

IMAGE DENOISING BY ADAPTIVE COMPRESSED SENSING RECONSTRUCTIONS AND FUSIONS

William Meiniel^{a,b}, Elsa Angelini^b, and Jean-Christophe Olivo-Marin^a

^aInstitut Pasteur, Unité d'Analyse d'Images Biologiques, CNRS UMR 3691, France

^bInstitut Telecom, Telecom ParisTech, CNRS LTCI, France

ABSTRACT

In this work, Compressed Sensing (CS) is investigated as a denoising tool in bioimaging. The denoising algorithm exploits multiple CS reconstructions, taking advantage of the robustness of CS in the presence of noise via regularized reconstructions and the properties of the Fourier transform of bioimages. Multiple reconstructions at low sampling rates are combined to generate high quality denoised images using several sparsity constraints. We present different combination methods for the CS reconstructions and quantitatively compare the performance of our denoising methods to state-of-the-art ones.

Keywords : Bioimaging, compressed sensing, denoising, total-variation, Fourier transform.

1. INTRODUCTION

The Compressed Sensing (CS) theory, introduced by^{1,2}, gives a solution to the problem of recovering a signal x of size N from an observation vector $y = \Phi x$ of size $M \ll N$, constituted by linear projections of x . Under the assumption that x has a sparse representation in some known linear projection space, CS theory states that the signal x can be recovered from the observation y made in an adequate projection space via convex optimization.

The purpose of image denoising is the recovery of an image x - that is called the ground truth image - from a degraded noisy observation y . A denoising algorithm is a process that takes as entry the observation y and creates an estimator \hat{x} of the ground truth x .

Several image denoising algorithms already exist in the literature, that consider different models of noise and exploit different properties of the images (see³ for a recent review). We can highlight the total-variation based filtering⁴ and the non-local means (NLM)⁵ as the basis of many current techniques. In this paper, we will apply our methods to images corrupted by a mixed Poisson-Gaussian noise, which corresponds to the reality of biological microscopy images.⁶

The use of Total Variation (TV) for denoising was proposed in⁴, and consists in minimizing the operator $\|x\|_{\text{TV}} = \sum_{p,q} \sqrt{\partial_h x(p,q)^2 + \partial_v x(p,q)^2}$ where $\partial_h x$ and $\partial_v x$ are the partial horizontal and vertical derivatives of the image x , and p, q are pixels coordinates. Convex optimization theory uses this TV operator in a wide and diverse range of applications.^{7,8} One of the main reasons of this popularity is the fact that TV-based optimization techniques tend to preserve sharp edges in the image. This framework is well adapted, in the context of biological microscopy, for the study of images containing simple cells with a stationary noise, but not necessarily for more complex tissues with non-stationary Poisson noise. In these cases, the TV-based methods tend to generate staircase artifacts, that we propose to attenuate in our method, while avoiding the use of Poisson TV with a data-fidelity term using an estimation of the local variance of the noise.^{9,10}

Inspired by the work of Marim et al.¹¹ and Meiniel et al.¹², we present a denoising method that takes advantage of the total-variation based regularization through multiple compressed sensing reconstructions. Following the idea of, for instance, ultrasound compounding technique¹³, our method combines multiple redundant but imperfect reconstructions. Modeling of noise properties is avoided via the introduction of a new image processing operator: the *variance map*, which infers edge locations and noise level directly from the series of CS reconstructions.

2. MULTIPLE COMPRESSED SENSING RECONSTRUCTIONS

Following the work of¹⁴, the Fourier Transform can be used as the observation projection operator Φ to recover a signal that is sparse in the spatial domain. Fourier coefficients from biological images acquired with a microscope can be measured via digital holography¹⁵ or optical Fourier Transform¹⁶. Our working assumption for biological samples is that the energy of the noise-free image is mostly concentrated in the low-frequency part of its Fourier domain (see¹⁷ p. 52).

Noisy observations y are randomly sampled in the Fourier domain, while most of the high-frequency coefficients are not selected. Since this selection is random, each draw gives different results, but they all bear some similarity. We generate several sampling patterns, that lead to several measurement subsets y_k , where $k \in \{1, \dots, R\}$. Each of these subsets is used to compute an estimator \hat{x}_k of the ground truth image x using CS reconstruction. Finally, the \hat{x}_k are aggregated through a fusion operator f , such that the final estimator $\hat{x} = f(\hat{x}_1, \dots, \hat{x}_k)$ is of higher quality than individual \hat{x}_k (e.g. corrects the staircase effect caused by the TV regularization). In Louchet and Moisan¹⁸, a similar concept is proposed, where a Monte Carlo Markov Chain implementing a Metropolis scheme is designed. In this model, the TV regularization is modeled in a Bayesian approach, and each element of the random process is an image constructed with respect to the posterior law of the model. Then, the images are combined through averaging into the final estimator.

2.1 Generation of multiple random subsets of samples

2.1.1 Undersampling of the Fourier coefficients

The denoising method is decomposed into three steps (see Fig. 1).

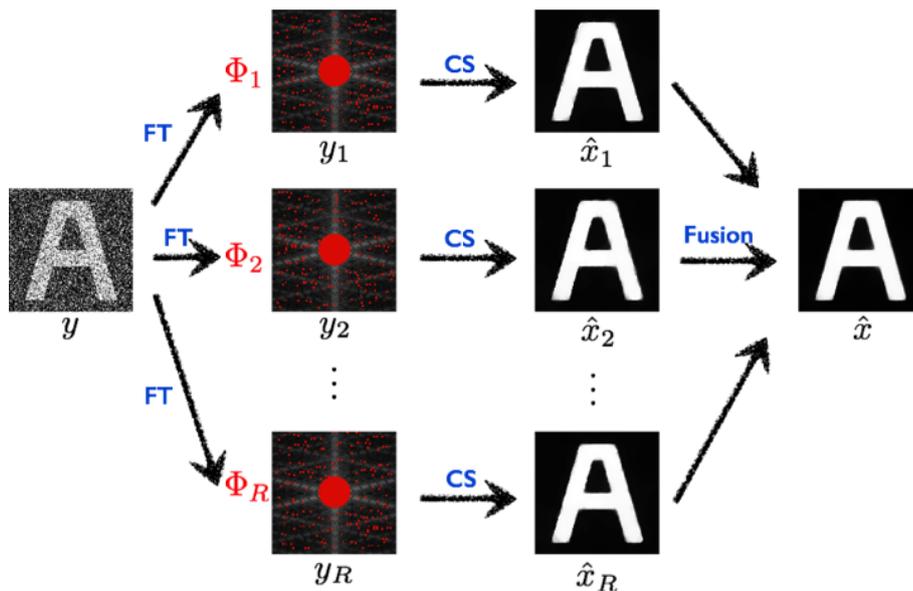


Figure 1: **Denoising scheme using several CS reconstructions.** From a noisy image y , a given number of CS measurement vectors y_k are generated, by taking the Fourier transform (FT) of y and selecting a subset of the Fourier coefficients. Then, each y_k is used to produce an estimator \hat{x}_k of the original image through a CS reconstruction scheme. Finally, all the \hat{x}_k are combined into an estimator \hat{x} .

The first step is the *generation of random subsets from measurement vectors y_k in the Fourier domain.*: As mentioned above, we perform a sub-sampling of the Fourier Transform of the noisy image y , following a random process in the high-frequency part of the Fourier Transform space ($\nu > \nu_c$). Gaussian or polynomial sampling are natural schemes to perform this kind of sampling, but LeMontagner in¹⁹ (p. 55) proposes a new design:

- i. First, keep all the low-frequency coefficients below a certain cut-off spatial frequency ν_c ,
- ii. Second, sample the Fourier coefficients above ν_c according to a uniform random process, such that the overall sampling rate of the Fourier Transform of y is τ .

The generation of these measurement vectors can be written as a linear transformation: $y_k = \Phi_k y$, where Φ_k is constructed by zeroing many coefficients in the 2D Fourier transform matrix corresponding to frequencies $\nu > \nu_c$. Φ_k follows the restricted isometry property²⁰ and thus is well suited for CS reconstructions.

Therefore, the two parameters for the sampling part are the cut-off frequency ν_c and the sampling rate τ . Their influence is studied in the Results section. Intuitively, the parameter ν_c , which is the radius of the fully sampled low-frequency area, is chosen so that it contains most of the relevant information of the natural image. We select its value so that the low-frequency area contains 90% of the energy of the whole image. This value corresponds, for the images we work with in this paper, to a circle containing from 3% up to 20% of the Fourier coefficients.

2.2 Generation of multiple reconstructions

We use the Compressed Sensing theory to recover the R estimators \hat{x}_k from the sub-sampled Fourier measures. Given each measurement vector y_k , we resolve a CS reconstruction problem using a TV regularization constraint:

$$\hat{x}_k = \arg \min_{x \in \mathbb{C}^N} \|x\|_{\text{TV}} \quad \text{s.t.} \quad \|\Phi_k x - y_k\|_2 \leq \epsilon \quad (1)$$

The reconstruction of Φ_k from a Fourier transform is performed very efficiently using convex optimization algorithms such as NESTA.⁸ This algorithm is based on the general framework developed in²¹ for the minimization of composite objective functions. It addresses the following constrained problem, with $f(x) = \|x\|_{\text{TV}}$ or $f(x) = \|\Psi^* x\|_1$ without specific requirements of the sparsity matrix Ψ^* , and using accelerated gradient descent with back-projection on the feasible set:

$$\arg \min_{x \in \mathbb{C}^N} f(x) \quad \text{s.t.} \quad \|\Phi_k x - y_k\|_2 \leq \epsilon \quad (2)$$

We have chosen this algorithm among many others, because of its trade-off between flexibility and execution speed (see¹⁹ (p. 37) for a detailed study on the comparison of different CS reconstruction algorithms for Fourier-based measures and TV regularization).

The parameters in this part of the algorithm are the number R of CS reconstructions and the weight ϵ of the data fidelity term to handle noise in the observations. Throughout the remainder of this paper, we set the value of R so that $\tau \times R = 1$.

The parameter ϵ is chosen as in⁸ : $\epsilon = \sigma \sqrt{\tau N + 2\sqrt{2\tau N}}$, where σ is an estimator of the Gaussian component of the noise. In the scenario of denoising as post-processing, it is estimated in a uniform background region of the image using cumulant method, matching the first four cumulants of y with the k -statistics of the samples (see²² for more details). As the noisy image is not available, we define an estimate \hat{y} of this image using the Fourier coefficients of each acquisition, combine them and compute the inverse Fourier transform. The error made for the computation of σ with this approach is fairly low ($\sim 5\%$).

2.3 Fusion of the reconstructions

Once the partial estimators are known, we propose to combine them through a *fusion operator*, to obtain a single estimator of the original image x :

$$\hat{x} = f(\hat{x}_1, \dots, \hat{x}_R)$$

Each of \hat{x}_R reconstructions provides an approximated denoised version of the original image. They differ due to the random sampling of Fourier coefficients but also share some structural coherence, sharing multiple common Fourier coefficients (especially in the low-frequencies) and spatial coherence thanks to the TV regularization. In previous works, we tested several fusion techniques, such as *non-local merging*, *weighted averaging*¹¹, or simple average of the \hat{x}_k ¹². We dedicate the next sub-sections to the introduction of different fusion operators f .

2.3.1 Fusion by averaging

Simple averaging of the \hat{x}_k has the advantage of being easy to compute and, moreover, of being an unbiased method for fusing the reconstructions without any bias between them. We define this fusion operator as follows:

$$\hat{x} = f(\hat{x}_1, \dots, \hat{x}_R) = \frac{1}{R} \sum_{k=1}^R \hat{x}_k \quad (3)$$

Results provided by this method are detailed in¹². The clear limitation of this approach is the smoothing of the image information due to variations within reconstructions with respect to the exact position of contours between structures. One legitimate question at this point is whether there is one reconstruction better than the other ones. We try to answer this question in the next section.

2.3.2 Oracle Fusion

In the particular case where the image has a known ground truth x , let us define the *oracle fusion* operator as:

$$\forall (p, q) \in \Omega, o_x(p, q) = \hat{x}_{k^*}(p, q) \quad \text{s.t.} \quad k^* = \arg \min_{k \in \{1, \dots, R\}} \left| \hat{x}_k(p, q) - x(p, q) \right| \quad (4)$$

where the algorithm 1 is described below in pseudo-code.

Algorithm 1: Oracle Fusion

Input : R partial reconstructions \hat{x}_k and ground truth image x .

Output: Estimator o_x .

```

1 for each pixel  $(p, q)$  do
2    $s_x(p, q) \leftarrow [\hat{x}_k(p, q) - x(p, q)]_{k=1, \dots, R}$ ;
3    $k^* \leftarrow \arg \min_{k \in \{1, \dots, R\}} s_x(p, q)$ ;
4    $o_x(p, q) \leftarrow \hat{x}_{k^*}(p, q)$ 
5 end

```

The results obtained with this method outperform the classical denoising algorithms (see Fig.2), as evidenced by the SSIM index. Hence, we set it as the reference method for estimating the quality of the fusion of Compressed Sensing reconstructions.

In order to test if this method follows a deterministic selection process, and if there is a way to aggregate the partial reconstructions in a method similar to the oracle one, without knowing the ground truth, we now take a closer look at the oracle method.

We notice that the oracle value is not representative of the statistics of the values among the R reconstructions. On a synthetic example with a noise-free image corrupted with mixed Gaussian-Poisson noise, we generated enough reconstructions for the oracle fusion to be almost perfect. Then, for each pixel, we computed the *samples vector*: $s_x(p, q) = [\hat{x}_1(p, q), \dots, \hat{x}_R(p, q)]$. The results of this experiment, shown on Fig.3, suggest that the oracle method does not follow a deterministic selection process, as it corresponds to outliers in the statistics.

It seems therefore difficult to search for an algorithm that would mimic the oracle selection of reconstructions. But if we compare the results given by the oracle fusion and the fusion by averaging, we notice that the latter algorithm performs poorly only at the edges of objects within the image (see Fig.4). Hence, we propose to improve the averaging fusion method by modifying the fusion rule at edges. We call this approach *spatially adaptive fusion* and describe it next.

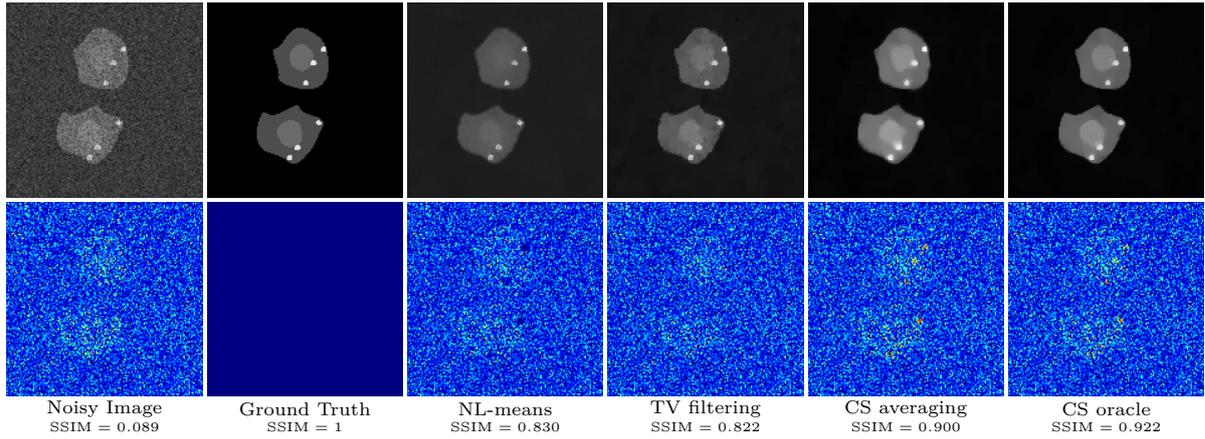


Figure 2: Result of the oracle fusion on the synthetic cell image (perturbed with a mixed Poisson-Gaussian noise with parameters $\sigma = 0.1$ and $\lambda = 0.02$). Comparison with classical denoising methods and with the fusion by averaging.

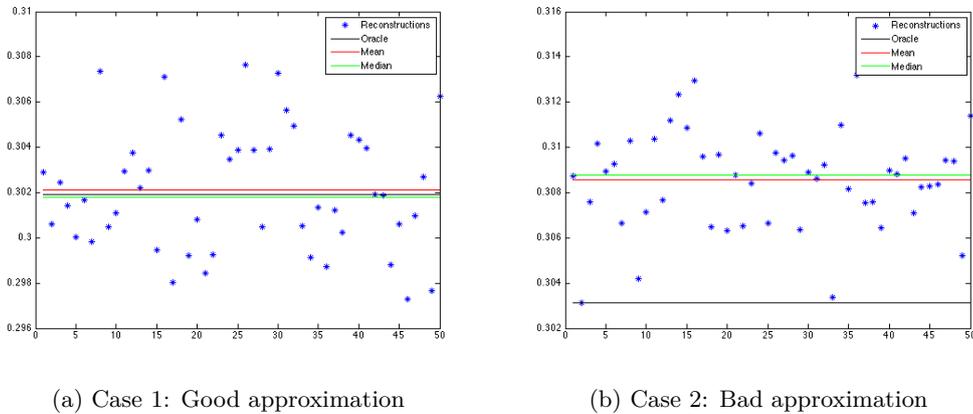


Figure 3: Comparison of the oracle fusion with the mean and the median of 50 reconstructions for two random pixels in the synthetic cell image. On the right-hand side image, it can be seen that the best value seems to be an outlier in the distribution of the sample values.

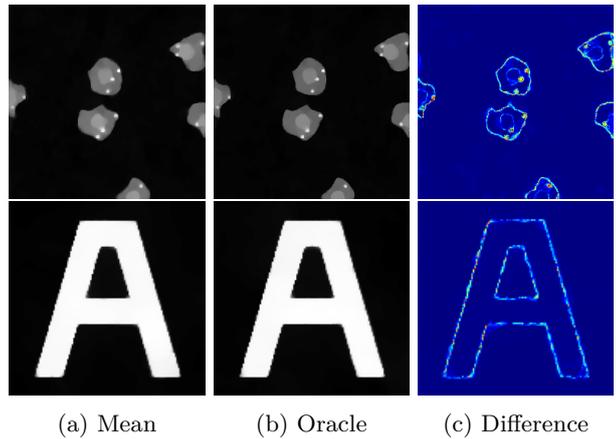


Figure 4: Comparison of oracle and averaging fusion results on two synthetic images (corrupted with a mixed Poisson-Gaussian noise of parameters $\sigma = 0.1$ and $\lambda = 0.02$) showing large differences at edge locations only.

2.3.3 Spatially-adaptive fusion

We introduce a spatially-adaptive fusion rule which varies within the image (i.e. between pixels). In order to do so, and since we do not have any ground truth, we propose to rely on an inference method to localize edge positions in the image being reconstructed, based on what we call the reconstruction *variance map*. The approach is similar to the works existing in the field of Anisotropic diffusion²³, where a prior detection of particular directions in the image is studied, in order to enhance the shapes.

Variance map: The so called *variance map* is derived as the classical standard deviation of a series of measures (see Fig. 5).

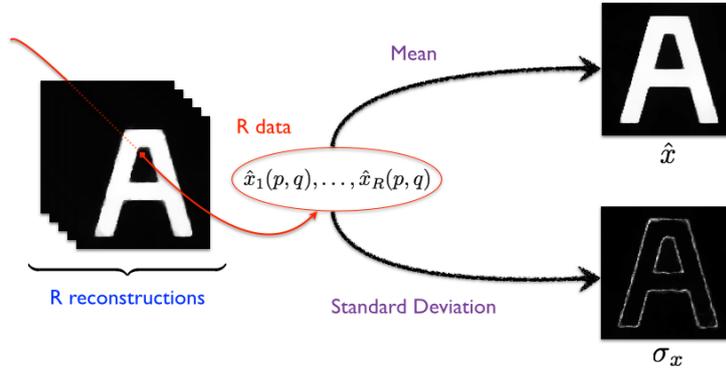
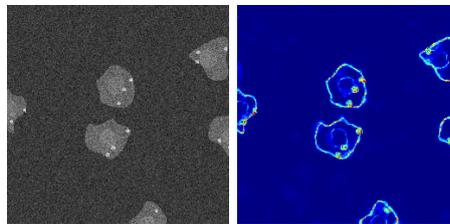


Figure 5: **Computation of the variance map.**

Keeping the same notations as above, we define the operator σ_x as follows:

$$\sigma_x = \sqrt{\frac{1}{R-1} \sum_{k=1}^R (\hat{x}_k - \hat{x})^2} \quad (5)$$

Applied on reconstructions from 2.3, this operator probes the pixel-wise disagreement between the partial CS reconstructions \hat{x}_k : high values correspond to pixels where the reconstructed gray level is highly dependent on the selection of the Fourier coefficients from the originally sampled signal y (i.e. the generation of y_k), while low values reveal consensus areas. An example is shown on Fig. 6. The shape of the variance map is very similar to the reconstruction errors induced by the averaging method, when using the oracle fusion as a reference (see Fig. 7). In other words, in order to recover the results given by the oracle method from the averaging reconstruction, the variance map gives a good indication of where discrepancies are highest. Moreover, the advantage of the variance map is that it does not require the knowledge of the ground truth of an image. This operator therefore allows us to work, even on images where the ground truth is not known.



(a) Synthetic cells (b) Variance map

Figure 6: **Variance map of the synthetic cell image. This map confirms our hypothesis that the overall reconstruction method performs very well on flat areas (cytoplasm and nucleus in the cells) but that the different reconstructions disagree around edges.**

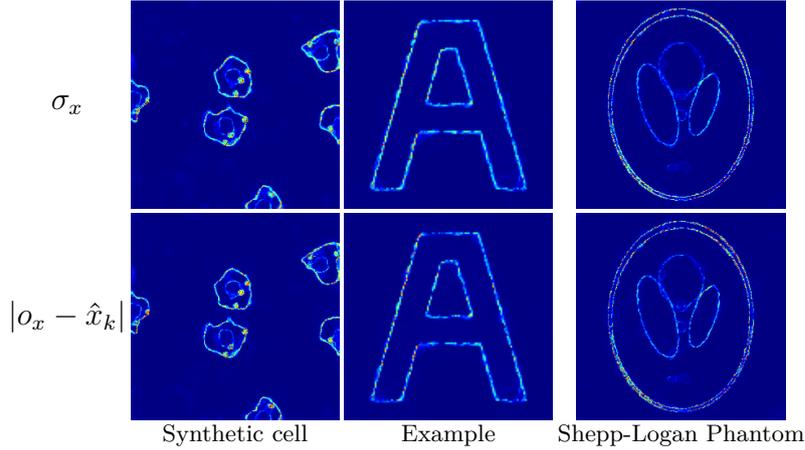


Figure 7: **Visual comparison between the variance map and the oracle-mean difference for several images.**

Next, we propose a new denoising scheme that takes advantage of the new information provided by the variance map, while taking into account the fusion method of 2.2. This will lead to a new class of fusion methods, that we will refer to as *spatially-adaptive fusion*.

Spatially-adaptive fusion operator

We propose to combine the reconstructions via averaging with an edge-preserving constraint, where the variance map acts as a weight. In what follows, and in order to simplify the notations, we adopt a matrix notation. The estimator is defined as:

$$\hat{x} = \sigma_x \circ y + (1 - \sigma_x) \circ \hat{x}_{\text{mean}} \quad (6)$$

where \circ represents the Hadamard product of two matrices, and σ here represents the normalized version of the variance map, so that the values of σ belong to $[0, 1]$.

This approach can be viewed as a barycenter estimate (quoting the term from¹⁸) between the original and the denoised images, where the confidence is weighted in the two images with the variance map. Indeed in¹⁸ the authors state that visual quality can be improved when noise and texture are not completely removed. Here we further guarantee that the original image information will be preserved at edges, which are the critical location for visual assessment.

Filtering method. In the cases when the noise is very strong or when the image has some particular properties, one may want to filter the noisy image prior to the adaptive reconstruction. In that case, if we denote by \mathcal{H} a local filter on an image (e.g. local gaussian filter, median filter, etc.), we define the filtered adaptive reconstruction as follows:

$$\hat{x} = \sigma_x \circ \mathcal{H}(y) + (1 - \sigma_x) \circ \hat{x}_{\text{mean}} \quad (7)$$

2.3.4 Other fusion methods

Here we give an overview of different fusion methods that we will compare with ours in the Results section.

Median, Min, Max For each pixel (p, q) in the image, define the *samples vector*, containing the values of each partial reconstruction at this pixel :

$$s_x(p, q) = \left[\hat{x}_1(p, q), \dots, \hat{x}_R(p, q) \right]$$

Then, for each pixel, select the median value of its statistics vector. This defines the median fusion:

$$\forall (p, q) \in \Omega, f(\hat{x}_1, \dots, \hat{x}_R)(p, q) = \text{med}(s_x(p, q))$$

The minimum fusion and the maximum fusion follow the same process, only replacing med by min or max operators.

Patch For this method, we need to define a neighborhood in the image. For instance, define the Von Neumann neighborhood of the pixel (p, q) in the image x as : $\mathcal{V}_x(p, q) = \{x, (p, q), x(p-1, q), x(p, q-1), x(p+1, q), x(p, q+1)\}$.

The idea of this fusion method is to select for each pixel the value of the element that is the most frequent in the neighborhood of this pixel in each partial reconstruction. Define the set where we want to seek for the most frequent value:

$$\mathcal{N}_{p,q} = \bigcup_{k=1}^R \mathcal{V}_{\hat{x}_k}(p, q)$$

Then, the final estimator is defined as follows:

$$\hat{x}(p, q) = \text{mode}(\mathcal{N}_{p,q})$$

3. RESULTS

When working on real images obtained in biological imaging, one major problem is the fact that we do not have access to the ground truth image; another problem is that the images present in the databases of the denoising literature do not fit with our hypotheses.

Therefore, in order to evaluate quantitatively the quality of our method, we have used a synthetic cell image generated by the SIMCEP simulator for fluorescent cell populations²⁴. We perturbed the ground truth image (Fig. 8) with a mixed Poisson Gaussian noise of parameters $\sigma = 0.1$ and $\lambda = 0.02$. In other words, if x is the ground truth image, make $z \sim \mathcal{P}\left(\frac{x}{\lambda}\right)$ follow a Poisson distribution, then the final noisy image is $y = \lambda z + b$, where $b \sim \mathcal{N}(0, \sigma^2)$ follows a Gaussian distribution.

On the synthetic images, we evaluate the quality of our results by using the Structural Similarity Index Measure (SSIM)²⁵. We have chosen this measure among others (MSE, SNR) because it corresponds better to the perception of the human eye.

We also assess the performances of our algorithm on real microscopy images of a HeLa cell, where we obtained a ground truth image using a much longer acquisition time than for the noisy version (Fig. 8) in order to smooth out the noise.

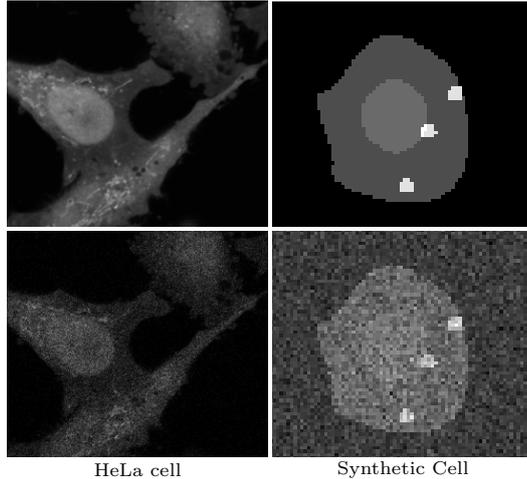


Figure 8: Natural and synthetic images used to assess the performances of our denoising algorithm. The synthetic image is perturbed with a mixed Poisson-Gaussian noise of parameters $\sigma = 0.1$ and $\lambda = 0.02$. The ground truth of the HeLa cell is obtained with a long acquisition time (7 min), and the noisy version is obtained with a fast acquisition time (2 sec). The conditions of the experiment were met to avoid the presence of photobleaching in this case (excitation with a LED, strong fluorescence marker, low exposure time).

3.1 Influence of the parameters

The algorithm that we propose in this paper has the advantage of using very few parameters. As stated earlier, there are only four of them: τ, ν_c, R and ϵ . Only τ and R have still to be determined, keeping in mind that we constrain $\tau \times R = 1$. In this part, we investigate the influence of these parameters on the averaging fusion operator. Similar comments can be made for other fusion operators as the median or the adaptive reconstruction.

We first fixed the number of reconstructions R to 5 and studied the influence of τ for the example image, for different values of ν_c . The results are shown on Fig. 9, where the SSIM of the reconstruction is measured for different values of τ and ν_c . For very low values of τ , the reconstruction is poor but, in the other hand, values of τ around 20% are enough to reach a quite optimal result.

On Fig. 10 we show the influence of the number of reconstructions R on the SSIM evaluation of the results of the method (with $\tau = 10\%$ and $\nu_c = 10\%$). The striking shape of the curve indicates that if multiple reconstructions perform better than just one, using too many reconstructions reduces the quality of the final image. This proves the degradation caused by the averaging of different images, that smooths the shapes in the image. For each processed image, there is an optimal value of R which never exceeds $R = 10$.

For the following experiments, we fix the values of τ and R respectively at 20% and 5, so that we meet the requirement that $\tau \times R = 1$.

3.2 Results on a synthetic image

In the case of the synthetic cell image, the ground truth is known. This allows to compare quantitatively the different fusion operators proposed in this paper with standard denoising algorithms. The results are shown in Fig. 11 and in Tab. 1.

As expected, the oracle method gives the best results, and the other methods have different characteristics. While the staircase effect is very strong on the TV, median and patch methods, it is smoothed out by the adaptive and averaging methods. The contrast inside the cell is poorly recovered by the NLM method. Visual contrast is better recovered on images processed with the patch method, but the small white objects are corrupted.

The method giving the best results in terms of SSIM - apart from the oracle reconstruction - is the spatially-adaptive reconstruction. Qualitatively, it also appears sharper than the other images. This result proves that the

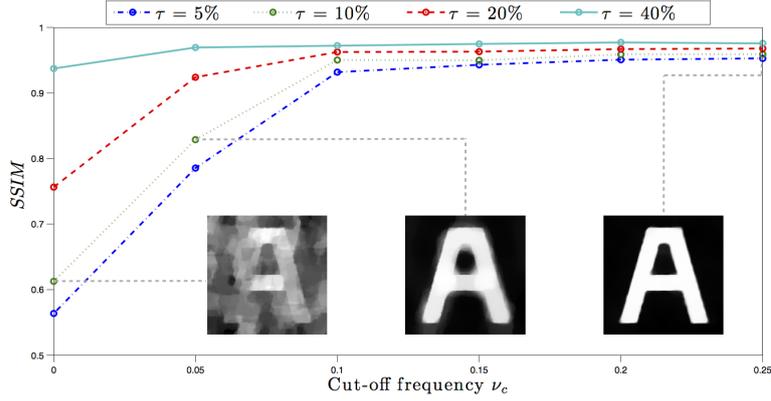


Figure 9: Denoising results via averaging of multiple CS reconstructions obtained on the example image for different values of the sampling parameters ν_c and τ . Note that the influence of τ is more visible for low values of ν_c but becomes marginal for higher values.

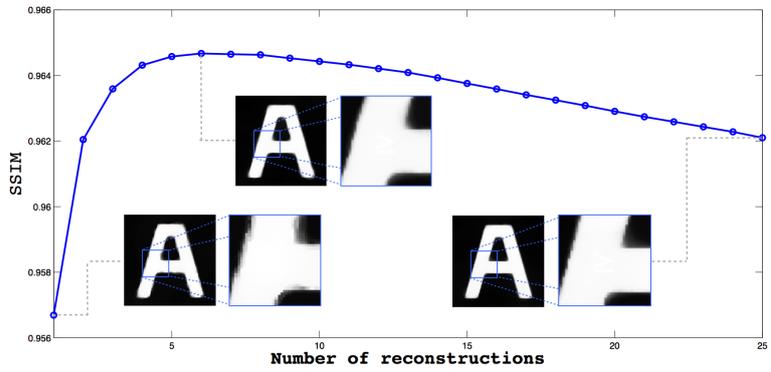


Figure 10: Denoising results via averaging of multiple CS reconstructions obtained on the example image for different values of the number of reconstructions R . The fusion of several reconstructions improves the quality of the algorithm, until a certain value of R (depending on the image we process). If we use more reconstructions that this value, the effect of the fusion degrades the result. In this example, $R = 6$ gives the best quantitative result.

	Averaging	Median	Patch	NLM	TV	Oracle	Adaptive	Adaptive (filtered)
Synthetic cell	0.903	0.904	0.867	0.831	0.838	0.928	0.906	0.902
Hela Cell	0.820	0.822	0.616	0.725	0.764	0.835	0.811	0.824

Table 1: SSIM Table for the Synthetic cell image corrupted with a mixed Poisson-Gaussian noise ($\sigma = 0.1$, $\lambda = 0.02$) and for the Hela cell. Cells in blue show, for each image, the best result among all different methods. The oracle fusion, which is the reference method, indeed shows the best results.

use of the variance map actually improves the reconstruction quality of the method proposed by¹². Finally, most of the methods proposed in this paper outperform the classical NLM and TV-filtering method for the synthetic cell image corrupted with a mixed Poisson-Gaussian noise.

3.3 Results on biological images

To quantitatively evaluate our algorithm on a real biological image, we acquired, with a confocal microscope, a Hela cell image with different exposure times, and labeled with a Cell Tracker CMFDA dye. Because of the different sources of noise present during the acquisition process⁶, these images are corrupted with mixed Poisson-Gaussian noise, which can be drastically reduced using a longer acquisition time. The dye that we used

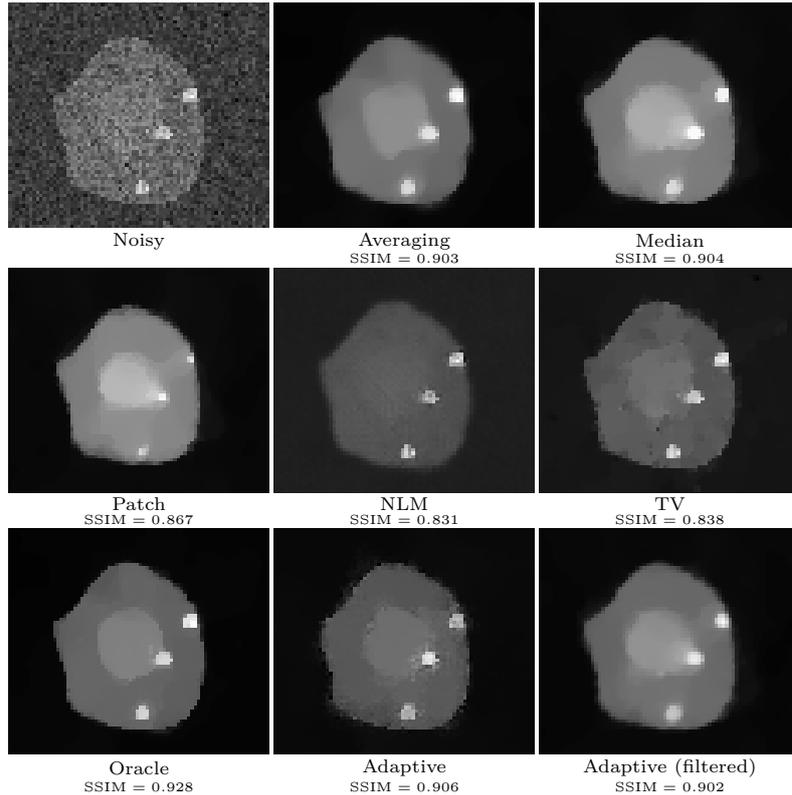


Figure 11: Denoising results on the synthetic cell image (perturbed with a mixed Poisson-Gaussian noise with parameters $\sigma = 0.1$ and $\lambda = 0.02$) obtained with the different methods proposed and the standard algorithms. For the proposed algorithm, we used $R = 5$ and $\tau = 20\%$. In this example, the ground truth is given, thus we can compute the SSIM of each reconstruction.

for this experiment is strong enough, and the laser weak enough, to avoid photobleaching.

We first acquired one image with a short acquisition time (2 seconds), and used this image as the "reference" noisy image y . Then we acquired the same field of view for a longer time (7 minutes), and used this other image as the ground truth x . The results of our algorithm are compared with standard denoising methods (see Fig. 12 and Tab. 1).

Except for the patch reconstruction method which reconstructs many outliers and behaves poorly in the background, the different fusion methods give similar results. The amount of noise removed in the process is larger than for the NLM and TV-filtering methods. This remark is confirmed with the SSIM values obtained for each reconstruction, our method outperforming the standard ones in each case. However, artifacts appear in the images and some outliers remain visible in most cases.

Similarly to the synthetic example, the adaptive method gives the best quantitative results. The filtered version looks smoother, as the outliers have completely disappeared. With the adaptive methods, we propose an alternative to classical denoising techniques.

4. CONCLUSION

In this paper, we have presented a new denoising method based on multiple Compressed Sensing reconstructions, and presents a new image processing operator: the variance map. The method gives results that are comparable with classical denoising methods, while relying on image reconstruction principles that could be tuned directly during an acquisition process. The proposed spatially-adaptive reconstruction algorithm has the advantage of using very few parameters, that are easy to tune, and makes no assumption on the noise present in the tested

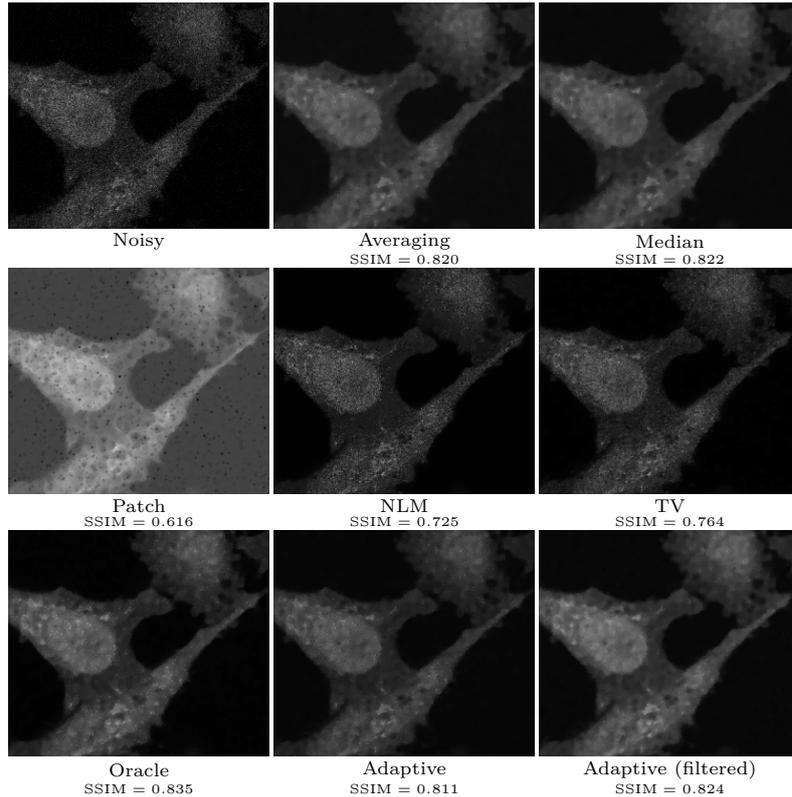


Figure 12: Denoising results on the HeLa cell image obtained with the different methods proposed and the standard algorithms. For the proposed algorithm, we used $R = 5$ and $\tau = 20\%$. In this example, we use as a ground truth the same image obtained with a longer acquisition time. Hence we can compute the SSIM of each reconstruction.

images. Along with the spatially-adaptive reconstruction, our method opens the possibility of creating many other denoising schemes, just by changing the fusion operator. We compared several of these fusion operators, but it would be interesting to look for even more diverse fusion schemes.

5. ACKNOWLEDGMENTS

William Meiniel would like to thank the Ecole Normale Supérieure de Cachan for his PhD fellowship, the Heffner Biomedical Imaging Laboratory at Columbia University for hosting him during part of this project, the GDR ISIS for partial funding, and Maria Manich from Institut Pasteur for the acquisition of the HeLa cell image. This work was partially funded through grants from the Agence Nationale de la Recherche (ANR-10-INSB-04-06 France Bioimaging and ANR-10-LABX-62-IBEID).

REFERENCES

- [1] Candès, E. J. et al., “Compressive sampling,” *Proceedings of the international congress of mathematicians* **3**, 1433–1452 (2006).
- [2] Candès, E. J., Romberg, J., and Tao, T., “Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information,” *IEEE Transactions on Information Theory* **52**(2), 489–509 (2006).
- [3] Milanfar, P., “A tour of modern image filtering: New insights and methods, both practical and theoretical,” *Signal Processing Magazine* **30**(1), 106–128 (2013).

- [4] Rudin, L. I., Osher, S., and Fatemi, E., “Nonlinear total variation based noise removal algorithms,” *Phys. D* **60**(1-4), 259–268 (1992).
- [5] Buades, A., Coll, B., and Morel, J.-M., “A review of image denoising algorithms, with a new one,” *Multiscale Modeling & Simulation* **4**(2), 490–530 (2005).
- [6] Le Montagner, Y., Angelini, E., and Olivo-Marin, J.-C., “An unbiased risk estimator for image denoising in the presence of mixed poisson-gaussian noise,” *Transactions on Image Processing* **23**(3), 1255–68 (2014).
- [7] Chambolle, A., “An algorithm for total variation minimization and applications,” *Journal of Mathematical Imaging and Vision* **20**(1-2), 89–97 (2004).
- [8] Becker, S., Bobin, J., and Candès, E., “Nesta: A fast and accurate first-order method for sparse recovery,” *SIAM Journal on Imaging Sciences* **4**(1), 1–39 (2011).
- [9] Mugnier, L., Fusco, T., and Conan, J.-M., “Mistral : a myopic edge-preserving image restoration method, with application to astronomical adaptive-optics-corrected long-exposure images,” *Journal of the Optical Society of America* **21**(10), 1541–54 (2004).
- [10] Le, T., Chartrand, R., and Asaki, T., “A variational approach to reconstructing images corrupted by poisson noise,” *Journal of Mathematical Imaging and Vision* **27**(3), 257–263 (2007).
- [11] Marim, M., Angelini, E., and Olivo-Marin, J.-C., “Denoising in fluorescence microscopy using compressed sensing with multiple reconstructions and non-local merging,” in [*International Conference of the IEEE EMBS*], (2010).
- [12] Meiniel, W., Le Montagner, Y., Angelini, E., and Olivo-Marin, J.-C., “Image denoising by multiple compressed sensing reconstructions,” in [*IEEE International Symposium on Biomedical Imaging (ISBI)*], (April 2015).
- [13] Tanter, M., Bercoff, J., Sandrin, L., and Fink, M., “Ultrafast compound imaging for 2-d motion vector estimation: application to transient elastography,” *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*. **49**(10), 1363–74 (2002).
- [14] Candès, E. and Tao, T., “Near-optimal signal recovery from random projections: Universal encoding strategies?,” *IEEE Transactions on Information Theory* **52**(12), 5406–25 (2006).
- [15] Marim, M., Angelini, E., Olivo-Marin, J.-C., and Atlan, M., “Off-axis compressed holographic microscopy in low-light condition,” *Optics Letters* **36**(1), 79–81 (2011).
- [16] Le Montagner, Y., Angelini, E., and Olivo-Marin, J.-C., “Comparison of reconstruction algorithms in compressed sensing applied to biological imaging,” in [*IEEE International Symposium on Biomedical Imaging (ISBI)*], (2011).
- [17] Marim, M., *A compressed Sensing framework for biological microscopy*, PhD thesis, Télécom ParisTech (2011).
- [18] Louchet, C. and Moisan, L., “Posterior expectation of the total variation model: properties and experiments,” *SIAM Journal on Imaging Sciences* **6**(4), 2640–84 (2013).
- [19] LeMontagner, Y., *Algorithmic solutions toward applications of compressed sensing for optical imaging*, PhD thesis, Télécom ParisTech (2013).
- [20] Candès, E., “The restricted isometry property and its implications for compressed sensing,” *Comptes Rendus Mathématique* **346**(9 - 10), 589 – 592 (2008).
- [21] Nesterov, Y., “Gradient methods for minimizing composite objective functions,” tech. rep., Université Catholique de Louvain (2007).
- [22] Rose, C. and Smith, M., [*Mathematical statistics with Mathematica*], ch. 7.2c: k-statistics: Unbiased estimators of cumulants, Springer-Verlag (2002).
- [23] Weickert, J., Ter Haar Romeny, B., and Viergever, M., “Efficient and reliable schemes for nonlinear diffusion filtering,” *IEEE Transactions on Image Processing* **7**, 398–410 (1998).
- [24] Ruusuvuori, P. et al., “Benchmark set of synthetic images for validating cell image analysis algorithms,” in [*16th European Signal Processing Conference*], (August 2008).
- [25] Wang, Z. and Bovik, A., “Mean squared error: Love it or leave it? a new look at signal fidelity measures,” *IEEE Signal Processing Magazine* **26**(1), 98–117 (2009).