in $\hat{\mu}(\mu_r)$, due to the correlation among partially overlapping patches, and propose a weighting scheme to take this into account. Duval et al. [7] interpret NLM denoising as a bias / variance dilemma and show that $\hat{\mu}(\mu_r)$ is biased even in absence of noise; reducing the number of neighbors decreases the cost of the filter and the bias at the same time, but it also increases the variance on $\hat{\mu}(\mu_r)$, leaving some residual noise in the image. State-of-the art algorithms like BM3D [4], BM3D-SAPCA [5], and NL-Bayes [10] use a reduced set of neighbors, but they also require additional, computationally intensive processing steps and continue to suffer from visible artifacts on sharp edges and in smooth regions [9]. Machine learning can be used to learn to combine a set of patches for denoising, but the pros and cons of using a set of NN patches have not been discussed [1].

3 Method

None of the existing methods focus their attention on the fact that the neighbors' selection criterion, combined with the fact that the reference patch μ_r is noisy, affects the prediction error of $\hat{\mu}(\mu_r)$. In the *Additional Material*, we resort to a toy problem to compute the bias in $\hat{\mu}(\mu_r)$ introduced by the NN search strategy and we show analytically that collecting neighbors with the SNN approach can dramatically reduce it. For reasons of space, we describe here only the guiding principle behind SNN, and invite the curious reader to read the *Additional Material* for the math details.

Let's assume that a set of N neighbors, $\{\gamma_k\}_{k=0..N-1}$, has to be collected for a noisy, reference patch μ_r . Fig. 2 shows the case of a two-dimensional patch. When N neighbors are collected through the NN approach (left panel), the average of the set of neighbors is clearly biased toward μ_r . In NLM denoising, this drawback shows up as the *noise-to-noise* matching issue, *i.e.*, residual noise correlated with the reference patch is still present in the filtered image.

The recipe for collecting neighbors through the SNN approach is different and illustrated in the right panel. First of all, we notice that, in presence of white, Gaussian noise with standard deviation σ , the expected squared distance between the noisy, reference patch and a noisy neighbor patch

is $E[\delta^2(\boldsymbol{\mu}_r, \boldsymbol{\gamma}_k)] = 2\sigma^2$. The SNN neighbors are then defined as the patches $\{\boldsymbol{\gamma}_k\}_{k=1..N_n}$ that minimize:

$$|\delta^2(\boldsymbol{\mu}_r, \boldsymbol{\gamma}_k) - o \cdot 2\sigma^2|, \quad (3)$$

where where the *offset* parameter o allows to continuously move from the traditional NN approach ($o = 0$) to the SNN approach ($o = 1$). The right panel in Fig. 2 shows that the average of the set of SNNs is more likely to lie close to the noise-free patch $\boldsymbol{\mu}$. In fact, because of the statistical distribution of the noisy patches, the chance to collect a neighbor $\boldsymbol{\gamma}_k$ on the left side of $\boldsymbol{\mu}_r$ (and thus close to $\boldsymbol{\mu}_r$) is high, while neighbors on the right side of $\boldsymbol{\mu}_r$ are unlikely to occur. In practice, the SNN criterion looks for similar patches having orthogonal realizations of the noise. This minimizes the *noise-to-noise* matching issue and it leads to more effective noise cancellation, at the price of a slightly higher variance of $\hat{\boldsymbol{\mu}}(\boldsymbol{\mu}_r)$ (due to the fact that the neighbors are generally farther from $\boldsymbol{\mu}_r$).

4 Results

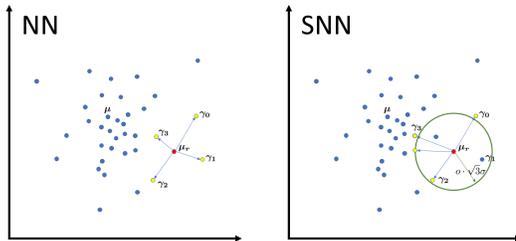


Figure 2: Left: collecting neighbors $\{\boldsymbol{\gamma}_k\}_{k=0..N-1}$ around the reference, noisy patch $\boldsymbol{\mu}_r$ through the traditional NN approach generates a neighbor set which is highly biased towards $\boldsymbol{\mu}_r$. The proposed SNN approach (right) collects neighbors close to the distance $o \cdot \sqrt{3}\sigma$ from $\boldsymbol{\mu}_r$, where σ is the noise standard deviation and o and additional *offset* parameter. The resulting set $\{\boldsymbol{\gamma}_k\}_{k=0..N-1}$ is less biased towards $\boldsymbol{\mu}_r$ and closer to the noise-free value $\boldsymbol{\mu}$.

For reasons of space, a detailed numerical evaluation of the effectiveness of the SNN schema for NLM denoising, obtained through a wide set of image quality metric (PSNR, SSIM[16], MSSSIM[17], GMSD[19], FSIM, and FSIM_C[21]), is reported only in the *Additional Material*. It is important noticing that these numerical evaluation confirms the superiority of the SNN approach, but we discuss here only the qualitative results.

The case of white, Gaussian noise is illustrated in Fig. 1. This shows a higher level of details when reducing the number of NN neighbors (from traditional NLM with 361 patches, $\text{NLM}_{0,0}^{361}$, to NLM with only 16 patches, $\text{NLM}_{1,0}^{16}$), coherently with the reduction of the NLM bias for a small set of neighbors, already explained in [7]. At the same time, the PSNR is decreasing, mostly because of the residual noise left in the flat areas (the skin of the girl, the wall with the

fresco in Fig. 1); this is how the *noise-to-noise* matching problem manifests in the NLM filtered image. Increasing the offset from $o = 0.0$ (NN) to $o = 1.0$ (SNN) increases the PSNR: SNN removes more noise in the flat areas (*e.g.*, the girl’s skin, the wall surface) compared to NN, thanks to the reduction of the bias introduced by the *noise-to-noise* matching issue. On the other hand, low-contrast edges (*e.g.*, the texture of the textile, the sun rays and the small details in the fresco) tend to be blurred by SNN. The best compromise between preserving low-contrasted details in the image and effectively smoothing the flat areas is obtained for $o = 0.8$, as shown in Fig. 1.

Having in mind denoising in practical situations, we also study the case of colored noise. In fact, images are generally acquired using a Bayer sensor. RGB data are obtained only after demosaicing, which introduces correlations between nearby pixels. In this case the NN approach is even more likely to collect false matches, amplifying low-frequency noise and demosaicing artifacts. Fig. 3 shows the comparison of state-of-the-art BM3D-CFA [6] to NLM with a NN or SNN approach in the case of an image corrupted by Gaussian noise ($\sigma = 20$) in the Bayer domain and then demosaiced. Visual inspection (Fig. 3) confirms that the NLM with NN continues to suffer from the *noise-to-noise* matching problem in the case of colored noise. NLM with SNN is comparable in terms of image quality to BM3D-CFA: our approach generates a slightly more noisy image, but with better preserved finer details (*e.g.*, see the eyelids in Fig. 3) and without introducing grid artifacts that are on the other hand visible for BM3D-CFA. Overall, our approach and BM3D-CFA result to be comparable in terms of image quality (as further demonstrated numerically in the *Additional Material*), but NLM with SNN has a much lower computational cost.

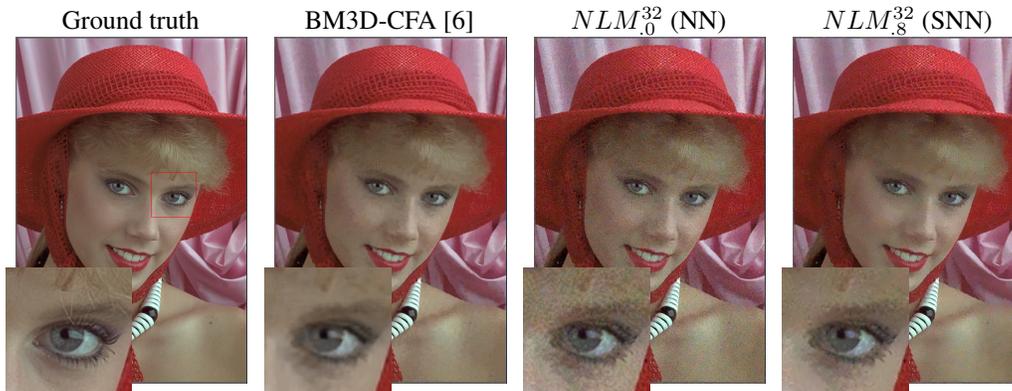


Figure 3: Comparison of different algorithms for colored noise denoising, $\sigma = 20$. The state-of-the-art BM3D-CFA [6] achieves the highest image quality metrics (see Fig. 5 in the Additional Material); a grid artifact is however visible in the white part of the eye. The traditional NLM algorithm, using $N_n = 32$ NNs (NLM_{0}^{32}), does not eliminate the colored noise introduced by demosaicing. Instead, when SNNs neighbors are used ($NLM_{.8}^{32}$), the colored noise is efficiently removed and quality at visual inspection is comparable to that of BM3D-CFA. In fact, our result appears less “splotchy” and a bit sharper. Better seen at 400% zoom.

Further numerical comparison against SDCT [20], BM3D [4] (in the case of white and colored noise), and BM3D-CFA [6] (in the case of colored noise) are reported in the *Additional Material*, where we also report and discuss the case of a “real” image.

5 Discussion

The NLM algorithm has been analyzed as a bias / variance dilemma: Duval et al. [7] showed that a reduction of the number of neighbors, reduces the bias in the estimate $\hat{\mu}_r$. The price to be paid is the introduction of the *noise-to-noise* matching issue, that shows up as splotchy, colored artifacts in the filtered images. In fact, since the NNs of any noisy reference patch μ_r lie close to it, their average is biased towards μ_r itself. We have shown here that this new source of bias is associated to the strategy adopted to collect the set of neighbors and demonstrated that the SNN approach can largely mitigate this bias both in theory and in practice.

When applied to NLM denoising and compared to the traditional NN strategy, the SNN approach produces images of quality similar to the original NLM in the flat areas, while keeping the visibility of the small details. In other words, SNN looks for neighbors whose noise realization is likely to be orthogonal to that of the reference patch. Averaging these patches is consequently more likely to effectively cancel out the noise.

The advantage of using the SNNs is even more evident in the case of colored noise, when *noise-to-noise* matching is more likely to occur. Visual inspection reveals that in this case NLM with SNN achieves an image quality comparable to the state-of-the-art BM3D-CFA [6], but at a lower computational cost (see Fig. 3). This case is indeed of practical importance, since denoising is performed after demosaicing in “real” images. Our experiments (see the *Additional Material*) suggest that NLM with SNN is close to BM3D-CFA also in the case of “real” images, where white balance, color correction, gamma correction and edge enhancement are applied after denoising, and they can significantly enhance the visibility of artifacts and residual noise.

It is worth mentioning that both the NN and SNN approaches to NLM denoising require the knowledge of the noise power, σ^2 , to properly set the filtering parameter h in Eq. (1) and guide the selection of the neighbors (in the case of SNN). In the case of “real” images, several effective noise estimation methods have been proposed [10], but it is also worthy mentioning that the noise distribution in real images is far more complicated than zero-mean Gaussian [8]. Even if a comprehensive analysis of the application of SNN to NLM denoising for real images goes beyond the scope of this paper, our approach can be effectively applied to “real” images through the application of a VST [8], and it achieves image quality far superior compared to traditional NLM and in practice comparable at visual inspection to the state-of-the-art BM3D-CFA approach [6], at a lower computational cost.

References

- [1] AHN, B., AND CHO, N. I. Block-matching convolutional neural network for image denoising. *CoRR abs/1704.00524* (2017).
- [2] BUADES, A., COLL, B., AND MOREL, J. M. A review of image denoising algorithms, with a new one. *Multiscale Model Sim.* (2005).
- [3] BUADES, A., COLL, B., AND MOREL, J.-M. Non-Local Means Denoising. *Image Processing On Line 1* (2011), 208–212.
- [4] DABOV, K., FOI, A., KATKOVNIK, V., AND EGIAZARIAN, K. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on Image Processing 16*, 8 (2007).
- [5] DABOV, K., FOI, A., KATKOVNIK, V., AND EGIAZARIAN, K. Bm3d image denoising with shape-adaptive principal component analysis. In *SPARS* (2009).
- [6] DANIELYAN, A., VEHVILAINEN, M., FOI, A., KATKOVNIK, V., AND EGIAZARIAN, K. Cross-color bm3d filtering of noisy raw data.
- [7] DUVAL, V., AUJOL, J.-F., AND GOUSSEAU, Y. On the parameter choice for the non-local means. *SIAM J. Imag. Sci.* (2011).
- [8] FOI, A. Clipped noisy images: Heteroskedastic modeling and practical denoising. *Signal Processing* (2009).
- [9] KNAUS, C., AND ZWICKER, M. Progressive image denoising. *IEEE TIP* (2014).
- [10] LEBRUN, M., COLOM, M., BUADES, A., AND MOREL, J. M. Secrets of image denoising cuisine. *Acta Numerica* (2012).
- [11] LOTAN, O., AND IRANI, M. Needle-match: Reliable patch matching under high uncertainty. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2016).
- [12] LOU, Y., FAVARO, P., SOATTO, S., AND BERTOZZI, A. Nonlocal similarity image filtering, 2009.
- [13] PONOMARENKO, N., JIN, L., IEREMEIEV, O., LUKIN, V., EGIAZARIAN, K., ASTOLA, J., VOZEL, B., CHEHDI, K., CARLI, M., BATTISTI, F., AND KUO, C.-C. J. Image database TID2013: Peculiarities, results and perspectives. *Signal Processing: Image Communication 30* (2015).
- [14] RICO MALVAR, LI-WEI HE, R. C. High-quality linear interpolation for demosaicing of bayer-patterned color images. In *International Conference of Acoustic, Speech and Signal Processing* (May 2004), Institute of Electrical and Electronics Engineers, Inc.
- [15] TSAI, Y.-T., STEINBERGER, M., PAJAK, D., AND PULLI, K. Fast ANN for high-quality collaborative filtering. *CGF* (2015).
- [16] WANG, Z., BOVIK, A. C., SHEIKH, H. R., AND SIMONCELLI, E. P. Image quality assessment: From error visibility to structural similarity. *IEEE TIP* (2004).
- [17] WANG, Z., SIMONCELLI, E., AND BOVIK, A. Multiscale structural similarity for image quality assessment. In *ACSSC* (2003).
- [18] WU, Y., TRACEY, B., NATARAJAN, P., AND NOONAN, J. P. Probabilistic non-local means. *IEEE Sig. Proc. Let. 20* (2013).
- [19] XUE, W., ZHANG, L., MOU, X., AND BOVIK, A. C. Gradient magnitude similarity deviation: A highly efficient perceptual image quality index. *IEEE TIP* (2014).
- [20] YU, G., AND SAPIRO, G. DCT Image Denoising: a Simple and Effective Image Denoising Algorithm. *IPOL* (2011).
- [21] ZHANG, L., ZHANG, D., MOU, X., AND ZHANG, D. Fsim: A feature similarity index for image quality assessment. *IEEE TIP* (2011).

On Nearest Neighbors in Non Local Means Denoising: Additional Material

1 Prediction error of NN, SNN

In this section we resort to a toy problem to compute analytically the bias and variance of the estimate of a noise-free patch, in case neighbor patches are collected through the traditional NN approach or through the proposed SNN method. We also demonstrate numerically the advantage of SNN over NN for the toy problem in hand.

1.1 Prediction error of NN

NLM denoising is a three step procedure: i) for each patch, the neighbors are identified; ii) neighbors are averaged through a weighted average procedure; iii) denoised patches are aggregated (since patches are partially overlapping, multiple estimates are averaged for each pixel). Collecting NN neighbors introduces a prediction error in step ii). To demonstrate this, we consider a 1×1 patch with noise-free value μ , and assume Gaussian noise¹ with variance σ^2 . We assume that N noisy replicas of the patch are available (black samples in Fig. 4a). For each reference patch, μ_r , we collect N_n neighbors, $\{\gamma_k\}_{k=1..N_n}$; γ_k and μ_r are Gaussian random variables with average μ and variance σ^2 , indicated by $G(\mu, \sigma^2)$, with pdf $1/\sigma \cdot \phi[(x - \mu)/\sigma]$, where $\phi(x) = 1/\sqrt{2\pi} \cdot e^{-0.5 \cdot x^2}$. We define $d = d(\mu_r)$ the range in which we are likely to find the N_n nearest neighbors of μ_r (red stars in Fig. 4a). More specifically, $d(\mu_r)$ satisfies:

$$\frac{1}{\sigma} \int_{\mu_r - d(\mu_r)}^{\mu_r + d(\mu_r)} \phi\left(\frac{x - \mu}{\sigma}\right) dx = \Phi\left[\frac{\mu_r + d(\mu_r) - \mu}{\sigma}\right] - \Phi\left[\frac{\mu_r - d(\mu_r) - \mu}{\sigma}\right] = \frac{N_n}{N}, \quad (4)$$

where $\Phi(x) = 1/2 \cdot [1 + \text{erf}(x/\sqrt{2})]$, and $\Phi[(x - \mu)/\sigma]$ is the cdf of a Gaussian random variable with mean μ and variance σ^2 . A closed form solution for $d(\mu_r)$ in Eq. (4) does not exist, but $d(\mu_r)$ can be computed numerically, for instance, using a bracketing technique (see Fig. 4c). From $d(\mu_r)$, we compute the expected value of the average of the N_n nearest neighbors around μ_r , which is $E[\hat{\mu}(\mu_r)]$. This is equivalent to step ii) of NLM, neglecting the weights in the weighted average. Since $\hat{\mu}(\mu_r)$ is a truncated Gaussian variable, bounded by $\mu_r \pm d(\mu_r)$, its expected value and variance are:

$$\alpha = \frac{[\mu_r - d(\mu_r)]}{\sigma}, \quad \beta = \frac{[\mu_r + d(\mu_r)]}{\sigma}$$

$$E[\hat{\mu}(\mu_r)] = \mu - \sigma \cdot \frac{\phi(\beta) - \phi(\alpha)}{\Phi(\beta) - \Phi(\alpha)} \quad (5)$$

$$\text{Var}[\hat{\mu}(\mu_r)] = \sigma^2 \cdot \left\{ 1 - \frac{\beta \cdot \phi(\beta) - \alpha \cdot \phi(\alpha)}{\Phi(\beta) - \Phi(\alpha)} + \left[\frac{\phi(\beta) - \phi(\alpha)}{\Phi(\beta) - \Phi(\alpha)} \right]^2 \right\} / N_n. \quad (6)$$

Fig. 4d shows $E[\hat{\mu}(\mu_r)]$ and $\text{Std}[\hat{\mu}(\mu_r)]$ for the case $\mu = 1$, $\sigma = 0.2$, when $N_n = 16$ neighbors are to be collected from a total of $N = 100$ samples. Since the NN neighbors $\{\gamma_k\}_{k=1..N_n}$ lie close to the noisy reference patch, μ_r , each estimate $\hat{\mu}(\mu_r)$ is biased towards μ_r . The bias grows almost linearly with $\mu - \mu_r$ and it saturates at approximately 5σ of distance from μ , as the set of neighbors becomes stable (since N is finite, the same 16 samples are found in the tail of the distribution). We compute

¹Our reasoning for a single-pixel patch easily generalizes to a larger patch, where each pixel is corrupted by zero-mean Gaussian noise.

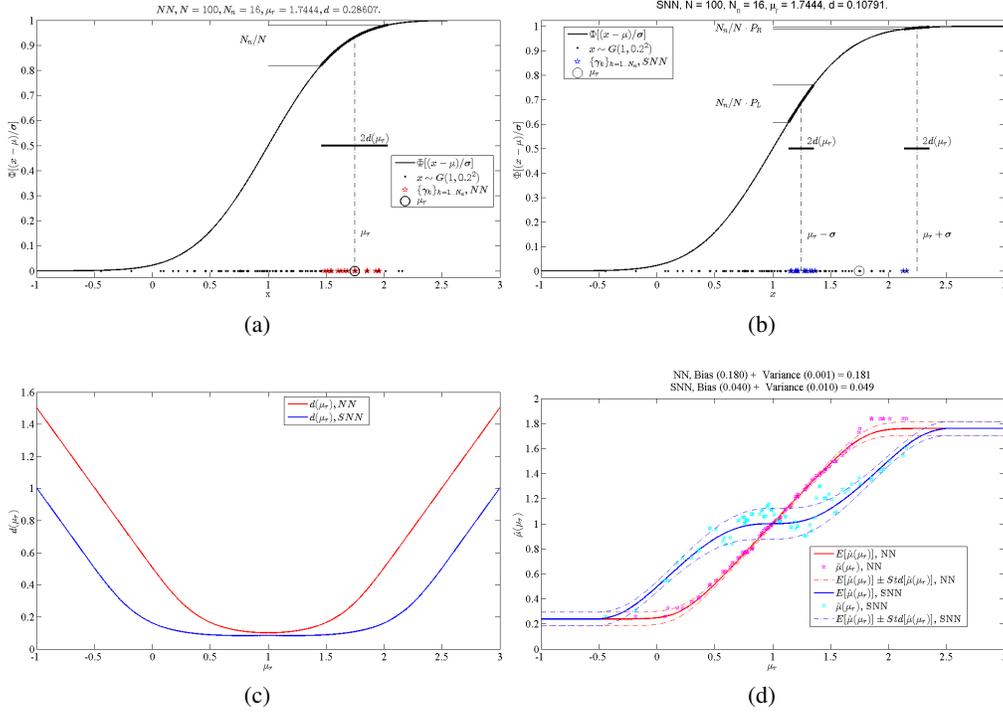


Figure 4: Panels (a) and (b) show respectively the NN and SNN strategy to collect $N_n = 16$ neighbors $\{\gamma_k\}_{k=1..N_n}$ of μ_r , for a 1×1 patch with noise-free value $\mu = 1$, corrupted by zero-mean Gaussian noise, $\sigma = 0.2$, from a set of $N = 100$ samples. In panels (a) and (b), $\mu_r = 1.7444$, far off $\mu = 1$. Its 16 NNs are in its immediate vicinity (red, in a), ultimately leading to a bad estimate. The 16 SNNs (blue, in b) are on average closer to the actual μ leading to a better estimate. Panel (c) shows the neighbor search intervals; $d(\mu_r)$ is generally smaller for SNN, since the search for neighbors occurs in this case in two ranges, on the left and right of μ_r . Panel (d) shows the expected value of the estimate, $E[\hat{\mu}(\mu_r)]$, and its standard deviation, $\text{Std}[\hat{\mu}(\mu_r)]$, as a function of μ_r ; $\hat{\mu}(\mu_r)$ are specific estimates of μ obtained from the black samples x in panels (a) and (b). SNN yields values closer to the actual $\mu = 1$, even when the noisy patch μ_r is far off the actual noise-free value.

the expected prediction error for estimating μ through $\hat{\mu}(\mu_r)$ by integrating the bias and variance terms as:

$$\begin{aligned}
 \text{MSE} &= \int_{-\infty}^{+\infty} \{ [E[\hat{\mu}(\mu_r)] - \mu]^2 + \text{Var}[\hat{\mu}(\mu_r)] \} \cdot \\
 &\quad \frac{1}{\sigma} \cdot \phi\left[\frac{\mu_r - \mu}{\sigma}\right] d\mu_r, \tag{7}
 \end{aligned}$$

and it is equal to 0.180 (bias) + 0.001 (variance) = 0.181 for the case in Fig. 4a. In practice, when NN neighbors are collected, samples with correlated noise are likely to be chosen. We call this the *noise-to-noise* matching problem: averaging NNs patches may amplify small correlations between the noise, without canceling it. In the context of image denoising, the *noise-to-noise* matching problems shows up as residual, colored noise in the filtered image (see the images filtered with $\text{NLM}_{0.0}^{16}$ in the main paper).

1.2 Statistics of patch distance

Before illustrating the SNN method to collect the neighbors of a reference patch, we analyze the statistical distribution of the distance between multi-dimension patches (similarly to [18]). This shows that, in presence of noise, the expected distance between the reference patch and its neighbors is not zero, which is the main observation behind the proposed SNN approach.

The search for similar patches is performed by computing the distance between μ_r and patches of the same size, but in different positions in the image. If, apart from noise, the reference patch μ_r and

its neighbor γ_k are two replica of the same patch, we get:

$$\delta^2(\boldsymbol{\mu}_r, \boldsymbol{\gamma}_k) = (2\sigma^2/P) \cdot \sum_{i=0}^{P-1} G(0, 1)^2. \quad (8)$$

Since the sum of P squared normal variables has a χ_P^2 distribution with P degrees of freedom, we have $\delta^2(\boldsymbol{\mu}_r, \boldsymbol{\gamma}_k) \sim (2\sigma^2/P) \cdot \chi_P^2$, and therefore:

$$E[\delta^2(\boldsymbol{\mu}_r, \boldsymbol{\gamma}_k)] = 2\sigma^2. \quad (9)$$

Thus, for two noisy replicas of the same patch, the expected squared distance is not zero. This has already been noticed and effectively employed since the original NLM paper [2] to compute the weights of the patches in the weighted average (see Eq. (1) in the main paper), giving less importance to patches at a squared distance larger than $2\sigma^2$, or to build a better weighting scheme, as in [18]. Nonetheless, to the best of our knowledge, it has never been employed as a driver for the selection of the neighbor patches, as we do here.

We can also establish a connection between the toy problem at hand and matching patches in general. In fact, for our toy problem, we can apply Fischer's approximation ($\sqrt{2\chi_P^2} \cong G(\sqrt{2P-1}, 1)$) to a single-pixel patch ($P=1$), and get:

$$\delta(\boldsymbol{\mu}_r, \boldsymbol{\gamma}) \cong \sigma \cdot \sqrt{2P-1} + G(0, \sigma^2) = \sigma + G(0, \sigma^2), \quad (10)$$

showing that the expected distance between two 1×1 noisy patches is approximately σ .

1.3 SNN

As an alternative to NN, and mostly inspired by Eq. (9), we propose collecting neighbors whose squared distance from the reference patch is instead close to its expectation. Thus SNNs are the patches $\{\boldsymbol{\gamma}_k\}_{k=1..N_n}$ that minimize:

$$|\delta^2(\boldsymbol{\mu}_r, \boldsymbol{\gamma}_k) - o \cdot 2\sigma^2|, \quad (11)$$

where we have introduced an additional offset parameter o , that allows to continuously move from the traditional NN approach ($o=0$) to the SNN approach ($o=1$). We assume $o=1$ for now.

1.4 Prediction error of SNN

To compare the prediction error of NN and SNN in our toy problem, we compute $E[\hat{\mu}(\mu_r)]$ and $\text{Var}[\hat{\mu}(\mu_r)]$, when SNN neighbors of μ_r are collected, *i.e.*, when samples are close to $\mu_r \pm o \cdot \sigma^2$. Fig. 4b illustrates the sampling of SNN neighbors. Sampling potentially occurs on both sides of μ_r , with a different chance of collecting neighbors on each side. The interval $d(\mu_r)$ where we expect to find the neighbors, has now to satisfy:

$$\begin{aligned} & \Phi\{[\mu_r - o \cdot \sigma + d(\mu_r) - \mu]/\sigma\} + \\ & - \Phi\{[\mu_r - o \cdot \sigma - d(\mu_r) - \mu]/\sigma\} + \\ & + \Phi\{[\mu_r + o \cdot \sigma + d(\mu_r) - \mu]/\sigma\} + \\ & - \Phi\{[\mu_r + o \cdot \sigma - d(\mu_r) - \mu]/\sigma\} = N_n/N, \end{aligned} \quad (12)$$

where the first two rows represent the probability to sample one neighbor in the left interval (P_L in Fig. 4b), whereas the term in the third and fourth row is for the right interval (P_R in Fig. 4b). Since $P_L > P_R$ here, the estimate $\hat{\mu}(\mu_r)$ will be likely moved towards the left, decreasing the prediction error. As for the NN case, $d(\mu_r)$ can be estimated through Eq. (12) using a bracketing technique. Fig. 4c shows the search intervals $d(\mu_r)$ for a Gaussian random variable, $\mu=1$, $\sigma=0.2$, when $N_n=16$ and $N=100$. Compared to the NN search range, the SNN one is smaller as it generally includes two search intervals. These two search ranges may collapse into one when μ_r belongs to the tails of the distribution, providing the same result as NN in this (rare) case.

²Notice that in our toy problem we collect the samples based on their expected distance from the noisy reference μ_r , computed through Eq. (10), while in SNN we use the squared distance. This simplification makes the toy problem mathematically manageable, without changing the intuition behind SNN.

The expected value of the average of the SNNs is the expected value of the mix of two truncated Gaussian random variables, on the left and right side of μ_r . After a few simplifications and using Eqs. (5 - 6) to compute the expected values (E_L, E_R) and variances ($\text{Var}_L, \text{Var}_R$) on the left and right side of μ_r , we have:

$$\begin{aligned}
\alpha_L &= [\mu_r - o \cdot \sigma - d(\mu_r) - \mu]/\sigma, \\
\beta_L &= [\mu_r - o \cdot \sigma + d(\mu_r) - \mu]/\sigma \\
\alpha_R &= [\mu_r + o \cdot \sigma - d(\mu_r) - \mu]/\sigma, \\
\beta_R &= [\mu_r + o \cdot \sigma + d(\mu_r) - \mu]/\sigma \\
P_L &= \Phi(\beta_L) - \Phi(\alpha_L), \quad P_R = \Phi(\beta_R) - \Phi(\alpha_R) \\
E[\hat{\mu}(\mu_r)] &= (P_L E_L + P_R E_R)/(P_L + P_R) \tag{13} \\
\text{Var}[\hat{\mu}(\mu_r)] &= \left\{ \frac{P_L(E_L^2 + \text{Var}_L) + P_R(E_R^2 + \text{Var}_R)}{P_L + P_R} + \right. \\
&\quad \left. - E[\hat{\mu}(\mu_r)] \right\} / N_n. \tag{14}
\end{aligned}$$

Eqs. (12), (13), and (14) boil down to the simple NN case in the case of overlapping intervals. Fig. 4d shows $E[\hat{\mu}(\mu_r)]$ and $\text{Std}[\hat{\mu}(\mu_r)]$ for the SNN case. The overall bias and variance for SNN are computed by integration as in Eq. (7), and they are respectively equal to 0.040 and 0.010, for a total MSE of 0.49. Compared to NN, SNN slightly increases the variance of $\hat{\mu}(\mu_r)$, but it drastically decreases the bias of the estimate, especially for those points close to μ (which are, by the way, the most frequent). In practice, the SNN criterion looks for similar patches having orthogonal realizations of the noise. This minimizes the *noise-to-noise* matching issue and it leads to more effective noise cancellation, at the price of a slightly higher variance of $\hat{\mu}(\mu_r)$.

2 Quantitative evaluation of the SNN approach for NLM denoising

We report here the extensive quantitative evaluation of the results achieved with the proposed SNN approach applied to NLM denoising, in the case of white Gaussian noise, colored noise, and in the case of a “real” image. Such evaluation has not been reported in the main paper for reasons of space.

2.1 White Gaussian noise

We test the effectiveness of the SNN schema on the Kodak image dataset [13]. We first consider denoising in case of white, zero-mean, Gaussian noise with standard deviation $\sigma = \{5, 10, 20, 30, 40\}$, and evaluate the effect of the offset parameter o and number of neighbors N_n . With $\text{NLM}_o^{N_n}$ we indicate NLM with N_n neighbors and an offset o . Each image is processed with NLM_o^{16} and for o ranging from $o = 0$ (*i.e.*, the traditional NN strategy) to $o = 1$; NLM_0^{361} and NLM_0^{900} indicate the original NLM, using all the patches in the search window, respectively for $\sigma < 30$ and $\sigma \geq 30$. The patch size, the number of neighbors N_n , and the filtering parameter h are the optimal ones in [2].

We resort to a large set of image quality metrics to evaluate the filtered images: PSNR, SSIM[16], MSSSIM[17], GMSD[19], FSIM and FSIM_C [21]. PSNR is proportional to the average squared error in the image; SSIM and MSSSIM take inspiration from the the human visual system, giving less importance to the residual noise close to edges. GMSD is a perceptually-inspired metric that compares the gradients of the filtered and the reference images. FSIM is the one that correlates better with the human judgement of the image quality, and takes into consideration the structural features of the images; FSIM_C also takes into account the color component.

Table 1 shows the average image quality metrics measured on the Kodak dataset, whereas visual inspection of Fig. 1 in the main paper reveals the correlation between these metrics and the artifacts introduced in the filtered images. The PSNR decreases when reducing the number of NN neighbors (see the first two rows for each noise level in Table 1), mostly because of the residual, colored noise left by NLM_0^{16} in the flat areas (like the skin of the girl and the wall with the fresco in Fig. 1 in the main paper); SSIM and MSSSIM show a similar trend, while GMSD, FSIM, and FSIM_C improve when the number of NN neighbors is reduced to 16. The improvement of these three perceptually-based image quality metrics is associated with the reduction of the NLM bias in the case of a small set of neighbors, clearly explained in [7]. At visual inspection, edges and structures are better preserved, (see the textile pattern and the sun rays in Fig. 1 in the main paper). Table 1 shows

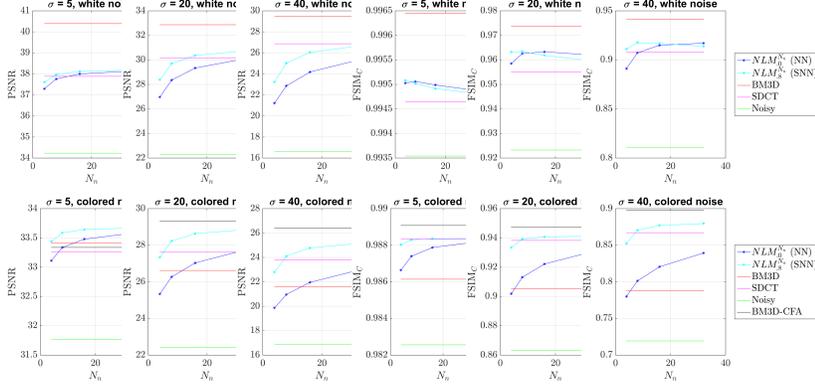


Figure 5: Average PSNR and FSIM_C measured on the Kodak dataset, as a function of the number of neighbors N_n , for NLM using the NN ($NLM_0^{N_n}$) and SNN ($NLM_{.8}^{N_n}$) approaches, compared to SDCT [20], BM3D [4], and BM3D-CFA [6] in the case of additive white Gaussian noise (first two rows) and colored noise (bottom rows). Notice that the number of patches used by BM3D and BM3D-CFA is constant and coherent with the optimal implementations described in [4] and [6].

the PSNR consistently increasing with the offset o , coherently with the decrease of the prediction error measured in our toy problem from $o = 0.0$ (NN) to $o = 1.0$ (SNN). Visual inspection (Fig. 1 in the main paper) confirms that SNN removes more noise in the flat areas (*e.g.*, the girl’s skin, the wall surface) compared to NN. For $o = 1.0$ the PSNR approaches that achieved by traditional NLM, which requires more neighbor patches and therefore a higher computational time. Even more interestingly, SSIM and MSSIM are maximized for $o = 0.9$, whereas the best GMSD, FSIM and FSIM_C are obtained for $o = 0.65$ or $o = 0.8$. Even if none of these metrics is capable of perfectly measuring the quality perceived by a human observer, they consistently suggest that filtering with the SNN approach generates images of superior, perceptual quality compared to the traditional NLM filtering. To summarize, sharp edges are reconstructed well by all the approaches. Low-contrast edges (*e.g.*, the texture of the textile, the sun rays and the small details in the fresco), are well preserved by the NN approach, which achieves better GMSD, FSIM and FSIM_C compared to traditional NLM, but $NLM_0^{1.6}$ also leaves higher, colored noise in the image (low PSNR, SSIM and MSSIM). These same low-contrast edges are oversmoothed by $NLM_{0.0}^{361}$ and by SNN when $o = 1$ ($NLM_{1.0}^{361}$), whereas SNN with an offset $o = 0.8$ ($NLM_{0.8}^{361}$) achieves the best compromise between preserving low-contrasted details in the image and effectively smoothing the flat areas, coherently with the numerical evaluation in Table 1. The proposed SNN strategy always outperforms NN (when using the same number of patches) in the image quality metrics, with the exception of very low noise ($\sigma = 5$), where GMSD, FSIM and FSIM_C are constant for $.0 < o < 0.65$. Quite reasonably, the advantage of SNN over NN is larger for a large σ : notice in fact that, for $\sigma \rightarrow 0$, NN and SNN converge to the same algorithm.

The first row of Fig. 5 shows the average PSNR and FSIM_C achieved by NLM using NN ($o = 0$) and SNN ($o = 0.8$) on the Kodak dataset, as a function of the number of neighbors N_n . Notice that the filter computational cost scales linearly (apart from the neighbors search part) with N_n . Compared to NN, the SNN approach always achieves a higher PSNR, since it smooths better the flat areas, as already shown in Fig. 1 in the main paper. FSIM_C is slightly lower for NN when $N_n \geq 16$, because SNN tends to oversmooth the edges, but when few neighbors are used the SNN approach clearly outperforms NN. The state-of-the-art denoising filter BM3D [4] produces images of superior quality, but at a much higher computational cost. We also compared the image quality achieved by our approach with the Sliding Discrete Cosine Transform filter (SDCT [20]), a denoising algorithm of comparable computational complexity based on signal sparsification and not using a NL approach. SDCT achieves a comparable PSNR, but generally a lower FSIM_C.

2.2 Colored noise

Although the case of white Gaussian noise is widely studied, it represents an ideal situation which rarely occurs in practice. Having in mind the practical application of denoising, we also study the more general case of colored noise. In fact, images are generally acquired using a Bayer sensor.

	PSNR	SSIM	MSSSIM	GMSD	FSIM	FSIM _C	σ
noisy	34.15	.9625	.9797	.0211	.9937	.9931	5
NLM ₀ ³⁶¹	38.34	.9851	.9921	.0130	.9941	.9939	
NLM ₀ ¹⁶	38.08	.9847	.9919	.0113	.9949	.9947	
NLM ₁ ¹⁶	38.08	.9847	.9919	.0113	.9949	.9947	
NLM _{.35} ¹⁶	38.08	.9847	.9919	.0113	.9949	.9947	
NLM _{.65} ¹⁶	38.14	.9850	.9920	.0113	.9949	.9947	
NLM _{.8} ¹⁶	38.23	.9854	.9922	.0115	.9948	.9946	
NLM _{.9} ¹⁶	38.28	.9854	.9923	.0118	.9946	.9945	
NLM _{1.0} ¹⁶	38.29	.9852	.9921	.0124	.9944	.9942	
noisy	28.13	.8744	.9332	.0667	.9763	.9743	10
NLM ₀ ³⁶¹	34.84	.9634	.9804	.0338	.9814	.9811	
NLM ₀ ¹⁶	33.96	.9580	.9773	.0269	.9861	.9856	
NLM ₁ ¹⁶	33.96	.9580	.9773	.0269	.9861	.9856	
NLM _{.35} ¹⁶	33.96	.9580	.9773	.0269	.9861	.9856	
NLM _{.65} ¹⁶	34.19	.9610	.9789	.0266	.9861	.9857	
NLM _{.8} ¹⁶	34.51	.9639	.9805	.0274	.9856	.9852	
NLM _{.9} ¹⁶	34.66	.9645	.9809	.0291	.9845	.9841	
NLM _{1.0} ¹⁶	34.71	.9635	.9803	.0319	.9828	.9824	
noisy	22.11	.6810	.8236	.1549	.9250	.9174	20
NLM ₀ ³⁶¹	31.18	.9109	.9505	.0749	.9475	.9468	
NLM ₀ ¹⁶	29.21	.8802	.9343	.0607	.9646	.9633	
NLM ₁ ¹⁶	29.21	.8802	.9343	.0607	.9646	.9633	
NLM _{.35} ¹⁶	29.21	.8802	.9343	.0607	.9646	.9633	
NLM _{.65} ¹⁶	29.75	.8948	.9420	.0572	.9653	.9641	
NLM _{.8} ¹⁶	30.45	.9086	.9493	.0583	.9631	.9621	
NLM _{.9} ¹⁶	30.81	.9119	.9511	.0635	.9585	.9576	
NLM _{1.0} ¹⁶	30.93	.9084	.9490	.0719	.9510	.9503	
noisy	18.59	.5331	.7243	.2138	.8687	.8536	30
NLM ₀ ⁹⁰⁰	29.11	.8631	.9197	.1071	.9154	.9144	
NLM ₀ ¹⁶	27.40	.8273	.9035	.0827	.9438	.9419	
NLM ₁ ¹⁶	27.40	.8273	.9035	.0827	.9438	.9419	
NLM _{.35} ¹⁶	27.40	.8273	.9035	.0827	.9438	.9419	
NLM _{.65} ¹⁶	27.48	.8311	.9055	.0817	.9441	.9422	
NLM _{.8} ¹⁶	28.15	.8543	.9175	.0793	.9439	.9422	
NLM _{.9} ¹⁶	28.69	.8665	.9233	.0850	.9382	.9368	
NLM _{1.0} ¹⁶	28.91	.8640	.9201	.0988	.9247	.9235	
noisy	16.09	.4280	.6414	.2501	.8157	.7937	40
NLM ₀ ⁹⁰⁰	27.71	.8184	.8893	.1303	.8856	.8844	
NLM ₀ ¹⁶	25.38	.7522	.8591	.1061	.9237	.9206	
NLM _{0.1} ¹⁶	25.38	.7522	.8591	.1061	.9237	.9206	
NLM _{0.2} ¹⁶	25.38	.7522	.8591	.1061	.9237	.9206	
NLM _{0.35} ¹⁶	25.38	.7522	.8591	.1061	.9237	.9206	
NLM _{0.5} ¹⁶	25.38	.7522	.8591	.1061	.9237	.9206	
NLM _{0.65} ¹⁶	25.49	.7580	.8624	.1043	.9243	.9213	
NLM _{0.8} ¹⁶	26.38	.7942	.8813	.0977	.9248	.9223	
NLM _{0.9} ¹⁶	27.11	.8153	.8910	.1037	.9168	.9148	
NLM _{1.0} ¹⁶	27.42	.8147	.8870	.1214	.8971	.8954	

Table 1: Average image quality index for denoising the Kodak image dataset, corrupted by zero-mean Gaussian noise, different standard deviation σ and different values of the offset parameter o . Bold numbers indicate the best results.

RGB data are obtained only after demosaicing, which introduces correlations between nearby pixels. A direct consequence of this is that the NN approach is even more likely to collect false matches, amplifying low-frequency noise as well as demosaicing artifacts. To study this, we mosaic the noisy images of the Kodak dataset, add noise in the Bayer domain, and subsequently demosaic them through the widely used Malvar algorithm [14]. Denoising is then performed through $NLM_{0.8}^{N_n}$, $NLM_{0.8}^{N_n}$ (for $N_n = \{4, 8, 16, 32\}$), BM3D [4], and SDCT [20]. All these filters require an estimate of the noise standard deviation in the RGB domain, obtained here considering that RGB data come from a linear combination (as in [14]) of independent samples in the Bayer domain. To compare with the state-of-the-art, we also filter the images with BM3D-CFA [6], a modified version of BM3D specifically tailored to denoise data in the Bayer domain.

The bottom rows of Fig. 5 show the average PSNR and $FSIM_C$ on the Kodak dataset. The proposed SNN approach is always more effective than NN in removing colored noise, both in terms of PSNR and $FSIM_C$. This is clearly related with the higher occurrence of *noise-to-noise* matching for NN in the case of colored noise. Visual inspection (Fig. 3 in the main paper) further confirms that the NN approach applied to NLM suffers from the *noise-to-noise* matching problem. NLM with SNN also generally outperforms SDCT and BM3D in terms of PSNR and $FSIM_C$. Since BM3D adopts a NN search selection strategy, it also suffers from the *noise-to-noise* matching problem in this case. Remarkably, for $\sigma \leq 20$, NLM coupled with SNN is better, in terms of PSNR, than the state-of-the-art BM3D-CFA. It is slightly inferior for higher noise levels, but at a lower computational cost. Visually, the result produced by NLM with SNN is comparable with that obtained by BM3D-CFA: our approach generates a slightly more noisy image, but with better preserved finer details (*e.g.*, see the eyelids in Fig. 3 in the main paper) and without introducing grid artifacts.

2.3 The case of a “real” image

Having in mind the practical application of SNN, it is worthy analyzing the peculiarities of image denoising for “real” images. We test different filtering techniques (SDCT, $NLM_{0.8}^{16}$, $NLM_{0.8}^{16}$, and BM3D-CFA) on a high resolution (2592×1944) image, captured at ISO 1200 with an NVIDIA Shield tablet. Notice that image resolution is much higher in this case, compared to the Kodak dataset. Furthermore, the noise distribution in the Bayer domain is far from the ideal, Gaussian distribution with constant variance. Therefore, we compute the sensor noise model and the corresponding Variance Stabilizing Transform (VST) as in [8] and apply the VST in the Bayer domain such that the noise distribution approximately resembles the ideal one. We filter the image through the SDCT, $NLM_{0.8}^{16}$, and $NLM_{0.8}^{16}$ denoising algorithms in the RGB domain, after applying the VST and Malvar demosaicing, which introduces correlations among pixels. The inverse VST is then applied to the data in the RGB domain. Differently from the other algorithms, BM3D-CFA is applied directly to raw data, after the VST and before the inverse VST. This represents an advantage of BM3D-CFA over the other approaches, that have been designed to deal with white noise, but are applied in this case to colored noise. Finally, after the denoising step, we apply color correction, white balancing, gamma correction, unsharp masking, coherently with the typical workflow of an Image Signal Processor; these steps can easily increase the visibility of artifacts and residual noise. A stack of 650 frames is averaged to build a ground truth image of the same scene.

	Noisy	SDCT	$NLM_{0.8}^{16}$ (NN)	$NLM_{0.8}^{16}$ (SNN)	BM3D-CFA
PSNR	21.2689	23.3715	24.0972	24.5500	25.2580
SSIM	0.8741	0.8889	0.9135	0.9261	0.9607
MSSSIM	0.7133	0.7773	0.8124	0.8450	0.8943
GMSD	0.1921	0.1499	0.1356	0.1145	0.0921
FSIM	0.9920	0.9927	0.9933	0.9940	0.9952
$FSIM_C$	0.9890	0.9899	0.9911	0.9921	0.9941

Table 2: Image quality metrics for the “real” image in Fig. 6, acquired with and NVIDIA Shield Tablet at ISO 1200 and denoised with different algorithms.

The results are shown in Fig. 6. Visual inspection confirms the superiority of $NLM_{0.8}^{16}$ over SDCT and $NLM_{0.8}^{16}$ even for the case of a real image. The $NLM_{0.8}^{16}$ filter preserves high frequency details in the image, but it also generates high frequency colored noise in the flat areas, because of the *noise-to-noise* matching problem. The SDCT filter produces a more blurry image with middle frequency colored noise in the flat areas. Coherently with the case of colored noise and the Kodak



Figure 6: Comparison of several patches from the image on the top, acquired with and NVIDIA Shield Tablet at ISO 1200, demosaiced and then denoised with the different algorithms. Color correction, white balancing, gamma correction and unsharp masking have also been applied after denoising. The ground truth image is obtained by averaging 650 frames. Better seen at 400% zoom.

dataset, our approach provides an image of quality comparable to BM3D-CFA; the NLM¹⁶ filter is slightly more noisy than BM3D-CFA, as measured by the image quality metrics in Table 2, but filtering requires a lower computational cost.

References

- [1] AHN, B., AND CHO, N. I. Block-matching convolutional neural network for image denoising. *CoRR abs/1704.00524* (2017).

- [2] BUADES, A., COLL, B., AND MOREL, J. M. A review of image denoising algorithms, with a new one. *Multiscale Model Sim.* (2005).
- [3] BUADES, A., COLL, B., AND MOREL, J.-M. Non-Local Means Denoising. *Image Processing On Line 1* (2011), 208–212.
- [4] DABOV, K., FOI, A., KATKOVNIK, V., AND EGIAZARIAN, K. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on Image Processing 16*, 8 (2007).
- [5] DABOV, K., FOI, A., KATKOVNIK, V., AND EGIAZARIAN, K. Bm3d image denoising with shape-adaptive principal component analysis. In *SPARS* (2009).
- [6] DANIELYAN, A., VEHVILAINEN, M., FOI, A., KATKOVNIK, V., AND EGIAZARIAN, K. Cross-color bm3d filtering of noisy raw data.
- [7] DUVAL, V., AUJOL, J.-F., AND GOUSSEAU, Y. On the parameter choice for the non-local means. *SIAM J. Imag. Sci.* (2011).
- [8] FOI, A. Clipped noisy images: Heteroskedastic modeling and practical denoising. *Signal Processing* (2009).
- [9] KNAUS, C., AND ZWICKER, M. Progressive image denoising. *IEEE TIP* (2014).
- [10] LEBRUN, M., COLOM, M., BUADES, A., AND MOREL, J. M. Secrets of image denoising cuisine. *Acta Numerica* (2012).
- [11] LOTAN, O., AND IRANI, M. Needle-match: Reliable patch matching under high uncertainty. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2016).
- [12] LOU, Y., FAVARO, P., SOATTO, S., AND BERTOZZI, A. Nonlocal similarity image filtering, 2009.
- [13] PONOMARENKO, N., JIN, L., IEREMEIEV, O., LUKIN, V., EGIAZARIAN, K., ASTOLA, J., VOZEL, B., CHEHDI, K., CARLI, M., BATTISTI, F., AND KUO, C.-C. J. Image database TID2013: Peculiarities, results and perspectives. *Signal Processing: Image Communication 30* (2015).
- [14] RICO MALVAR, LI-WEI HE, R. C. High-quality linear interpolation for demosaicing of bayer-patterned color images. In *International Conference of Acoustic, Speech and Signal Processing* (May 2004), Institute of Electrical and Electronics Engineers, Inc.
- [15] TSAI, Y.-T., STEINBERGER, M., PAJAK, D., AND PULLI, K. Fast ANN for high-quality collaborative filtering. *CGF* (2015).
- [16] WANG, Z., BOVIK, A. C., SHEIKH, H. R., AND SIMONCELLI, E. P. Image quality assessment: From error visibility to structural similarity. *IEEE TIP* (2004).
- [17] WANG, Z., SIMONCELLI, E., AND BOVIK, A. Multiscale structural similarity for image quality assessment. In *ACSSC* (2003).
- [18] WU, Y., TRACEY, B., NATARAJAN, P., AND NOONAN, J. P. Probabilistic non-local means. *IEEE Sig. Proc. Let. 20* (2013).
- [19] XUE, W., ZHANG, L., MOU, X., AND BOVIK, A. C. Gradient magnitude similarity deviation: A highly efficient perceptual image quality index. *IEEE TIP* (2014).
- [20] YU, G., AND SAPIRO, G. DCT Image Denoising: a Simple and Effective Image Denoising Algorithm. *IPOL* (2011).
- [21] ZHANG, L., ZHANG, D., MOU, X., AND ZHANG, D. Fsim: A feature similarity index for image quality assessment. *IEEE TIP* (2011).