Denoising of Microscopy Images: A Review of the State-of-the-Art, and a New Sparsity-Based Method

William Meiniel[ID], Jean-Christophe Olivo-Marin[ID], *Fellow, IEEE*, and Elsa D. Angelini[ID], *Senior Member, IEEE*

Abstract—This paper reviews the state-of-the-art in denoising methods for biological microscopy images and introduces a new and original sparsity-based algorithm. The proposed method combines total variation (TV) spatial regularization, enhancement of low-frequency information, and aggregation of sparse estimators and is able to handle simple and complex types of noise (Gaussian, Poisson, and mixed), without any *a priori* model and with a single set of parameter values. An extended comparison is also presented, that evaluates the denoising performance of the thirteen (including ours) state-of-the-art denoising methods specifically designed to handle the different types of noises found in bioimaging. Quantitative and qualitative results on synthetic and real images show that the proposed method outperforms the other ones on the majority of the tested scenarios.

Index Terms—Bioimaging, sparsity, image denoising, total variation, mixed Poisson-Gaussian model.

I. INTRODUCTION

LOW-light conditions and short exposure times are current strategic goals of research in biological imaging for longer observations and lesser degradation of specimens, but raise major challenges as they significantly increase the noise influence and degrade image quality. Two different avenues are being investigated currently to tackle these challenges: (1) super-resolution imaging such as STED, PALM/STORM or SIM [1], which rely on multiple full-field acquisition but result in computational issues for data streaming and computational resources required for image reconstruction, and (2) dedicated denoising methods based on mixed Gaussian-Poisson noise models, which are specific to low noise levels.

This paper presents a review of a selection of 12 state-of-the-art denoising methods together with an original sparsity-based denoising algorithm that are benchmarked and tested on the different types of strong and complex noise (Gaussian, Poisson and mixed) found in microscopic images. While all the methods have been originally designed for general denoising purpose, most of them have been adapted afterwards to restore images acquired with deterministic or Poisson measures as well as with additive or multiplicative noise components. Even if they have not been designed specifically for microscopic imaging, the selected methods represent, in our view, the best performing denoising options currently available, with the additional advantage that the respective codes were made available and are distributed by their authors. A detailed presentation is given in Sections II and III, while results are provided in Section IV.

In the remainder of this introduction we review in detail noise properties in microscopic imaging. In Section II we review 12 state-of-the-art denoising methods [2]–[13] that we tested against our method introduced in Section III. We present in Section IV comparative results on simulated and real microscopic images obtained in low-light conditions. Parameters used for each denoising method have been optimized for these conditions, and are reported along with the code sources for reproducibility. Indeed, to further promote reproducible research and build a collective know-how, we share all the image data and denoising results of this *in extenso* study with the community towards two goals: documenting parameter optimization approaches on shared codes for image denoising, and building up a database of noisy and ground-truth microscopic images to further test denoising methods on a common and shared benchmark (http://icy.bioimageanalysis.org/matlab/state-of-the-art-denoising).

A. Noise Models in Microscopic Image Processing

In the context of microscopy, there are three main sources of noise, inherent to the digital microscopy sensing apparatus.

The dark noise corresponds to the electronic noise generated by thermal agitation of electrons. More precisely, vibrations of silicon atoms in the camera sensor's substrate liberate electrons even when no incident photon is detected. High-quality sensors have a cooler to reduce dark noise. In terms of image models,

the dark noise follows a Poisson distribution $\mathcal{P}(\lambda_d)$, where λ_d represents the average dark flux. This noise affects the data model in the background, making it stochastic rather than null, and following a Poisson distribution.

The photon noise, or shot noise, is generated from the statistical fluctuations of the number of photons sensed at a given exposure level. Due to the stochastic nature of photon emissions, photon noise is inherent in all optical signals. In terms of signal modeling, we assume an average photon flux λ_p, and what is recorded is a number of photon counts that follows a Poisson distribution with parameter λ_p. This noise affects the signal model inside the structures being imaged, making it stochastic rather than deterministic, and following a Poisson distribution.

The readout noise is mainly generated by the non-perfectness of the output amplifier during the process of converting charges into voltages. In terms of image modeling, this noise is usually described as an additive component with a zero-mean Normal distribution.

The level of noise depends on the exposure time, experimental conditions affecting the sensors such as the temperature, or other parameters like the fluorescence of the structures being imaged in the case of fluorescence microscopy.

Three models are mainly used in microscopic image denoising [14]. In general terms, we note $x \in \mathbb{R}^N$ the true image, where N is the number of pixels in the image, and $y \in \mathbb{R}^N$ the observation of x, corrupted by noise. We also denote by Ω the domain of the image.

1) Additive White Gaussian Noise: The additive white Gaussian noise is the standard model used in most of the image processing literature (see [14] for instance). It represents the readout noise in digital microscopy, modeled as an additive random component to each true measure, with independent temporal values drawn from a zero-mean and constant variance Gaussian probability distribution function. Mathematically, the model is written $y = x + n$, where $n \in \mathbb{R}^N$ follows a zero-mean Gaussian distribution, $n \sim \mathcal{N}(0, \sigma^2)$. In other words, for each pixel coordinates $s \in \Omega$,

$$\mathbb{P}\left(y(s) = g\right) = \frac{1}{\sqrt{2\pi\sigma^2}}e^{-\frac{(x(s)-g)^2}{2\sigma^2}}, \tag{1}$$

where g is a gray level value and σ is the standard deviation of the noise model.

2) Poisson Noise: Poisson noise is actually the result of observing a deterministic signal through a detector that samples it as a Poisson process, i.e. dark noise and photon noise, with inherent fluctuations in the measured values and signal-dependent levels of fluctuations in the final pixel measure. Poisson noise is the main source of fluctuations in microscopy. The observed measures follow a discrete probability distribution that expresses the probability of a given number of events to occur in a fixed interval of time. Mathematically, the model describing an image perturbed with Poisson noise can be written as $y(s) = \mathcal{P}(\lambda_s)$, where $\mathcal{P}$ is a Poisson random process of intensity parameter λ_s. In other words, for each pixel coordinate $s \in \Omega$,

$$\mathbb{P}\left(y(s) = g\right) = \frac{\lambda_s^g\, e^{-\lambda_s}}{g!}. \tag{2}$$

The denoising process aims at estimating the underlying intensity value $\lambda_s = x(s)$.

The mean and variance of a Poisson law are both equal to λ_s and vary across the image domain. The brighter parts of an image have a higher mean value but also a higher variance, and therefore higher noise level. The parameter λ_s is directly linked to the quantity of photons arriving at each pixel. In microscopy, it depends on the exposure time used to acquire the image y, but also on the fluorescence level in case of fluorescence microscopy.

3) Mixed Poisson-Gaussian Noise: A more realistic noise model in microscopy [14], [15] consists in a combination of both Poisson and Gaussian noise, which is commonly called Mixed Poisson-Gaussian (MPG) noise. This model takes into account the three main sources of noise (dark noise, shot noise and readout noise). Mathematically, we write $y(s) = \gamma \mathcal{P}(\lambda_s) + n$, where γ is a gain constant that modulates the predominance of the contribution of the Poisson noise in the model. $\mathcal{P}$ is a Poisson random process and n is a zero-mean Gaussian random variable. In other words, assuming that the Poisson and Gaussian random processes are independent, for each pixel coordinate $s \in \Omega$ (cf. [16]),

$$\mathbb{P}\left(y(s) = g\right) = \frac{e^{-\lambda_s}}{\sqrt{2\pi\sigma^2}} \times \sum_{p=0}^{+\infty} \frac{\lambda_s^p}{p!}e^{-\frac{(\gamma p - g)^2}{2\sigma^2}}. \tag{3}$$

II. STATE-OF-THE-ART DENOISING METHODS

We first present a review of a broad selection of 12 state-of-the-art denoising methods that are tested in Section IV. While all the methods have been designed for general denoising purposes, most of them have been adapted afterwards to recover images acquired with deterministic or Poisson measures and with additive or multiplicative noise components. Even if they have not been particularly designed for microscopic imaging, they represent, in our view, the best performing denoising options currently available. Our selection was also driven by the availability of the codes to run them, shared or distributed by their authors. The list of reviewed methods is summarized in Table I, and sources for the codes in Tables II-III.

A. Total Variation-Based Methods

Total Variation (TV)-based methods have been introduced to infer, in an optimization framework, the most probable original noiseless image given a noisy observation and a noise model, as well as an image model involving the nature of the measure (deterministic or Poisson) and the spatial regularity of the structures in the image (via TV semi-norm suited for piecewise constant structures). We review the standard TV filter and a subset of four TV-based denoising methods designed for Poisson measures and complex noise models. In all reviewed methods, the TV norm metric $\|\cdot\|_{TV}$ refers to the isotropic TV semi-norm defined in [17] as

$$\|x\|_{TV} = \sum_{s\in\Omega} \sqrt{(\partial_h x(s))^2 + (\partial_v x(s))^2}, \tag{4}$$

where ∂_h and ∂_v represent the horizontal and vertical discrete derivative operators.

TABLE I

REVIEW OF THE TESTED DENOISING METHODS

Acronym	Ref.	Method	A priori noise model	Signal model
TV	[2]	Convex optimization	Additive	Deterministic
TV-ICE	[3]	Iterative local update of the pixel values	Additive	Deterministic
TV-MAP Poisson	[4], [61]	Bayesian	Poisson	Probabilistic
Poisson EM-TV	[5]	Convex optimization	Poisson	Probabilistic
MIDAL	[6]	Variable splitting convex optimization	Multiplicative	Deterministic
NLM	[7]	Non-local filtering	Additive	Deterministic
NLM-Poisson	[8]	Non-local filtering	Poisson	Probabilistic
NLPCA	[9]	Principal component analysis	Poisson	Probabilistic
PNLW	[10]	Non-local filtering	Poisson	Probabilistic
Analysis K-SVD	[11]	Dictionary learning	Additive	Deterministic
Wavelet thresh.	[12]	Wavelet domain filtering	Additive	Deterministic
BM3D	[13]	Non-local wavelet-based filtering	Additive	Deterministic
FSR		Fusion of several sparsity-based reconstructions	High-frequency	Deterministic

TABLE II

SOURCES OF CODE IMPLEMENTATION AND PARAMETER VALUES OF THE DIFFERENT DENOISING METHODS FOR THE *Synth* IMAGE. NOTE THAT ALL THE *In-House* CODES HAVE BEEN WRITTEN IN MATLAB. IN ADDITION, CODES WITH A SYMBOL "*" ARE USING MEX FILES

Method	Code source	Parameters
TV	in-house, using NESTA* (see Table III(a))	$\epsilon = 65.53$
TV-ICE	in-house, based on [3]	$\lambda = 0.2$, $n_{\text{iter}} = 10$ (convergence reached)
TV-MAP Poisson	Abergel website (see Table III(b))	$\lambda = 0.7$, $n_{\text{iter}} = 200$
Poisson EM-TV	Sawatzky himself	$w = 256^2$, $\beta = 0.002$
MIDAL	Bioucas Dias himself	$M = 2$, $\lambda_{\text{mid}} = 4$, $n_{\text{iter}} = 200$
NLM	Peyré Matlab toolbox (see Table III(c))	Patch size = 4 pixels, $\alpha = 0.1$
NLM-Poisson*	Deledalle website (see Table III(d))	Patch size for pre-estimate = 4 pixels
NLPCA	Salmon website (see Table III(e))	$P = 10$, $K = 14$, $\ell = 4$
PNLW	Bindilatti himself	$\gamma = 0.32$, $\beta = 20$, Patch size = 11
Analysis K-SVD	Elad website (see Table III (f))	$N_{patches} = 1000$, $n_{patch} = 5$
Wavelet thresh.	in-house	$\lambda_{\text{wav}} = 0.2$, wavelet depth = 10, wavelet type = haar
BM3D*	Dabov website (see Table III(g))	$\sigma_b = 25$
FSR	in-house, using NESTA*	$R = 3$, $\nu_c = 0.3$, $\tau = 0.2$, $\epsilon = 11.59 (\gamma_\epsilon = 1)$

TABLE III

WEBSITES FOR SOURCE CODES

Ref.	Method	Author	Link
(a)	NESTA	Becker	http://statweb.stanford.edu/~candes/nesta/
(b)	TV-MAP Poisson	Abergel	http://www.math-info.univ-paris5.fr/~rabergel/tvdenoise-poissonian.html
(c)	NLM	Peyré	http://www.mathworks.com/matlabcentral/fileexchange/13619-toolbox-non-local-means
(d)	NLM-Poisson	Deledalle	http://www.math.u-bordeaux.fr/~cdeledal/poisson_nlmeans.php
(e)	NLPCA	Salmon	http://josephsalmon.eu/code/index_codes.php?page=NLPCA
(f)	Analysis K-SVD	Elad	http://www.cs.technion.ac.il/~elad/software/
(g)	BM3D	Dabov	http://www.cs.tut.fi/~foi/GCF-BM3D/index.html#ref_software

1) TV-Filtering: Image filtering and denoising using TV was introduced by Rudin *et al.* [2]. The proposed method assumes that an image is composed of a small number of piecewise constant sets, and a small number of discontinuities and edges. The denoised image is obtained as a minimizer of the following functional:

$$\hat{x}_{\text{TV}} = \arg\min_{x \in \mathbb{R}^N} \|x\|_{\text{TV}} \quad \text{s.t.} \quad \|x - y\|_2 \le \epsilon \quad (5)$$

where y is the input noisy image, $\hat{x}_{\text{TV}}$ is the inferred true image, and $\epsilon > 0$ is a real constant that is proportional to the variance of the noise in the image. Note that this approach was designed for an additive Gaussian noise model. This denoising technique is very popular and extremely powerful on piecewise constant images, but suffers from the well known *staircase effect* in large structures of almost but non-exactly constant intensity values.

2) TV-ICE: Before introducing the TV-ICE model, we need to define the TV-LSE method, that Louchet and Moisan proposed in [18] to reduce the staircase effect. They reformulate the TV optimization problem using a least square error (LSE) criterion on a Bayesian formulation of the problem. They define image sets $\mathcal{E}_\mu$, and prior probability density functions p_β of the true image x on these sets as follows, with parameter $\beta > 0$:

$$\forall \mu \in \mathbb{R}, \quad \mathcal{E}_\mu = \left\{ x \in \mathbb{R}^\Omega, \sum_{s \in \Omega} x(s) = \mu|\Omega| \right\}$$

and

$$p_\beta(x) = \frac{1}{Z_\beta} e^{-\beta \|x\|_{\text{TV}}}, \quad \text{where} \quad Z_\beta = \int_{\mathcal{E}_0} e^{-\beta \|x\|_{\text{TV}}} \, dx.$$

In the case of *Gaussian noise*, they write $y = x + n$, with $n \sim \mathcal{N}(0, \sigma_n^2)$. In a Bayesian framework, they formulate the

posterior density function (pdf) $p(x|y)$ as:

$$p(x|y) = \frac{1}{Z} \exp\left(-\frac{E_\lambda(x)}{2\sigma_n^2}\right) \qquad (6)$$

where

$$E_\lambda(x) = \|x - y\|_2^2 + \lambda \|x\|_{\text{TV}},$$

$\lambda = 2\beta\sigma_n^2$, and Z is a normalizing factor ensuring that $x \mapsto p(x|y)$ is a probability density function on $\mathbb{R}^\Omega$. Here images are modeled as finite sets of pixels with real values. As the authors explain, regular TV filtering generates the *maximum a posteriori* (MAP) estimator of this pdf, and can introduce staircase effects in the solution with globally maximal probability. As an alternative, they propose to maximize the pdf using a LSE criterion. This leads to the following LSE estimate:

$$\hat{x}_{\text{LSE}} = \frac{\displaystyle\int_{\mathbb{R}^\Omega} \exp\left(-\frac{E_\lambda(x)}{2\sigma_n^2}\right) \cdot x \, dx}{\displaystyle\int_{\mathbb{R}^\Omega} \exp\left(-\frac{E_\lambda(x)}{2\sigma_n^2}\right) dx} = \mathbb{E}_{x \sim \pi}(x) \qquad (7)$$

with the posterior distribution:

$$\pi(x) = \exp\left(-\frac{E_\lambda(x)}{2\sigma_n^2}\right) \Big/ \int_{\mathbb{R}^\Omega} \exp\left(-\frac{E_\lambda(x)}{2\sigma_n^2}\right) dx$$

The authors propose an algorithmic solution to Eq (7) using a Monte Carlo Markov Chain (MCMC) algorithm with a Metropolis scheme. The denoised pixels are weighted averages over pixels from multiple possible image candidates.

Louchet and Moisan proposed in [3] the TV-Iterated conditional expectation (TV-ICE) as a variant of TV-LSE to improve computational efficiency. They replace the posterior mean (obtained via the MCMC model) by an iterated conditional marginal mean. This leads to the definition of the following recursion operator that converges towards its fixed point: $\forall s \in \Omega, p \geq 0$,

$$x^{p+1}(s) = \mathbb{E}_{x \sim \pi}\left[x(s)|x\,(\Omega\backslash\{s\}) = x^p\,(\Omega\backslash\{s\})\right]$$

with:

$$x^{p+1}(s) \xrightarrow[p \to \infty]{} \hat{x}_{\text{ICE}}(s)$$

where $\mathbb{E}_{x \sim \pi}$ is defined in Eq. (7).

From an algorithmic point of view, the denoised pixels are posterior means, conditionally to the values of all other pixels.

Both the TV-LSE and TV-ICE methods show strong performance in terms of reducing the staircase effect and the TV-ICE technique shows great improvement in terms of computation time. The main drawback of these approaches is their sensitivity to the noise model and level. If the noise model is not additive Gaussian, or if the noise variance σ_n^2 is very large, denoising performance is greatly degraded for both methods. The authors have recently extended in [4] the TV-ICE denoising framework to Poisson noise using the Poisson pdf in the data fidelity term. We denote by $\hat{x}_{TV-ICE-Poisson}$ the estimator obtained following the TV-ICE process with a Poisson data fidelity. Using TV regularization to infer the posterior Poisson likelihood raises issues that have been the focus of several works [5], [19]–[21]. Options are direct gradient descent with TV approximation via the divergence operator, following a log-likelihood variational approach, that leads to computational difficulties due to the nonlinearity of the model [19], the Expectation-Maximization (EM) applied to KL-divergence which is slow and introduces "checkboard" artefacts [5] or applying directly the split Bregman method [20], [21] which is complex to implement and has a high computational cost. In this work, we chose to test the Poisson EM-TV approach detailed below, because it offers the best compromise between sophistication and computational complexity.

3) Poisson EM-TV: The TV-based denoising scheme proposed by Rudin *et al.* [2] was formulated for deterministic measures corrupted with additive Gaussian noise. The Poisson-TV variational model, derived from the negative log likelihood function, is:

$$\tilde{x}_{\text{P-TV}} = \arg\min_{x \in \mathbb{R}^N} \sum_{s \in \Omega} \Big(x(s) - y(s)\log[x(s)]\Big) + \lambda\|x\|_{TV},$$

where Ω is the image domain and λ is the positive regularization parameter, as before. Note that in a Poisson statistics model, $x(s), y(s) > 0, \forall s \in \Omega$. This model is strongly nonlinear in the data fidelity term, which leads to issues in the computation of minimizers. Authors in [5] proposed a formulation of the solution of the Poisson-TV regularization problem alternating between expectation maximization (EM) and TV regularization.

Alternating EM inference and TV regularization, along with a simplification of the data fidelity term with a second order Taylor approximation, leads to the following modified TV-filter model:

$$\hat{x}_{\text{P-EM-TV}} = \arg\min_{x \in \mathbb{R}^N} \frac{1}{2}\left\|\frac{x - y}{\sqrt{y}}\right\|_2^2 + \lambda\|x\|_{\text{TV}}. \qquad (8)$$

The proposed approach to solve the minimization problem is based on a dual approach using a characterization of subgradients of total variations as divergences of vector fields with constrained norm, resulting in a projected gradient algorithm. Examples on images with Poisson statistics are provided in [5] using a weighted ℓ_2 data-fidelity model.

4) MIDAL: The Multiplicative Image Denoising by Augmented Lagrangian (MIDAL) algorithm proposed by [6] was specifically designed to remove multiplicative noise components from a piecewise-constant field. The noise is assumed to be positive and to follow a known probability distribution (here Gamma distribution for speckle noise) with mean 1 and standard deviation σ_n. The method is designed in three steps: Log-transform of the signal and formulation of the unconstrained MAP estimator with TV regularization; reformulation as a constrained problem via variable splitting; and solving via an augmented Lagrangian method (in their implementation the alternating direction method of multipliers (ADMM) [22] was used). The algorithmic formulation, along with the noise model and the log-transform, leads to an optimization problem with guaranteed convergence and a unique solution. The noisy image is modeled as $y = nx$ and log-transformed to manipulate $\log(y) = \log(n) + \log(x)$, which we write

$\underline{y} = \underline{n} + \underline{x}$. Note that this approach is used in many techniques that aim at denoising images corrupted with multiplicative noise [23]–[25]. The new variable $\underline{x}$ is split into a pair of variables $(u, \underline{x})$. In the case where $n \sim \text{Gamma}(1, \sigma_n)$, the proposed regularized MAP estimator becomes:

$$\hat{x}_{MIDAL} = \exp(\hat{\underline{x}})$$
$$(\hat{\underline{x}}, \hat{u}) = \underset{\underline{x}, u}{\arg \min} \, L(\underline{x}, u)$$
$$\text{s.t. } \underline{x} = u,$$

with

$$L(\underline{x}, u) = \frac{1}{\sigma_n^2} \sum_{s \in \Omega} \left(x_s + e^{y_s - x_s} \right) + \lambda ||u||_{TV}.$$

where λ is a regularization parameter. This method requires the *a priori* knowledge of the noise standard deviation σ_n and the authors showed that λ should decrease as the noise level increases for optimal denoising performance.

B. Non-Local Filtering

Non-local filtering constitutes a very popular and powerful family of denoising methods, exploiting the fusion of patches from over the whole image instead of just filtering locally. They are based on the assumption that real images are non-locally repetitive (i.e. self-similar). We review here some of the most popular non-local denoising methods, which have generated among the best results in the case of complex non-additive noise.

1) NLM: The original non-local means (NLM) algorithm was introduced by Buades *et al.* in [7]. The denoising principle is that, for each noisy patch, there exists, somewhere in the same image, other patches representing the same structure, shape or texture, that can then be aggregated altogether to remove the noise. Formally, the denoised image is generated via the following operation:

$$\forall s \in \Omega, \quad \hat{x}_{\text{NLM}}(s) = \frac{\sum_{t \in \Omega} w_{s,t} y(t)}{\sum_{t \in \Omega} w_{s,t}} \quad (9)$$

with $w_{s,t} = \exp\left(-\frac{F_{s,t}}{h^2}\right)$,

$$F_{s,t} = \sum_{b \in \mathcal{N}} f\left(y(s+b), y(t+b)\right),$$

where h is a filtering parameter, $\mathcal{N}$ is a *patch* (i.e. a neighborhood of fixed size and centered around the pixel of interest), and f is a similarity metric used to compare noisy patches. In the case of additive white Gaussian noise, f is chosen as the ℓ_2 norm and the denoised estimate is computed as:

$$\forall s \in \Omega, \quad \hat{x}_{\text{NLM}}(s) = \frac{1}{Z(s)} \sum_{t \in \Omega} y(t) e^{-\frac{\|y(\mathcal{N}_s) - y(\mathcal{N}_t)\|_2^2}{h^2}}$$

with $Z(s) = \sum_{t \in \Omega} e^{-\frac{\|y(\mathcal{N}_s) - y(\mathcal{N}_t)\|_2^2}{h^2}}$

where $\mathcal{N}_k$ denotes a neighborhood of fixed size and centered at the pixel k, and h is a parameter that controls the decay of

the weights as a function of the ℓ_2 norm (note that we used notations from [7]).

For microscopic images, if there are many similar cells on a background, the NLM algorithm will perform very well [7]. However, in the case of images with only few cells, with singular objects, or with Poisson measures, the method is less suited, suffering from the lack of similar patches. It also has problems with large areas separated by a sharp edge, called the *noise halo* effect in [8] and observed in our results as well.

2) NLM-Poisson: Deledalle *et al.* have proposed in [8] a formulation of the NLM to denoise images corrupted by Poisson noise. They add two modifications to Eq. (9), where $\hat{x}_{NLM-Poisson}$ is defined as $\hat{x}_{NLM}$, but:

$$\text{with } w_{s,t} = \exp\left(-\frac{F_{s,t}}{h_1^2} - \frac{G_{s,t}}{h_1^2}\right),$$
$$F_{s,t} = \sum_{b \in \mathcal{N}} f\left(y(s+b), y(t+b)\right)$$
$$\text{and } G_{s,t} = \sum_{b \in \mathcal{N}} g\left(\hat{\theta}(s+b), \hat{\theta}(t+b)\right)$$

with θ an *a-priori* estimate of the noise free image $\hat{x}$ and (h_1, h_2) two filtering parameters. Using an image prior has been suggested as a means of improving NLM performance in images with very low SNR. It is typically inferred with local convolution filters such as averaging or Gaussian filters. The following similarity metrics are proposed:

$$\forall s, t \in \Omega$$
$$f\left(y(s), y(t)\right) = y(s) \log y(s) + y(t) \log y(t)$$
$$- \left(y(s) + y(t)\right) \log\left(\frac{y(s) + y(t)}{2}\right).$$
$$g\left(\hat{\theta}(s), \hat{\theta}(t)\right) = \left(\hat{\theta}(s) - \hat{\theta}(t)\right) \log \frac{\hat{\theta}(s)}{\hat{\theta}(t)}.$$

which correspond to the Poisson likelihood ratio and the Kullback-Leibler divergence metrics. The parameters (h_1, h_2) influence greatly the denoising performance. The authors propose a method for tuning them using a Poisson unbiased risk estimator (PURE) criterion and its optimization via Newton's gradient descent. These parameters end up being mainly influenced by the level of noise in the image. This approach suffers from two limitations: first, the parameter tuning phase makes the method quite slow to run; second, the noise variance parameter of the method is estimated automatically, and hence the user cannot enforce manually the level of regularization, which is needed for most methods in case of strong noise variance.

3) NLPCA: Building up on non-local (NL) patch aggregation and Poisson likelihood measures, Salmon *et al.* have proposed in [9] to denoise groups of patches via joint optimization of dictionaries and projection coefficients, which constitute the principal component analysis (PCA) part of the NLPCA denoising method. The algorithm is decomposed into five steps: aggregation of patches for all pixels, clustering of patches, denoising of patches in individual clusters via dictionary decomposition, fusion of denoised patches within

clusters and backprojection of denoised patches to form the denoised image. Denote by $Y \in \mathbb{R}^{N \times P}$ the collection of N overlapping patches (size $\sqrt{P} \times \sqrt{P}$) of the noisy image y. Y is organized so that $Y_{i,:}$ denotes the ith patch in a row-vector form. The first step of the method consists in clustering patches into K clusters using some K-means type of algorithm (e.g. Poisson K-means), that enable working with an *a priori* fixed small dictionary size within each cluster Y^k of size N_k (in their work they use $l = 4$ elements per dictionary and $K = 14$ clusters). The second step of the method is to define dictionaries $V_k \in \mathbb{R}^{l \times P}$ and corresponding projection coefficients $U_k \in \mathbb{R}^{N_k \times l}$ in an exponential form so that positivity of the denoised estimates $\hat{x}_k$ of patches in the k^{th} cluster is guaranteed. Using the Poisson likelihood, the kth PCA-based partial estimator $\hat{x}_k \in \mathbb{R}^{N_k}$ is obtained via the following optimization:

$$\hat{x}_k = \exp(U_k^* V_k^*),$$
$$\text{with } (U_k^*, V_k^*) \in \underset{(U,V) \in \mathbb{R}^{N \times l} \times \mathbb{R}^{l \times P}}{\arg \min} L(U_k, V_k),$$
$$\text{where } L(U_k, V_k) = \sum_{i=1}^{N_k} \sum_{j=1}^{P} \exp(U_k V_k)_{i,j} - Y_{i,j}^k (U_k V_k)_{i,j}$$

Finally, the denoised image $\hat{x}$ is obtained by fusing all the partial estimators for a given pixel (averaging is used in the implementation of the authors). They also investigated a refined version called NLSPCA, where S stands for *sparse*, adding a ℓ_1 penalty term on the coefficients U in the loss function $L(U, V)$. This method requires two parameters: number of patch clusters and number of dictionary elements per cluster. This method was designed specifically for low-photon counting statistics and ends up being very sensitive to the quality of the patch clustering step. This observation led the authors to recommend to perform this clustering on a pre-denoised version of the image in case of high noise. Results can suffer from smoothing artifacts as illustrated in our results on synthetic images.

4) PNL Wiener: The Wiener filter [26] is a very popular tool for image restoration in the presence of noise. It was derived as the optimal linear filter solution to the linear minimum mean-squared-error (LMMSE) problem. In the presence of white Gaussian noise, its expression is very simple while leading to high-quality denoising performance. Bindilatti *et al.* [10], [27] have proposed to adapt the Wiener filter for images with Poisson statistics, exploiting a non-local weighted parameter estimation. The Poisson Non-local Wiener estimator (PNLW) is obtained as follows:

$$\forall s \in \Omega, \quad \hat{x}_{PNLW}(s) = \bar{y}(s) + \left(\frac{\sigma_y^2(s) - \bar{y}(s)}{\sigma_y^2(s)} \right)^{\beta} \times (y(s) - \bar{y}(s))$$

where $\bar{y}(s) = \sum_{t \in \mathcal{N}_s} w_{s,t} y(t)$ and $\sigma_y^2 = \left(\sum_{t \in \mathcal{N}_s} w_{s,t} y(t)^2 \right) - \left(\sum_{t \in \mathcal{N}_s} w_{s,t} y(t) \right)^2$ are non-local estimators of the mean and variance of y, $\mathcal{N}_s$ is a search region centered on the pixel s, and β is a parameter introduced to control *a priori* the denoising strength. The non-local weights $w_{s,t}$ are obtained

using the Kullback-Leibler distance d_{KL} as:

$$w_{s,t} = \frac{1}{W_s} exp \left(\frac{-\sum_{q \in \mathcal{N}} d_{KL}(\lambda_{\mathcal{N}_t}(q), \lambda_{\mathcal{N}_s}(q))}{\gamma \lambda(s)^2} \right)$$

where λ are the approximate pixel-wise Poisson parameter values, obtained via non-local denoising of y in [27], $\mathcal{N}$ is a patch neighborhood, and corresponding patch values of λ centered at pixel s are denoted $\lambda_{\mathcal{N}_s}$. W_s is a normalization parameter, so that the $w_{s,t}$ sum up to 1, and γ is a method parameter controlling the rate of decay of the weights. The two parameters of the method are empirically tuned in [10] and [27].

C. Sparse Filtering

1) Analysis K-SVD: Introduced in [28] and [29], the K-Singular Value Decomposition (K-SVD) method exploits the theory of dictionary learning to denoise images. The principle is the following: extract a large number ($N_{patches} = 1000$ in this paper) of fixed size patches ($n_{patches} = 5 \times 5$ pixels in their implementation) of the noisy image y, stored in a matrix Y. Then, denoise each patch using sparse coding [30] on a given dictionary D. Finally, average the overlapping patches to get a denoised estimator $\hat{x}_{KSVD}$. Several algorithms have been developed following the same workflow, with different dictionary constructions or optimization methods in the sparse coding step (see for instance [31] or [32] in the context of Poisson data). In this work, we propose to study the Analysis K-SVD method [11], which improves the original K-SVD method by including both backward greedy (BG) and optimized backward greedy (OPG) algorithms, and a penalty function to take into account some prior information given on the data, like TV sparsity for example.

Given a training set Y of patches obtained from the noisy image y, and an initial dictionary D_0, the method constructs a dictionary D adapted to the image, and a denoised estimator $\hat{x}_{A-KSVD}$ of the true image x. The algorithm uses an iterative scheme, each iteration containing two steps: first, find for each patch Y_i which rows of the current dictionary $D_{current}$ represent it in the sparsest way, and second, update each row of $D_{current}$ to represent in the sparsest way each of the patches associated to it. Once convergence is reached, overlapping patches are aggregated to reconstruct $\hat{x}_{A-KSVD}$. Although this technique shows high performances for large sets of images, and has the versatility to be adapted to different types of noise and different regularization constraints, it is not really adapted to routine work in biological microscopy. Indeed, the learning step is extremely long, almost 4 minutes for one image and 1000 patches only in the training set, which is not compatible with real life experiments. In addition, quality of the results depends strongly on the training set, hence limiting capacity in the case of extremely noisy data.

2) Wavelet Soft-Thresholding: Wavelet soft-thresholding for denoising was introduced by [12], exploiting wavelet functions to decompose a noisy image into subbands at various scales and making the assumption that the noise is mainly encoded in the smaller coefficients. Formally, if $(\mathcal{W}, \mathcal{W}^{-1})$ denote the wavelet transform direct and inverse operators (i.e. analysis and synthesis), then the denoised estimate $\hat{x}_{\text{WST}}$ of the noisy

observation y is obtained as:

$$\hat{x}_{\text{WST}} = \mathcal{W}^{-1}\left(D\left[\mathcal{W}(y)\right]\right) \tag{10}$$

where D is the soft-thresholding operator defined as:

$$D[w] = \text{sign}(w)\ \max(|w| - \lambda_{\text{wav}}, 0) \tag{11}$$

with λ_{wav} the threshold parameter which is typically set based on *a priori* knowledge of the noise standard deviation σ_n and the dimension N of the image. Examples include the universal VisuShrink threshold [33], or the more evolved SureShrink procedure [34]. We used the SureShrink procedure in our experiments. This method is straightforward to implement and extremely fast. Results provided by wavelet soft-thresholding are often of good quality, but suffer from major artifacts, in particular around edges of piecewise constant structures, where oscillations tend to appear.

3) BM3D: Also exploiting non-local patches, the Block-Matching and 3D Filtering (BM3D) method was introduced by Dabov *et al.* in [13] and remains the most used method, often serving as state-of-the-art for all imaging modalities. The method is decomposed into three steps, pixel-wise: grouping similar patches from the original noisy image in stacks, denoising patch stacks in some transform domain (process called collaborative filtering), and aggregation of denoised patches to form the denoised image via weighted-averaging. The pipeline is actually run twice: a first time to estimate a pre-denoised image using wavelet transform and hard thresholding on noisy patches as collaborative filtering and a second time using Wiener filtering on both noisy and pre-denoised patches as collaborative filtering. The method involves several parameters that control the build up of stacks of similar patches and the thresholds of wavelet coefficients. As reported in [35], empirical optimal value for the wavelet threshold is quite insensitive to noise level but parameters controlling the patch stacks have optimal values that are greatly variable across noise levels. Note that other methods have been developed recently following the multi-layer structure approach of BM3D, combining it for instance with connected random fields (see [36]).

D. Variance Stabilizing Transforms

As an alternative to process the signal directly as a Poisson signal, it is possible to use a variance-stabilizing transform (VST) to transform the observed signal. Indeed, for Poisson signals, the variance of $y(s)$ is proportional to the intensity λ_s. It is possible to stabilize the variance of y, noted σ_y via a transformation $z = T(y)$ such that σ_z is asymptotically constant (e.g. $\sigma_z = 1$), irrespective of the values of λ_s.

The most used form of VST transform is: $T(y) = b \times sign(y+c) \times |y+c|^{1/2}$, where $c \in \mathbb{R}$ controls the convergence rate of z toward a Normal distribution with a stable variance as λ_s becomes large [37]. The most popular VST is the Anscombe transform for which $c = \sqrt{3/8}$, and has been used in recent works like [36]. In addition, some work has been done to adapt the Anscombe transform to the case where an image is perturbed with a mixed Poisson-Gaussian noise model [38]. In both cases, the inverse transform required after

denoising the transformed version is a source of instabilities and artefacts that have a major impact on the quality of the denoising process, and tend to oversmooth images in the case of low photon acquisitions [39].

III. PROPOSED FUSION OF SPARSE RECONSTRUCTIONS (FSR) DENOISING ALGORITHM

In the context of this work, which is centered on cellular bioimaging, we make the assumption that these images are piecewise constant (e.g. images consisting of cells with nuclei, cytoplasm and small objects, on a uniform background). The TV semi-norm (see Eq.(4)) is well suited for such images, with minimal values on constant-intensity areas. Based on the work in [40], we propose a method that combines three denoising strategies: local averaging (via an operator denoted $\mathcal{H}(y(s))$), TV-filtering to generate partial estimators $\hat{x}_k$, and a linear minimal error estimator for global TV optimization. Aggregation of noisy (or partial) estimators has been studied in many image denoising (or enhancement) approaches. Beyond naive averaging, the linear minimum mean-squared-error (LMMSE) estimator has optimal properties when working with correlated random observations or independent noisy observations with a white Gaussian additive noise model, but relies on the inference of covariance parameters between observations and the ground-truth signal. Therefore LMMSE performance is greatly affected by the type and level of the actual noise corrupting the observation, and is non-practical when working with spatially-varying noise. It has however raised interest in the context of non-local patch-based denoising [10], [13], [41]–[43]. In [43], LMMSE is used to aggregate multiple estimators of a pixel value from denoised patches, while taking into account the fact that estimators have different reliability levels (denoising confidence), depending on the similarity of patches inferred on the noisy image. In our case, all the estimators have the same *a priori* confidence, but this confidence varies in space, hence the choice to use spatially varying weigths $\alpha_k(s)$ for the linear combination operator.

We first formulate our general denoising estimator as:

$$\hat{x}_*(s) = \alpha_0(s) \times \mathcal{H}(y(s)) + \mathcal{F}(\{\hat{x}_k(s)\}_{k>0}) \tag{12}$$

If we set $\mathcal{F}$ to be linear (i.e. $\mathcal{F}(\{\hat{x}_k(s)\}_k) = \sum_{k>0} \alpha_k(s) \times \hat{x}_k(s)$) and using the same weights on all partial estimators, we can further simplify our estimator to:

$$\hat{x}_*(s) = \alpha_0(s) \times \mathcal{H}(y(s)) + \alpha_1(s) \times \hat{x}_{\text{mean}}(s)$$

where $\forall s \in \Omega$, $\hat{x}_{\text{mean}}(s) = \dfrac{1}{R}\displaystyle\sum_{k=1}^{R}\hat{x}_k(s)$.

Further simplifying Eq. (12), we set $\alpha_0(s) = \alpha(s)$ and $\alpha_1(s) = 1 - \alpha(s)$ with $\alpha(s) \in [0, 1]$, which leads to:

$$\hat{x}_*(s) = \alpha(s) \times \mathcal{H}(y(s)) + (1 - \alpha(s)) \times \hat{x}_{\text{mean}}(s)$$

The choice of $\alpha(s)$ is made so as to get $\|\hat{x}_*(s)\|_{TV}$ closer to the ground-truth TV than $\|\hat{x}_{\text{mean}}\|_{TV}$ via the following

computation, developed here in 1D for the sake of simplicity of notations:

$$\|\hat{x}_*(s)\|_{TV}^{1D}$$
$$= \left| \partial\alpha(s) \times \mathcal{H}(y(s)) + \alpha(s) \times \partial\mathcal{H}(y(s)) \right.$$
$$\left. + \partial(1-\alpha(s)) \times \hat{x}_{\mathrm{mean}}(s) + (1-\alpha(s)) \times \partial\hat{x}_{\mathrm{mean}}(s) \right|$$
$$= \left| \partial\hat{x}_{\mathrm{mean}}(s) + \alpha(s) \times (\partial\mathcal{H}(y(s)) - \partial\hat{x}_{\mathrm{mean}}(s)) \right.$$
$$\left. + \partial\alpha(s) \times (\mathcal{H}(y(s)) - \hat{x}_{\mathrm{mean}}(s)) \right|$$

The first term corresponds to $\|\hat{x}_{\mathrm{mean}}\|_{TV}$. Three assumptions are used to get $\|\hat{x}_*(s)\|_{TV}$ closer to the ground-truth TV than $\|\hat{x}_{\mathrm{mean}}\|_{TV}$: (1) $\hat{x}_{\mathrm{mean}}(s)$ is an approximation of the ground-truth with equal or higher contrast; (2) $\alpha(s)$ is close to a Dirac function with non-zero values only at true edge locations; (3) $\mathcal{H}(y(s))$ generates an approximation of the ground-truth with reduced contrast. Under these assumptions, the second term decreases the first term by a factor $\alpha(s)$ which is non-zero only at edges, while the third term cancels where $\alpha(s)$ is non-zero thanks to the contrast difference $\partial\mathcal{H}(y(s)) < \partial\hat{x}_{\mathrm{mean}}(s)$ and the Dirac derivative having opposite signs across edges. Assumptions are met if we use respectively (1) TV-regularized CS-based reconstructions from the noisy observation to generate the $\hat{x}_k(s)$; (2) a local-averaging filter for $\mathcal{H}$; (3) $\alpha(s) = \hat{x}_{\mathrm{std}}^*(s) = \sqrt{\hat{x}_{\mathrm{var}}(s)}/\|\sqrt{\hat{x}_{\mathrm{var}}(s)}\|_{\infty}$, with $\hat{x}_{\mathrm{var}}$, called the variance map, defined as: $\forall s \in \Omega$,

$$\hat{x}_{\mathrm{var}}(s) = \frac{1}{R}\sum_{k=1}^{R}(\hat{x}_k(s) - \hat{x}_{\mathrm{mean}}(s))^2$$

Regarding the choice (1), we exploit the denoising capacity of the Compressed Sensing (CS) framework [44]–[46] when passing a noisy observation through a Fourier-based sensing operator Φ defined as a random partial Fourier transform and using TV-regularization. If we denote $\mathcal{F} = \left(\frac{1}{\sqrt{N}}\exp\left[\frac{-2i\pi pq}{N}\right]\right)_{0 \leq p,q \leq N-1}$, the discrete Fourier transform matrix, then the sensing operator is written $\Phi = \Sigma\mathcal{F}$, where $\Sigma \in \{0,1\}^{M \times N}$ is a selection matrix, and M is the number of coefficients selected among the total N pixels of the image. We need to define the sampling rate $\tau = \frac{M}{N} \in [0,1]$ as an input parameter (see [47] for a detailed study of the suitable methods of selection of coefficients in the Fourier domain). A denoised estimator $\hat{x}$ of the true signal x can be obtained from the noisy observation y by solving the following convex optimization problem:

$$\hat{x} = \underset{x \in \mathbb{C}^N}{\mathrm{argmin}}\|x\|_{\mathrm{TV}} \text{ s.t. } \|\Phi x - \tilde{y}\|_2 \leq \epsilon \qquad (13)$$

where $\tilde{y} = \Phi y$ is the collection of subsampled Fourier coefficients of the noisy image y, and ϵ is a scalar input parameter related to the energy of the corrupting noise (discussed below). Note that, following the work of [44] and [48], the CS framework guarantees convergence to a stable and accurate estimate $\hat{x}$ of x under the assumption of an additive noise model. One strength of the proposed method is to exploit multiple estimators to yield high-quality results for various types of noise.
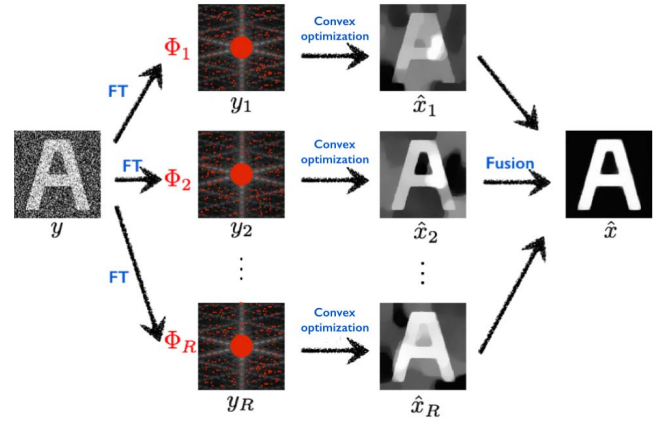

Fig. 1. Generation of partial estimators. From a noisy image y, a given number of low-sampled measurement vectors y_k are generated, by taking the Fourier transform (FT) of y and selecting a subset of the Fourier coefficients (Φ_k). Then, each y_k is used to produce an estimator $\hat{x}_k$ of the original signal through a convex optimization reconstruction scheme. Finally, all the $\hat{x}_k$ are combined into an estimator $\hat{x}$ using one of the proposed fusion techniques.
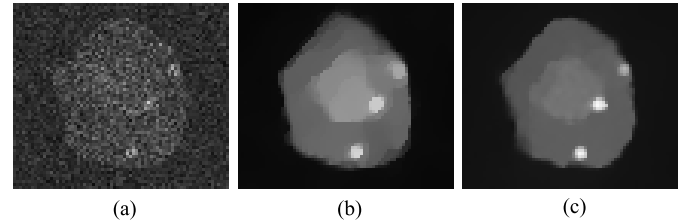

(a)   (b)   (c)

Fig. 2. Illustration of the visual quality of the partial estimators $\hat{x}_k$ and their fusion with our method. The left image is a corrupted version of the *Synth* image using mixed Poisson-Gaussian noise (see Fig.3). The middle image shows one of the $\hat{x}_k$ obtained by solving Eq. (15). The right image is obtained from $R = 3$ reconstructions, using the LWC fusion method. (a) y, PSNR: 13.7. (b) $\hat{x}_k$, PSNR: 27.9. (c) $\hat{x}_{LWC}$, PSNR: 29.3.

This leads to our proposed Linear Weighted Combination (LWC) denoised estimator defined as, $\forall s \in \Omega$:

$$\hat{x}_{LWC}(s) = \hat{x}_{\mathrm{std}}^*(s) \times \mathcal{H}(y(s)) + (1 - \hat{x}_{\mathrm{std}}^*(s)) \times \hat{x}_{\mathrm{mean}}(s)$$
$$(14)$$

where y is the noisy observation and $\mathcal{H}$ is the median filter.

We now detail the generation of the multiple partial estimators $\hat{x}_k$ and introduce two alternative fusion strategies, implemented for comparison purpose and to test if the general denoising concept introduced in Eq. (12) can be extended to alternative weighting strategies.

A. Generation of Multiple Sparsity-Based Partial Estimators

The proposed aggregation of CS-based partial estimators is composed of three steps (see Fig.1), with generated estimators illustrated in Fig. 2:

1) Generation of Random Measurements Vectors in the Fourier Domain: In the case of microscopic images, information concerning details of the structures is located in the high-frequency Fourier coefficients which are much more degraded by noise than the low-frequency coefficients which encode structural information of the objects inside the image [49]. The generation of random sampling patterns in the Fourier

domain and its impact on the reconstruction has been studied extensively in [47] and [50] in the context of Compressed Sensing. In this paper, we consider a dense low-pass sampling (selection of all the coefficients with frequency $\nu < \nu_c$), completed with a random sampling of high-frequency coefficients, with an overall sampling rate of $\tau \in [0, 1]$ (see [51]). From the single noisy observation y, we generate R sets of measures via linear operation: $y_k = \Phi_k y$, where $\Phi_k = \Sigma_k \mathcal{F}$ is constructed by zeroing through random measurement matrices Σ_k most coefficients in the 2D Fourier transform matrix $\mathcal{F}$. The influence and the setting of the three parameters τ, ν_c and R are studied in [52]. We emphasize here that the total quantity of samples used is not exactly $R \times \tau$, as the information in the central disk with radius ν_c is redundant. In this paper, we used $R = 3$, $\tau = 0.2$, $\nu_c = 0.3$, for a total sampling rate of $\sim 35\%$.

2) Reconstruction of Partial Estimators Through Convex Optimization: We solve the following TV-based convex optimization problems to recover R partial estimators $\hat{x}_k$ from the measurement vectors y_k:

$$\hat{x}_k = \underset{x \in \mathbb{R}^N}{\arg\min} \|x\|_{\text{TV}} \quad \text{s.t.} \quad \|\Phi_k x - y_k\|_2 \leq \epsilon \quad (15)$$

where ϵ is a noise-dependent parameter defined in [48] as: $\epsilon = \gamma_\epsilon \sigma_n \sqrt{\tau N + 2\sqrt{2\tau N}}$. In this formula, σ_n represents the standard deviation of the additive Gaussian component of the noise, which we obtain using the cumulant method [53] on a uniform region of the noisy image (e.g. the background) and γ_ϵ is a gain that we introduce to handle specific image types (see below Section IV). We used the NESTA algorithm [48] to solve the optimization problem, based on the recommendations made in the comparative study done in [54].

B. Alternative Fusions Of Sparse Reconstructions (FSR)

The partial estimators obtained by solving (15) are already denoised versions of the initial image y, as illustrated on Fig. 2. Qualitatively, these estimates are different random draws of the same density law, corresponding to TV-constrained estimations of the ground-truth image. We detail here two alternative fusion approaches to use in Eq. (12) with $\alpha_0 = 0$ for comparison and generalization purposes.

1) Exponentially Weighted Aggregate: In [55] Exponentially Weighted Aggregate (EWA) is an aggregation method that was designed to reach optimal average risk, in the context of non-parametric statistical regression. It has been used for image denoising in [56] to fuse "weakly" denoised patches into a single one. Using the formulation from [56], for each partial estimator $\hat{x}_k$, we denote $r_k(s) = |y(s) - \hat{x}_k(s)|^2 - \sigma_n^2$ the risk of the estimator $\hat{x}_k$ at pixel s, assuming a zero-mean additive white Gaussian noise with standard deviation σ_n.

The EWA aggregator is defined as:

$$\forall s \in \Omega, \quad \hat{x}_{\text{EWA}}(s) = \sum_{k=1}^R \theta_k(s) \hat{x}_k(s)$$

with

$$\theta_k(s) = \frac{\exp(-|r_k(s)|/\beta)\pi_k(s)}{\sum_{i=1}^R \exp(-|r_i(s)|/\beta)\pi_i(s)}$$

where $\beta > 0$ is typically called a temperature parameter and $\{\pi_i(s)\}_{i=1,.,R}$ is drafted from a probability distribution function $\pi(s)$ and is used to put prior weights on the estimators. In our work we define π as a uniform random distribution and $\beta = 0.01$. This aggregator considers the partial estimators $\hat{x}_k$ as independent observations of the same true image x, and aggregates them using exponential weights. Note that, as r_k decreases, the corresponding weight θ_k increases, which meets the MSE-minimization objective.

2) Fourier Burst Accumulation: In [57], the authors propose a new image deblurring method based on what they call Fourier burst accumulation. In the context of burst mode acquisition of multiple image frames with a camera, they reconstruct a single deblurred image by aggregating a burst of images via taking what is less blurred of each frame to build an image that is sharper and less noisy than all the images in the burst. To do so, they compute a weighted average of the Fourier coefficients of the images in the burst set. Their working assumption is that the partial estimators are non-noisy but blurred and that the blurring kernels are all positive and with unit norm, leading to no amplification of spectral values. In our context, the partial estimators $\hat{x}_k$ are derived by sacrificing some high-frequency components and using TV spatial regularization. We can therefore make the same assumptions. Formally, the FBA aggregator is defined as:

$$\forall s \in \Omega, \quad \hat{x}_{\text{FBA}}(s) = \mathcal{F}^{-1}\left(\sum_{k=1}^R w_k(\zeta) \circ \mathcal{F}(\hat{x}_k)(\zeta)\right)(s),$$

$$(16)$$

$$w_k(\zeta) = \frac{\left|\mathcal{F}(\hat{x}_k)(\zeta)\right|^p}{\sum_{l=1}^R \left|\mathcal{F}(\hat{x}_l)(\zeta)\right|^p} \quad (17)$$

where p is a non-negative integer, $\mathcal{F}$ represents the Fourier transform operator and ζ the frequency index. The weight w_k controls the contribution, at each frequency, of reconstruction k in the final reconstruction $\hat{x}_{FBA}$. This method prioritizes, among all reconstructions, the one that contains most information at a given frequency. The parameter p emphasizes the predominance of the highest Fourier coefficient value. The authors suggest to set its value in the interval $[7 .. 30]$, where the reconstruction error becomes minimal in terms of Mean Squared Error (MSE). In our work, we empirically identified $p = 15$ as the value leading to optimal results.

IV. Experiments and Results

A. Experimental Set Up

1) Dataset: We performed three sets of experiments to compare the state-of-the-art and proposed denoising methods, using the following images, illustrated in Fig.3:

1) (*Pure*) We generated a bi-color flat image (values 0.25 and 0.75, of size 100×100) corrupted with mixed Poisson-Gaussian noise, to evaluate the denoising power of each method on pure noise fields, before testing them on more complex images. We display the reconstructed images in Fig. 4 for $\lambda = 9$ and $\sigma = 0.15$.

2) (*Synth.*) We used a synthetic cell image [58], composed of a constant background and several cells (containing a
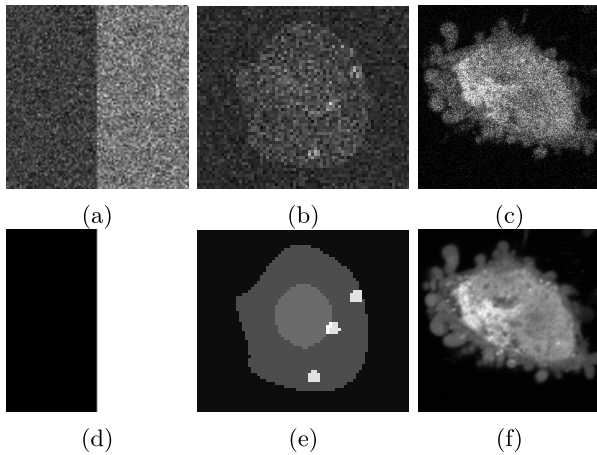
Fig. 3. Images used for the evaluation of the denoising methods: (a, d) *Pure.*, (b, e) *Synth.*, (c, f) *Hela*. The synthetic images are illustrated with realistic mixed Poisson-Gaussian noise with parameters $\lambda = 9$ and $\sigma = 0.15$. (a) MPG noise. (b) MPG noise. (c) Short exp. (d) Clean. (e) Clean. (f) Long exp.

cytoplasm, nucleus and several intracellular objects). For visibility, we display on Figures a zoom on a single of these cells. Matrix size is (256×256) pixels. We tested each method on the *Synth* cell image in the case of white Gaussian noise, pure Poisson noise and mixed Poisson-Gaussian noise, for several noise levels ($0.05 \leq \sigma \leq 0.2$ and $6 \leq \lambda \leq 100$). We compare in Fig. 5 denoising performance for all methods, when varying the noise level using mixed Poisson-Gaussian noise. We display in Fig. 6 the reconstructed images for $\lambda = 9$ and $\sigma = 0.15$, corresponding to the visual characteristics of the noisy *Hela* acquisition.

3) (*Hela*) We imaged, with a confocal microscope, a HeLa cell image labeled with a Cell Tracker CMFDA dye. This marker is strong enough, and the laser weak enough, to avoid photobleaching in this experiment. Because of the different sources of noise present during the acquisition process, these images are corrupted with mixed Poisson-Gaussian noise, which decreases as we increase exposure time. To obtain a noisy realistic test image, i.e. a very noisy observation dominated by photon noise, we first acquired an image with a short exposure time (0.1 second). To obtain a ground truth noise-free image, we then acquired a second image with a long time exposure (10 seconds), using the fact that the accumulation of a large number of photons results in a high quality image with almost no noise, at the expense however of the sample integrity (see results on Fig. 7, and find additional results on other HeLa cell images in the supplementary files).

2) Parameterization of the Methods: We summarize in Table II the code sources and parameterization details for each method. When possible, we used the code provided by the authors for the state-of-the-art methods. We list in Table III the websites used to retrieve the codes, when available. For the synthetic images, we used the ground-truth noise information to set *a priori* input noise parameters. For the real microscopic image, since we do not know the actual characteristics of the
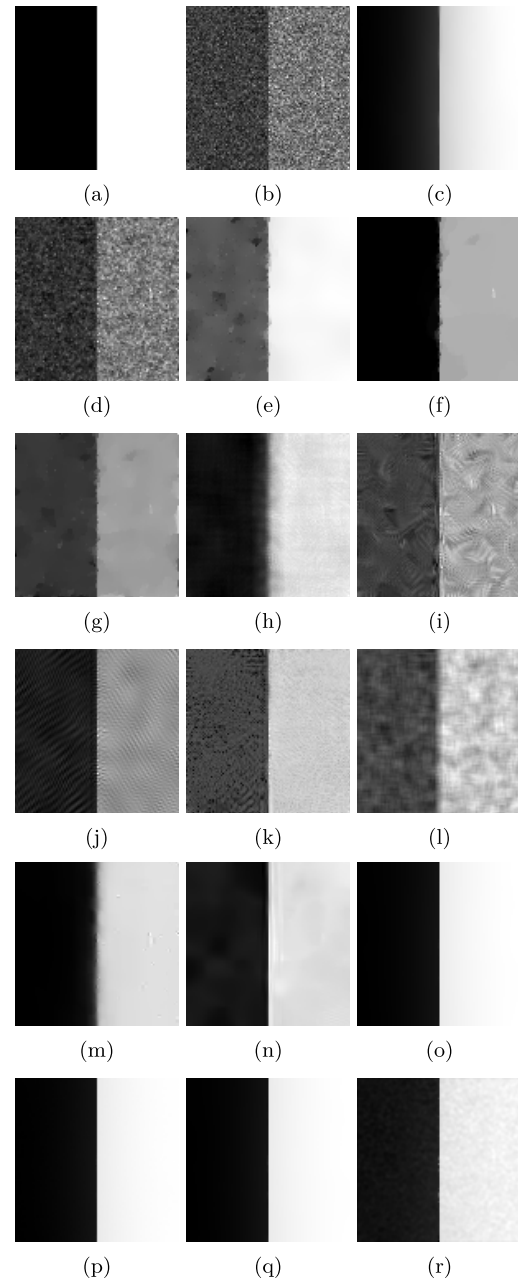

Fig. 4. Comparison of the denoising methods on the *Pure.* image, with values rescaled to [0, 1]. The image y is corrupted with a mixed Poisson-Gaussian noise of parameters $\lambda = 9$ and $\sigma = 0.15$. (a) Ground truth. (b) Noisy. PSNR: 10.81. SSIM: 0.255. (c) TV. PSNR: 23.82. SSIM: 0.916. Comp. time: 1.3s. (d) TV-ICE. PSNR: 12.40. SSIM: 0.410. Comp. time: 1.7s. (e) TV-MAP Poisson. PSNR: 19.65. SSIM: 0.908. Comp. time: 0.55s. (f) Poisson EM-TV. PSNR: 25.87. SSIM: 0.975. Comp. time: 0.56s. (g) MIDAL. PSNR: 15.48. SSIM: 0.886. Comp. time: 0.36s. (h) NLM. PSNR: 21.36. SSIM: 0.889. Comp. time: 1.0s. (i) NLM-Poisson. PSNR: 15.35. SSIM: 0.744. Comp. time: 0.24s. (j) NLPCA. PSNR: 17.54. SSIM: 0.725. Comp. time: 1.2s. (k) PNLW. PSNR: 17.62. SSIM: 0.655. Comp. time: 1.8s. (l) Analysis K-SVD. PSNR: 13.60. SSIM: 0.574. Comp. time: 277s. (m) Wavelet thresh. PSNR: 21.74. SSIM: 0.943. Comp. time: 5.3s. (n) BM3D. PSNR: 23.52. SSIM: 0.974. **Comp. time: 0.18s.** (o) FSR-Mean. PSNR: 32.52. SSIM: 0.996. Comp. time: 1.8s. (p) FSR-EWA. **PSNR: 32.93. SSIM: 0.997.** Comp. time: 1.9s. (q) FSR-FBA. PSNR: 32.65. SSIM: 0.996. Comp. time: 2.1s. (r) FSR-LWC. PSNR: 22.11. SSIM: 0.884. Comp. time: 1.9s.

noise, we estimated the noise variance inside a patch in the background of the image using the cumulant approach [53]. For each method, we optimized the different parameters
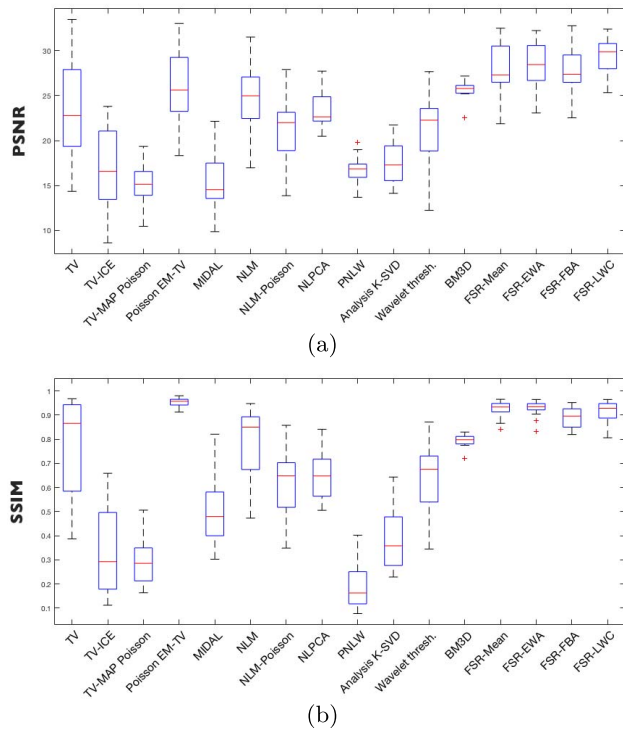
Fig. 5. Box plots of the PSNR and SSIM measures obtained with each denoising method, over 16 levels of mixed Poisson-Gaussian noise ($0.05 \leq \sigma \leq 0.2$ and $6 \leq \lambda \leq 100$). Each box plot emcompasses the PSNR or SSIM values over all noise levels. Blue rectangles englobe 50% of the measures, black tails represent first and last quartiles, horizontal red bars indicate the mean values, and red crosses represent outliers. (a) PSNR values. (b) SSIM values.

in order to reconstruct the best image with respect to the PSNR metric. The optimized parameters vary with the type of image and with the noise model. We report in Table II the parameter values tuned on the *Synth* image perturbed with white Gaussian noise with $\sigma = 0.1$.

3) Denoising Performance Metrics: We evaluate and compare denoising performance on two criteria: visual quality and computation performance. Visual quality is evaluated via measuring PSNR which is the most commonly used metric in the image denoising literature and SSIM [59] which is advocated to better reflect the human eye perception. Computation times are reported running the denoising experiments on a PC workstation 2.93 GHz Quad-core CPU with 8GB of RAM. All codes were implemented in Matlab©, some (indicated with a * in Table II) exploiting compiled mex files.

B. Results

Visual results for the three types of images are summarized in Fig. 4 for the *Pure* noise image, in Fig. 6 for the *Synth* cell image both corrupted with mixed Poisson Gaussian noise ($\sigma = 0.15$, $\lambda = 9$) and in Fig. 7 for the experimental *Hela* image. Comments on the visual results are gathered in Table IV. Several artifacts seen on the *Pure* and *Synth* images replicate findings in [42], including the *rare patch effect* at edges and presence of patch borders in flat regions of the image.
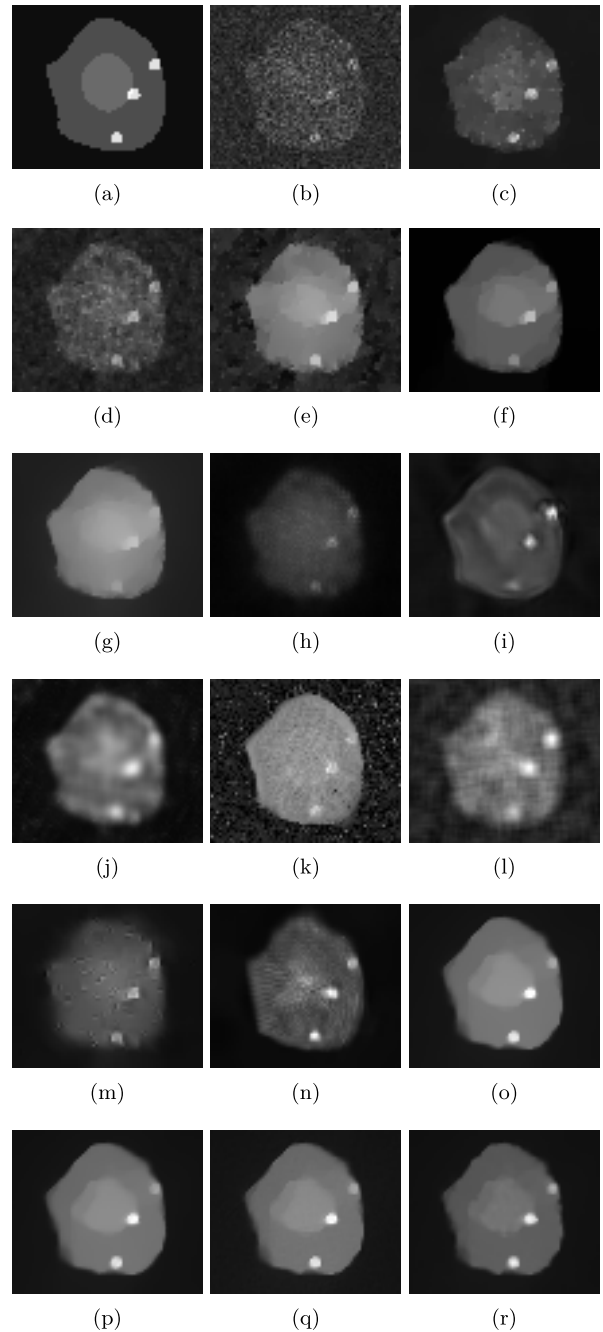

Fig. 6. Comparison of the denoising methods on the *Synth.* cell image - zoom on a single cell, with values rescaled to $[0, 1]$. The image y is corrupted with a mixed Poisson-Gaussian noise of parameters $\lambda = 9$ and $\sigma = 0.15$. (a) Ground truth. (b) Noisy. PSNR: 13.71. SSIM: 0.112. (c) TV. PSNR: 23.79. SSIM: 0.770. Comp. time: 1.1s. (d) TV-ICE. PSNR: 16.85. SSIM: 0.337. Comp. time: 10.6s. (e) TV-MAP Poisson. PSNR: 15.09. SSIM: 0.296. Comp. time: 2.3s. (f) Poisson EM-TV. PSNR: 26.01. SSIM: 0.903. Comp. time: 2.4s. (g) MIDAL. PSNR: 15.22. SSIM: 0.493. **Comp. time: 0.7s**. (h) NLM. PSNR: 24.44. SSIM: 0.777. Comp. time: 15.3s. (i) NLM-Poisson. PSNR: 21.26. SSIM: 0.626. Comp. time: 44.2s. (j) NLPCA. PSNR: 23.47. SSIM: 0.652. Comp. time: 15.9s. (k) PNLW. PSNR: 16.81. SSIM: 0.191. Comp. time: 27.9s. (l) Analysis K-SVD. PSNR: 17.47. SSIM: 0.378. Comp. time: 480s. (m) Wavelet thresh. PSNR: 25.97. SSIM: 0.651. Comp. time: 5.0s. (n) BM3D. PSNR: 25.67. SSIM: 0.795. Comp. time: 1.2s. (o) FSR-Mean. PSNR: 28.13. SSIM: 0.925. Comp. time: 14.2s. (p) FSR-EWA. PSNR: 28.63. *SSIM: 0.927*. Comp. time: 16.8s. (q) FSR-FBA. PSNR: 27.93. SSIM: 0.888. Comp. time: 14.3s. (r) FSR-LWC. **PSNR: 29.30**. SSIM: 0.911. Comp. time: 14.3s.
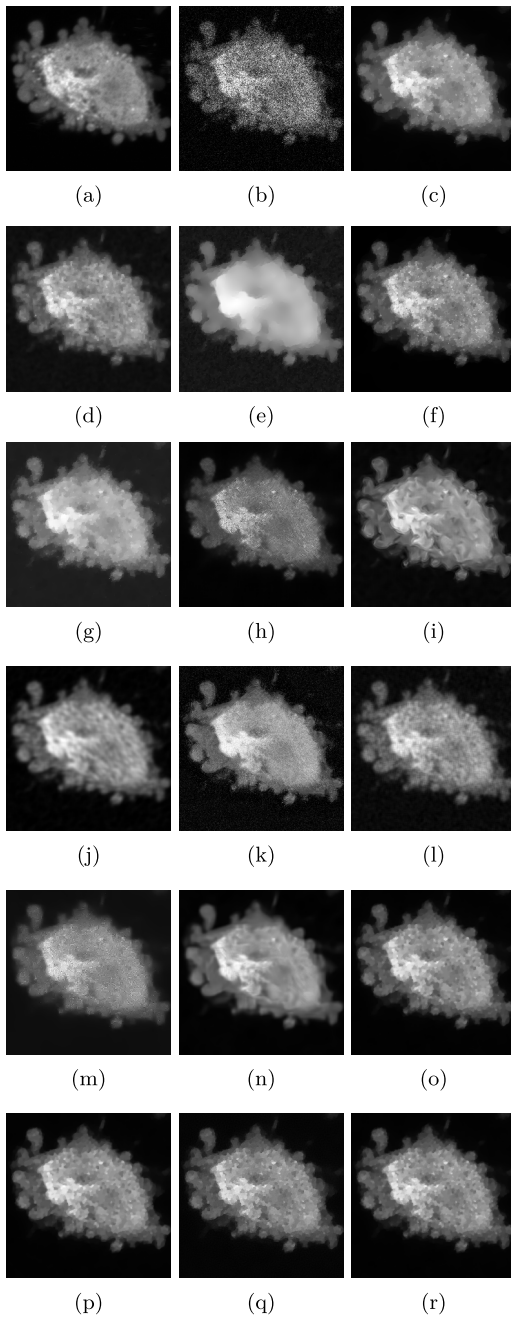
Fig. 7. Comparison of the denoising methods on the *Hela.* image, with values rescaled to [0, 1]. (a) Long exposure. (b) Short exposure. PSNR: 17.72. SSIM: 0.094. (c) TV. PSNR: 29.99. SSIM: 0.832. Comp. time: 4.9s. (d) TV-ICE. PSNR: 24.95. SSIM: 0.391. Comp. time: 62.7s. (e) TV-MAP Poisson. PSNR: 19.92. SSIM: 0.337. Comp. time: 11.1s. (f) Poisson EM-TV. PSNR: 30.53. SSIM: 0.753. Comp. time: 16.4s. (g) MIDAL. PSNR: 18.88. SSIM: 0.309. Comp. time: 10.4s. (h) NLM. PSNR: 27.56. SSIM: 0.804. Comp. time: 73.6s. (i) NLM-Poisson. PSNR: 28.46. SSIM: 0.571. Comp. time: 11.1s. (j) NLPCA. PSNR: 30.19. SSIM: 0.793. Comp. time: 46.7s. (k) PNLW. PSNR: 24.30. SSIM: 0.254. Comp. time: 82.8s. (l) Analysis K-SVD. PSNR: 26.62. SSIM: 0.473. Comp. time: 1208s. (m) Wavelet thresh. PSNR: 25.45. SSIM: 0.394. Comp. time: 4.8s. (n) BM3D. PSNR: 30.63. SSIM: 0.827. **Comp. time: 3.5s** (o) FSR-Mean. PSNR: 31.07. SSIM: 0.837. Comp. time: 33.1s. (p) FSR-EWA. PSNR: 31.05. SSIM: 0.831. Comp. time: 35.5s (q) FSR-FBA. PSNR: 29.73. SSIM: 0.638. Comp. time: 33.3s. (r) FSR-LWC. **PSNR: 31.13**. **SSIM: 0.842**. Comp. time: 33.3s

In the captions of the reconstructed synthetic images, we report the PSNR and SSIM values obtained for each method as well as computation times. We highlight in bold the best PSNR and SSIM values obtained with our method, and the fastest method among all, which is MIDAL for *Synth.* and BM3D for *Pure* and *Hela.*

We report in Fig. 5 results from testing on the *Synth.* image the robustness of denoising performance over 16 different levels of mixed Poisson-Gaussian noise ($0.05 \leq \sigma \leq 0.2$ and $6 \leq \lambda \leq 100$), with noise-level specific optimized parameters (for the state of the art methods) or keeping parameters fixed (our proposed method). We observe that our proposed method (with various fusion operators) always compares well with the best-performing state-of-the-art method, in terms of PSNR and SSIM, over all noise levels, is robust to noise level (i.e. smaller variations of performance metrics) and have no outlier cases. We can see that TV and Poisson EM-TV methods return the best maxima (corresponding to low noise level cases) but that the performance is not consistent over stronger noise levels. The FSR-LWC method returns the best average PSNR value. Finally, the Poisson EM-TV method returns high SSIM values, but a large variability in terms of PSNR, due to strong localized artefacts, that affect PSNR measures but not SSIM.

Results obtained on the *Pure* image (cf. Fig.4) enable to clearly capture visual characteristics of individual denoising methods. Images have been rescaled to [0, 1], so that the ground-truth appears black and white instead of grey (true values being 0.25 and 0.75). Consequently, the loss of contrast present in some results is due to the presence of few outliers that control the overall rescaling. We decided to emphasize the presence of such outliers, observed in many methods, via such rescaling. As expected, the TV filter gives very satisfactory results on this type of pure-noise images, as we enforced the regularization parameter to obtain a TV-sparse result. Similarly, wavelet-based images, using Haar wavelets, return close to state-of-the-art results. On the other hand, the non-local methods fail to denoise this image, as there is no texture to exploit here. Overall, this figure provides visual evidence on how each non-local denoising method deals with noise and suffers from generated artefactual structures.

Regarding the *Synth* image, some methods like NLM, NLPCA, PNLW, KSVD, wavelet thresholding, and BM3D fail to return a good quality image. Our proposed method performs well with all fusion operators, with a visual quality comparable to TV and Poisson EM-TV. MIDAL and TV-ICE have intermediate quality, with contrast between structures poorly recovered. The TV results show some clear staircase artefacts, not present in Poisson EM-TV results nor with our method. Our method returns the highest PSNR value with EWA fusion operator, and runs in about ~14s while Poisson EM-TV runs in 2.4s. But Poisson EM-TV requires the setting of a maximum number of iterations which is empirical and requires manual adjustment when performing tests on various types and level of noises.

On the *Hela* image, all methods return interpretable images, but TV-ICE, NLM and wavelet thresholding have the lowest visual quality, with clear loss of structural details. Other methods are all comparable and again our proposed method returns the highest PSNR value but this time with the LWC fusion operator. In terms of computation time, BM3D is the fastest algorithm with ~4s, while our method runs in ~30s.

TABLE IV
HIGHLIGHTS OF VISUAL RESULTS OBTAINED WITH THE TESTED DENOISING METHODS

Acronym	Ref.	Quality	Artefacts
TV	[2]	Very efficient on piecewise constant images.	Staircase effects on real images.
TV-ICE	[3]	Efficient on small structures.	Deals poorly with wide homogeneous areas.
TV-MAP Poisson	[4],[61]	Outperforms classical TV-ICE on wide areas.	Loss of contrast on $Synth$ image.
Poisson EM-TV	[5]	Outperforms TV, especially for Poisson noise.	Staircase effect.
MIDAL	[6]	Overall good performance.	Loss of texture and modification of overall contrast.
NLM	[7]	High performance on textured images.	Loss of contrast on $Synth$ image and appearance of texture patterns on $Synth$ image.
NLM-Poisson	[8]	Very efficient result on $Synth$ image with Poisson noise.	Performance degraded in presence of Gaussian noise.
NLPCA	[9]	Efficient on $Hela$ image.	Visible textures on $Synth$ image, blurred structures on $Synth$ image.
PNLW	[10]	Efficient on pure Poisson noise.	Loss of contrast on mixed noise.
Analysis K-SVD	[11]	High performance on textured images.	Extremely slow, not adapted to high noise.
Wavelet thresh.	[12]	Intermediate performance on $Hela$ image.	Artifacts around edges.
BM3D	[13]	Overall best on $Hela$.	Visual artifacts on the $Synth$ image.
FSR		Overall high performance over all image types and noise models.	Small staircase effects and blur at the edges of $Synth.$ image.

More generally, methods that have been designed specifically for Poisson noise, such as NLM-Poisson, NLPCA and PNLW, perform poorly on images corrupted with mixed Poisson-Gaussian noise. This observation confirms the fact that such methods make strong assumption on the noise model, and that the mathematical approach chosen to solve the denoising problem is then less robust to other types of noise. However, these methods have the particularity of being non-local, and give good results on biological images that contain textures.

V. DISCUSSION AND CONCLUSION

In this work, we have introduced a sparsity-based denoising method, combining partial reconstructions via dedicated fusion operators, which surpassed 12 state-of-the-art denoising methods in terms of PSNR in three sets of experiments on microscopic images corrupted with high and complex noise. Our proposed method systematically provided excellent quantitative and qualitative results both on synthetic and real images, with the advantages on being based on a single set of parameter values and a simple additive noise model to set a parameter that weights TV regularization versus data fidelity in the CS cost metric. Our method shares several attributes with the tested state-of-the-art methods. The use of multiple sparsity-based reconstructions is akin to wavelet-thresholding, which combines a low-pass approximation of the noisy observation with filtered high-pass sparse details. Indeed, our Fourier sampling scheme ensures the presence of a low-pass approximation, enriched during the aggregation procedure with denoised details reconstructed via CS and TV-regularization. We do not exploit a multi-scale structure in our proposed method, since we are able to gather multiple versions of the details at the initial resolution, but this could be envisaged easily via the exploitation of multiple values of the cutoff frequency ν_c in a future extension. Denoising power

in our method relies on the use of TV regularization, as in many of the best-performing denoising methods tested in this work. The proposed FSR method is not limited to images with patch redundancies across the field of view (as for NLM and BM3D), but it does exploit redundant spatial information from multiple overlapping sub-samplings of the Fourier domain during fusion. It does not rely on an *a priori* model of the type of noise that corrupts the observation, as required by some specialized methods tested.

Finally, we tested the concept of using an *a priori* estimate of denoised image to guide the denoising in the *EWA* fusion operator. Such oracle is used in the NLM-Poisson method which is based on Stein's unbiased risk estimate (SURE) minimization [60]. Single parameterization and robustness to noise types are highly relevant in the field of microscopic imaging, as such images are perturbed with highly variable levels and types of mixed Poisson-Gaussian noise which are complex to model *a priori* or via a Poisson likelihood term to optimize. Another benefit of our approach is the use of only a subset of the samples in the Fourier domain (typically 35%), which could lead to great potentials in data compression or data sub-sampling during acquisition. The proposed FSR method is highly adaptive and can be optimized in the future via investigations of alternative fusion strategies, Fourier sampling strategies and more sophisticated exploitation of the variance map. Finally, to facilitate the work of the community and build a collective know-how, we share all the image data and denoising results of this *in extenso* study with the community (http://icy.bioimageanalysis. org/matlab/state-of-the-art-denoising) toward two goals: documenting parameter optimization approaches on shared codes for some of the denoising techniques, and build-up a database of synthetic and real microscopic images with ground-truth to further test denoising methods on a shared benchmark.

References

[1] L. Schermelleh, R. Heintzmann, and H. Leonhardt, "A guide to super-resolution fluorescence microscopy," *J. Cell Biol.*, vol. 190, no. 2, pp. 165–175, 2010.

[2] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Phys. D, Nonlinear Phenomena*, vol. 60, nos. 1–4, pp. 259–268, 1992.

[3] C. Louchet and L. Moisan, "Total variation denoising using iterated conditional expectation," in *Proc. 22nd IEEE Eur. Signal Process. Conf. (EUSIPCO)*, Sep. 2014, pp. 1592–1596.

[4] R. Abergel, C. Louchet, L. Moisan, and T. Zeng, "Total variation restoration of images corrupted by Poisson noise with iterated conditional expectations," in *Proc. Int. Conf. Scale Space Variational Methods Comput. Vis.*, 2015, pp. 178–190.

[5] A. Sawatzky, C. Brune, J. Müller, and M. Burger, "Total variation processing of images with Poisson statistics," in *Computer Analysis of Images and Patterns (CAIP)*. Berlin, Germany: Springer, 2009, pp. 533–540.

[6] J. M. Bioucas-Dias and M. A. T. Figueiredo, "Multiplicative noise removal using variable splitting and constrained optimization," *IEEE Trans. Image Process.*, vol. 19, no. 7, pp. 1720–1730, Jul. 2010.

[7] A. Buades, B. Coll, and J.-M. Morel, "A review of image denoising algorithms, with a new one," *Multiscale Model. Simul.*, vol. 4, no. 2, pp. 490–530, 2005.

[8] C.-A. Deledalle, F. Tupin, and L. Denis, "Poisson NL means: Unsupervised non local means for Poisson noise," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Sep. 2010, pp. 801–804.

[9] J. Salmon, Z. Harmany, C.-A. Deledalle, and R. Willett, "Poisson noise reduction with non-local PCA," *J. Math. Imag. Vis.*, vol. 48, no. 2, pp. 279–294, 2014.

[10] A. A. Bindilatti, M. A. C. Vieira, and N. D. A. Mascarenhas, "Poisson Wiener filtering with non-local weighted parameter estimation using stochastic distances," *Signal Process.*, vol. 144, pp. 68–76, Mar. 2018.

[11] R. Rubinstein, T. Peleg, and M. Elad, "Analysis K-SVD: A dictionary-learning algorithm for the analysis sparse model," *IEEE Trans. Signal Process.*, vol. 61, no. 3, pp. 661–677, Feb. 2013.

[12] D. L. Donoho, "De-noising by soft-thresholding," *IEEE Trans. Inf. Theory*, vol. 41, no. 3, pp. 613–627, May 1995.

[13] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE Trans. Image Process.*, vol. 16, no. 8, pp. 2080–2095, Aug. 2007.

[14] A. Boyat and B. Joshi, "A review paper: Noise models in digital image processing," *Signal Image Process., Int. J.*, vol. 6, no. 2, pp. 63–75, 2015.

[15] P. Sarder and A. Nehorai, "Deconvolution methods for 3-D fluorescence microscopy images," *IEEE Signal Process. Mag.*, vol. 23, no. 3, pp. 32–45, May 2006.

[16] A. Jezierska, C. Chaux, J.-C. Pesquet, and H. Talbot, "An EM approach for Poisson-Gaussian noise modeling," in *Proc. IEEE Eur. Signal Process. Conf. (EUSIPCO)*, Sep. 2011, pp. 2244–2248.

[17] A. Chambolle, "An algorithm for total variation minimization and applications," *J. Math. Imag. Vis.*, vol. 20, no. 1, pp. 89–97, 2004.

[18] C. Louchet and L. Moisan, "Total variation denoising using posterior expectation," in *Proc. 16th IEEE Eur. Signal Process. Conf. (EUSIPCO)*, Aug. 2008, pp. 1–5.

[19] T. Le, T. Chartrand, and T. J. Asaki, "A variational approach to reconstructing images corrupted by Poisson noise," *J. Math. Imag. Vis.*, vol. 27, no. 3, pp. 257–263, Apr. 2007.

[20] S. Setzer, G. Steidl, and T. Teuber, "Deblurring Poissonian images by split Bregman techniques," *J. Vis. Commun. Image Represent.*, vol. 21, no. 3, pp. 193–199, Apr. 2010.

[21] M. A. T. Figueiredo and J. M. Bioucas-Dias, "Restoration of Poissonian images using alternating direction optimization," *IEEE Trans. Image Process.*, vol. 19, no. 12, pp. 3133–3145, Dec. 2010.

[22] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, Jan. 2011.

[23] L. Rudin, P.-L. Lions, and S. Osher, "Multiplicative denoising and deblurring: Theory and algorithms," in *Geometric Level Set Methods in Imaging, Vision, and Graphics*. New York, NY, USA: Springer, 2003, pp. 103–119.

[24] Y.-M. Huang, M. K. Ng, and Y.-W. Wen, "A new total variation method for multiplicative noise removal," *SIAM J Imag. Sci.*, vol. 2, no. 1, pp. 20–40, 2009.

[25] S. Durand, J. Fadili, and M. Nikolova, "Multiplicative noise removal using L1 fidelity on frame coefficients," *J. Math. Imag. Vis.*, vol. 36, no. 3, pp. 201–226, 2010.

[26] S. M. Kay, *Fundamentals of Statistical Signal Processing, Volume 1: Estimation Theory*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1993.

[27] A. A. Bindilatti and N. D. A. Mascarenhas, "A nonlocal Poisson denoising algorithm based on stochastic distances," *IEEE Signal Process. Lett.*, vol. 20, no. 11, pp. 1010–1013, Nov. 2013.

[28] M. Elad and M. Aharon, "Image denoising via sparse and redundant representations over learned dictionaries," *IEEE Trans. Image Process.*, vol. 15, no. 12, pp. 3736–3745, Dec. 2006.

[29] M. Aharon, M. Elad, and A. Bruckstein, "*K*-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Trans. Signal Process.*, vol. 54, no. 11, pp. 4311–4322, Nov. 2006.

[30] H. Lee, A. Battle, R. Raina, and A. Y. Ng, "Efficient sparse coding algorithms," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 19. 2007, pp. 801–808.

[31] R. Rubinstein, M. Zibulevsky, and M. Elad, "Efficient implementation of the K-SVD algorithm using batch orthogonal matching pursuit," *CS Technion*, vol. 40, no. 8, pp. 1–15, 2008.

[32] L. Ma, L. Moisan, J. Yu, and T. Zeng, "A dictionary learning approach for Poisson image deblurring," *IEEE Trans. Med. Imag.*, vol. 32, no. 7, pp. 1277–1289, Jul. 2013.

[33] D. L. Donoho and I. M. Johnstone, "Ideal spatial adaptation by wavelet shrinkage," *Biometrika*, vol. 81, no. 3, pp. 425–455, 1994.

[34] D. L. Donoho and I. M. Johnstone, "Adapting to unknown smoothness via wavelet shrinkage," *J. Amer. Statist. Assoc.*, vol. 90, no. 432, pp. 1200–1224, 1995.

[35] M. Lebrun, "An analysis and implementation of the BM3D image denoising method," *Image Process. On Line*, vol. 2, pp. 175–213, Aug. 2012.

[36] S. A. Haider *et al.*, "Fluorescence microscopy image noise reduction using a stochastically-connected random field model," *Sci. Rep.*, vol. 6, p. 20640, Feb. 2016.

[37] B. Zhang, M. J. Fadili, J.-L. Starck, and J.-C. Olivo-Marin, "Multiscale variance-stabilizing transform for mixed-Poisson-Gaussian processes and its applications in bioimaging," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, vol. 6. Oct. 2007, pp. VI-233–VI-236.

[38] M. Makitalo and A. Foi, "Optimal inversion of the generalized Anscombe transformation for Poisson-Gaussian noise," *IEEE Trans. Image Process.*, vol. 22, no. 1, pp. 91–103, Jan. 2013.

[39] R. Willett, "Multiscale analysis of photon-limited astronomical images," in *Proc. Stat. Challenges Modern Astron. IV ASP Conf. Ser.*, vol. 371. 2007, p. 247.

[40] M. Marim, E. Angelini, and J.-C. Olivo-Marin, "Denoising in fluorescence microscopy using compressed sensing with multiple reconstructions and non-local merging," in *Proc. IEEE Int. Conf. Eng. Med. Biol. Soc. (EMBS)*, Aug./Sep. 2010, pp. 3394–3397.

[41] P. Chatterjee and P. Milanfar, "Patch-based near-optimal image denoising," *IEEE Trans. Image Process.*, vol. 21, no. 4, pp. 1635–1649, Apr. 2012.

[42] J. Salmon and Y. Strozecki, "Patch reprojections for non-local methods," *Signal Process.*, vol. 92, no. 2, pp. 477–489, 2012.

[43] C. V. Angelino, E. Debreuve, and M. Barlaud, "Patch confidence k-nearest neighbors denoising," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Sep. 2010, pp. 1129–1132.

[44] E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.

[45] E. J. Candés, J. K. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Commun. Pure Appl. Math.*, vol. 59, no. 8, pp. 1207–1223, 2006.

[46] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.

[47] Y. LeMontagner, "Algorithmic solutions toward applications of compressed sensing for optical imaging," Ph.D. dissertation, Dept. Signal Image Process., Télécom ParisTech, Paris, France, 2013.

[48] S. Becker, J. Bobin, and E. J. Candès, "NESTA: A fast and accurate first-order method for sparse recovery," *SIAM J. Imag. Sci.*, vol. 4, no. 1, pp. 1–39, 2011.

[49] Z. Wang and G. R. Arce, "Variable density compressed image sampling," *IEEE Trans. Image Process.*, vol. 19, no. 1, pp. 264–270, Jan. 2010.

[50] C. Poon, "On the role of total variation in compressed sensing," *SIAM J. Imag. Sci.*, vol. 8, no. 1, pp. 682–720, 2015.

[51] Y.-C. Kim, S. S. Narayanan, and K. S. Nayak, "Accelerated three-dimensional upper airway MRI using compressed sensing," *Magn. Reson. Med.*, vol. 61, no. 6, pp. 1434–1440, 2009.

[52] W. Meiniel, E. Angelini, and J.-C. Olivo-Marin, "Image denoising by adaptive compressed sensing reconstructions and fusions," in *Proc. SPIE*, vol. 9597, p. 95970X, Sep. 2015.

[53] C. Rose and M. D. Smith, *MathStatica: Mathematical Statistics With Mathematica*. Heidelberg, Germany: Springer, 2002, pp. 437–442.

[54] Y. Le Montagner, E. Angelini, and J.-C. Olivo-Marin, "Comparison of reconstruction algorithms in compressed sensing applied to biological imaging," in *Proc. IEEE Int. Symp. Biomed. Imaging, Nano Macro (ISBI)*, Mar./Apr. 2011, pp. 105–108.

[55] A. Dalalyan and A. B. Tsybakov, "Aggregation by exponential weighting, sharp PAC-Bayesian bounds and sparsity," *Mach. Learn.*, vol. 72, nos. 1–2, pp. 39–61, 2008.

[56] C. Kervrann, "PEWA: Patch-based exponentially weighted aggregation for image denoising," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 2150–2158.

[57] M. Delbracio and G. Sapiro, "Hand-held video deblurring via efficient Fourier aggregation," *IEEE Trans. Comput. Imag.*, vol. 1, no. 4, pp. 270–283, Dec. 2015.

[58] P. Ruusuvuori, A. Lehmussola, J. Selinummi, T. Rajala, H. Huttunen, and O. Yli-Harja, "Benchmark set of synthetic images for validating cell image analysis algorithms," in *Proc. IEEE Eur. Signal Process. Conf. (EUSIPCO)*, Aug. 2008, pp. 1–5.

[59] Z. Wang and A. C. Bovik, "Mean squared error: Love it or leave it? A new look at signal fidelity measures," *IEEE Signal Process. Mag.*, vol. 26, no. 1, pp. 98–117, Jan. 2009.

[60] C. M. Stein, "Estimation of the mean of a multivariate normal distribution," *Ann. Statist.*, vol. 9, no. 6, pp. 1135–1151, 1981.

[61] R. Abergel, "Several mathematical models and fast algorithms for image processing," Ph.D. dissertation, Dept. Appl. Math., Univ. Paris-Descartes, Paris, France, 2016.

William Meiniel received the M.Sc. degree in applied mathematics from ENS Cachan, France, in 2014. From 2014 to 2018, he was with the BioImage Analysis Unit, Institut Pasteur, France, and the Image-Data-Signal Department, Télécom ParisTech, Institut Mines-Télécom, Paris. His research interests are focused on compressed sensing, image denoising, fluorescence microscopy, and superresolution.

Jean-Christophe Olivo-Marin (F'12) received the Ph.D. and H.D.R. degrees in optics and signal processing from the Institut d'Optique Théorique et Appliquée, University of Paris-Orsay, France. He was a Staff Scientist with the European Molecular Biology Laboratory, Heidelberg, from 1990 to 1998. He was the Chair of the Cell Biology and Infection Department, Institut Pasteur, from 2010 to 2014, where he was the CTO and the Director of the Center for Innovation and Technological Research from 2015 to 2018. He is currently the Head of the BioImage Analysis Unit, Institut Pasteur, and the Director of the Institute Carnot Pasteur Microbes et Santé. His research interests are in image analysis of multidimensional microscopy images, computer vision and motion analysis for cellular dynamics, and in mathematical approaches for biological imaging. He is a fellow of SPIE, an IEEE Signal Processing Society Distinguished Lecturer. He was the Chair of the IEEE SPS Bio Imaging and Signal Processing Technical Committee from 2009 to 2011, the Chair of the IEEE International Symposium on Biomedical Imaging Steering Committee from 2014 to 2016, the General Chair of the IEEE International Symposium on Biomedical Imaging in 2008. He was a Senior Area Editor of the IEEE SIGNAL PROCESSING LETTERS from 2013 to 2015.

Elsa D. Angelini (M'01–SM'12) received the B.Sc. degree in engineering from the Ecole Centrale de Nantes, France, in 1996, and the Ph.D. degree in biomedical engineering from Columbia University, USA, in 2003. She is currently the Senior Data Scientist and a Co-Leader of the ITMAT Data Science Group, NIHR Imperial Biomedical Research Centre, Imperial College London, U.K., and an Associate Professor (on leave) with the Department of Image-Data-Signal at Telecom ParisTech, Universite Paris-Saclay, France. She also holds an affiliated appointment with the Department of Biomedical Engineering, Columbia University, and the Department of Biomedical Engineering, as a Co-Director of the Heffner Biomedical Imaging Laboratory. Her primary area of investigation is medical and biological image processing. She is a member of the EMBS Technical Committee on BioImaging and Image Processing and a past member of the IEEE SPS Technical Committee on BioImaging and Signal Processing. She was the General Chair of the IEEE ISBI'15 conference. She also served in the organizing committees of ISBI'08 and MICCAI'08. She is currently the Vice President for Technical Activities of the IEEE EMB Society. She is also the Co-Chair of the SPIE Medical Imaging-Image Processing Conference. She served as an Associate Editor for the IEEE TBME. She also serves as an Associate Editor for the IEEE JBHI.