



HAL
open science

Fluctuation-based deconvolution in fluorescence microscopy using plug-and-play denoisers

Vasiliki Stergiopoulou, Subhadip Mukherjee, Luca Calatroni, Laure Blanc-Féraud

► **To cite this version:**

Vasiliki Stergiopoulou, Subhadip Mukherjee, Luca Calatroni, Laure Blanc-Féraud. Fluctuation-based deconvolution in fluorescence microscopy using plug-and-play denoisers. *Scale Space and Variational Methods in Computer Vision*, May 2023, Sardinia, Italy. pp.498-510, 10.1007/978-3-031-31975-4_38 . hal-04038738

HAL Id: hal-04038738

<https://hal.science/hal-04038738v1>

Submitted on 21 Mar 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Fluctuation-based deconvolution in fluorescence microscopy using plug-and-play denoisers

Vasiliki Stergiopoulou¹, Subhadip Mukherjee², Luca Calatroni¹, and Laure Blanc-Féraud¹

¹ UCA, CNRS, INRIA, Laboratoire I3S, 06903, Sophia-Antipolis France
vasiliki.stergiopoulou@i3s.unice.fr, calatroni@i3s.unice.fr,
blancf@i3s.unice.fr

² Department of Computer Science, University of Bath, UK
sm2467@cam.ac.uk

Abstract. The spatial resolution of images of living samples obtained by fluorescence microscopes is physically limited due to the diffraction of visible light, which makes the study of entities of size less than the diffraction barrier (around 200 nm in the x - y plane) very challenging. To overcome this limitation, several deconvolution and super-resolution techniques have been proposed. Within the framework of inverse problems, modern approaches in fluorescence microscopy reconstruct a super-resolved image from a temporal stack of frames by carefully designing suitable hand-crafted sparsity-promoting regularisers. Numerically, such approaches are solved by proximal gradient-based iterative schemes. Aiming at obtaining a reconstruction more adapted to sample geometries (e.g. thin filaments), we adopt a plug-and-play denoising approach with convergence guarantees and replace the proximity operator associated with the explicit image regulariser with an image denoiser (i.e. a pre-trained network) which, upon appropriate training, mimics the action of an implicit prior. To account for the independence of the fluctuations between molecules, the model relies on second-order statistics. The denoiser is then trained on covariance images coming from data representing sequences of fluctuating fluorescent molecules with filament structure. The method is evaluated on both simulated and real fluorescence microscopy images, showing its ability to correctly reconstruct filament structures with high values of peak signal-to-noise ratio (PSNR).

Keywords: Fluorescence microscopy · Image deconvolution · Variational regularisation · Proximal algorithms · Plug-and-Play regularisation.

1 Introduction

In optical microscopy, the highest achievable spatial resolution is governed by some fundamental physical laws related to light propagation and is therefore limited. According to Rayleigh’s criterion, the resolution of an optical microscope is defined as the smallest resolvable distance, i.e. the smallest distance between two point sources so that they can be distinguished in the image. For conventional fluorescence microscopes, this distance is approximately equal to 200 nm in the lateral (x - y) plane. In order to resolve sub-cellular structures of

size smaller than this barrier, several deconvolution and super-resolution techniques have emerged in the literature. Originally developed in the applied fields of chemistry, biology, and biophysics, such techniques can be naturally described in more mathematical terms as regularisation approaches for solving the ill-posed inverse problem considered. A big family of approaches achieving nanometric resolution (around 20 nm of lateral resolution) is known as *Single Molecule Localisation Microscopy* (SMLM) techniques (see [24] for a review). These methods rely on the use of sparse regularisation approaches for reconstructing frames of a temporal sequence of acquisitions where only a few molecules are active at a time. A more hardware-based super-resolution technique achieving a resolution of approximately 60 – 100 nm, is the *Stimulated Emission Depletion* (STED) microscopy approach [12] where the optical blur function (the microscope Point Spread Function, PSF) is depleted by means of suitable devices. While effective, these techniques show several drawbacks: for example, SMLM has long acquisition times while STED requires highly expensive commercial tools. Furthermore, both approaches require special (and expensive) fluorescent molecules able to support the high laser power required.

For overcoming such limitations, different types of approaches exploiting the independent stochastic temporal fluctuations of distinct fluorescent emitters became popular over the last decade. Such approaches rely on the use of both standard microscopes and fluorescent molecules and represent therefore a powerful class of approaches for applications. Some of these approaches are: the *Super-resolution Optical Fluctuation Imaging* (SOFI) approach [5] where the lack of correlation between distinct emitters is exploited by analyzing high-order statistics, the *Super-Resolution Radial Fluctuations* (SRRF) [11] microscopy, where super-resolution is achieved by calculating at each frame the degree of local symmetry and, finally, the *Sparsity-based Super-resolution Correlation Microscopy* (SPARCOM) [25] which models the sparse distribution of the fluorescent molecules via the use of a convex ℓ_1 regularisation applied on the emitters' covariance matrix. To improve the performance of SPARCOM, the Covariance-based ℓ_0 super-resolution microscopy with intensity estimation (COLORME) method [27,29] has been proposed to estimate both molecule positions and intensities, which is a valuable piece of information in several applications, such as, e.g., 3D imaging [28], by means of a two-step procedure relying on hand-crafted sparsity promoting regularisers. The approach further estimates noise statistics and background terms (containing out-of-focus molecules). Both SPARCOM and COLORME rely on the minimization of non-smooth (and possibly non-convex) functionals, for which tailored proximal optimization algorithms [19] based on soft- and hard-thresholding rules have thus been considered.

The applicability of these approaches to more complex geometries is limited due to the hand-crafted sparsity they enforce which creates biases (i.e., punctuated structures) in the reconstruction. This is particularly limiting when continuous curvilinear structures are desired, which is the case in several biological applications. For that, suitable regularisers can indeed be defined [18], with the major limitation of remaining tailored to particular shapes only. With the

intent of developing a flexible regularisation approach suited to adapt to different geometries, we present in the following a data-driven optimization-inspired technique relying on the use of the so-called Plug-and-Play (PnP) approaches [31], which, over the last decade have been proved to represent an efficient framework for solving inverse image restoration problems, see [15] for a review. In this framework, the regulariser is parameterised by a deep neural network that can be trained on simulated data implicitly characterised by desired structures of interest, thus better capturing/promoting their shape after training. Our primary motivations behind using the PnP approach are three-fold: (i) Training the denoiser is independent of the imaging forward operator, which makes the pre-trained denoiser applicable even if the forward operator undergoes some changes, without having to retrain the denoiser from scratch. (ii) The training problem does not require pairs of measured and ground-truth images, unlike supervised machine learning approaches. To train the denoiser, one only needs high-quality ground-truth images and their noisy counterparts (with additive Gaussian noise). (iii) The PnP approaches are rooted in proximal point algorithms, so their convergence can be rigorously studied using results from fixed-point theory and/or convex analysis. This leads to better interpretability of the reconstruction algorithm and results in a principled way of combining knowledge about imaging physics with the available training data.

Contributions: In this work, we leverage the framework of PnP approaches with convergence guarantees [3,13,14] to show good empirical performance on the inverse problem of fluctuation-based image deconvolution presented, e.g., in [25,27,29]. In Section 2, we review the recent advances in the field of PnP approaches for inverse problems, pointing out the convergent scheme we employ. In Section 3, we formulate the covariance-based deconvolution model and formulate its PnP extension, which we called PnP-COLORME in the following. In Section 4 we report several numerical results on both simulated and real data where the advantages of using the PnP reconstruction model are shown in comparison to its model-based counterpart.

2 Plug-and-Play approaches for inverse problems

A standard approach for solving ill-posed inverse problems in imaging consists in solving the optimisation problem:

$$\hat{\mathbf{x}} \in \arg \min_{\mathbf{x} \in \mathbb{R}^{n^2}} \mathcal{F}(\Psi \mathbf{x}; \mathbf{y}) + \lambda \mathcal{R}(\mathbf{x}), \quad \lambda > 0, \quad (1)$$

where, for observed data $\mathbf{y} \in \mathbb{R}^{n^2}$ (being the vectorisation of a 2D image of size $n \times n$) and model operator $\Psi \in \mathbb{R}^{n^2 \times n^2}$, \mathcal{F} denotes a (smooth) data fidelity term and \mathcal{R} a regularisation term encoding prior knowledge on the solution $\hat{\mathbf{x}} \in \mathbb{R}^{n^2}$. Depending on the available prior information (such as sparsity, gradient smoothness, etc.), tailored hand-crafted functions \mathcal{R} can be used. In most cases, \mathcal{R} is non-smooth, and proximal algorithms [19] can be used for solving (1). We

recall that the proximity operator of parameter $\tau > 0$ of a proper, convex and non-smooth function \mathcal{R} is defined by:

$$\text{prox}_{\tau\mathcal{R}}(\mathbf{z}) = \arg \min_{\mathbf{x} \in \mathbb{R}^{n^2}} \mathcal{R}(\mathbf{x}) + \frac{1}{2\tau} \|\mathbf{z} - \mathbf{x}\|_2^2, \quad \mathbf{z} \in \mathbb{R}^{n^2}. \quad (2)$$

Solving (2) corresponds to solving the problem of denoising an image $\mathbf{z} \in \mathbb{R}^{n^2}$ corrupted by an additive white Gaussian noise (AWGN) of constant variance equal to τ . Within a proximal gradient algorithm, (2) can thus be interpreted as a denoising step of the gradient descent iteration $\mathbf{z}^k = \mathbf{x}^k - \tau \nabla \mathcal{F}(\Psi \mathbf{x}^k; \mathbf{y})$ at each iteration $k \geq 0$. This observation inspired the authors of [31] to develop the framework of PnP priors, whose main idea consists in replacing $\text{prox}_{\tau\mathcal{R}}(\cdot)$ with an off-the-shelf image denoiser $D_\sigma(\cdot)$ depending on a parameter $\sigma > 0$ corresponding to a regularisation functional \mathcal{R} whose explicit definition is often not available. In a Bayesian framework, it is indeed possible to explicitly relate Gaussian minimum mean-squared error (MMSE) denoisers $D_\sigma(\cdot)$ with the (unknown) image prior $p(\cdot)$ one would like to model [17] using the Tweedie’s identity: $\sigma^2 \nabla \log p_\sigma(\mathbf{x}) = D_\sigma(\mathbf{x}) - \mathbf{x}$, where p_σ is the convolution of p with a Gaussian smoothing kernel of bandwidth $\sigma > 0$, which makes p_σ smoother (namely, Lipschitz differentiable) than p under mild conditions. As observed in [20, Eq. 74-75], considering a (Gaussian) denoiser residual is in fact a good approximation to the score of the image prior. Along with their Bayesian interpretability, another advantage of PnP approaches is that they allow the use of advanced image denoising models, e.g. denoisers parameterised by convolutional neural networks (CNNs), within the iterative scheme, with impressive representational capabilities. In most cases, the CNN image denoiser D is trained to perform denoising on some pairs of clean-noisy images and can be used afterward for more-general inverse problems (e.g. deblurring, super-resolution, etc.), see [15] for a review. Some state-of-the-art denoisers include image-dependent filtering algorithms such as Block-Matching & 3D filtering (BM3D) [4], Denoising Convolutional Neural Networks (DnCNN) [34] and Dilated-Residual U-Net (DRUNET) deep learning network [33].

These denoisers are typically used in iterative proximal schemes (see, e.g., [16] for a FISTA-type PnP scheme), although they can be flexibly used in other algorithms such as, e.g., the Alternate Directions Method of Multipliers (ADMM) [2], the Douglas-Rachford Splitting (DRS) [6] and the Half-Quadratic Splitting (HQS) [8]. For all these algorithms, a corresponding PnP version can indeed simply be obtained as described above. PnP versions of proximal algorithms have been used to solve image restoration problems such as for example PnP-PGD in [30], PnP-ADMM and PnP-DRS in [21,22] and PnP-HQS in [34,33,3].

In [21] an explicit regularisation by denoising (RED) strategy was designed in terms of an explicit function $\mathcal{R}(\cdot)$ defined, for generic image denoiser D , by:

$$\mathcal{R}(\mathbf{x}) := \frac{1}{2} \mathbf{x}^T (\mathbf{x} - D(\mathbf{x})).$$

Under conditions of local homogeneity, non-expansiveness, and Jacobian symmetry, D was shown to be indeed equivalent to a gradient step on \mathcal{R} [20], that is,

$D(\mathbf{x}) = \mathbf{x} - \nabla \mathcal{R}(\mathbf{x})$. However, as shown in [20], such requirements are unrealistic on the widely-used denoisers mentioned above, as they do not have symmetric Jacobians. In order to overcome this limitation, in [3,13,14], the authors proposed to formulate, similar to RED, a gradient step denoiser of the form:

$$D_\sigma(\mathbf{x}) = \mathbf{x} - \nabla \mathcal{R}_\sigma(\mathbf{x}), \quad (3)$$

where $\mathcal{R}_\sigma : \mathbb{R}^{n^2} \rightarrow \mathbb{R}$ is a scalar function parameterised by a neural network $N_\sigma : \mathbb{R}^{n^2} \rightarrow \mathbb{R}^{n^2}$. Interestingly, under mild structural assumption on D_σ , the authors are able to prove sound convergence guarantees for the underlying non-convex optimisation problem defined in terms of a non-trivial (but explicit) regularisation function $\mathcal{R}(\cdot)$. In the following section we specify the particular problem we are interested in and discuss its PnP extension based on the strategy discussed above.

3 Deconvolution via sparse auto-covariance analysis

We consider the following image formation model considered, e.g., in [25,29,27] to describe, for $t = 1, \dots, T$, $T > 0$, a video of temporal acquisitions $\mathbf{y}_t \in \mathbb{R}^{n^2}$ by standard fluorescent microscopes of true images $\mathbf{x}_t \in \mathbb{R}^{n^2}$:

$$\mathbf{y}_t = \Psi \mathbf{x}_t + \mathbf{b} + \mathbf{n}_t. \quad (4)$$

In (4), $\Psi \in \mathbb{R}^{n^2 \times n^2}$ is a (known) convolution operator associated with the system point spread function (PSF), $\mathbf{b} \in \mathbb{R}^{n^2}$ is a background term and $\mathbf{n}_t \in \mathbb{R}^{n^2}$ is the realisation at time t of an i.i.d. Gaussian noise random vector with unknown variance $s \geq 0$, i.e. $\mathbf{n} \sim \mathcal{N}(0, s \mathbf{Id})$. We look for a deconvolved image $\mathbf{x} \in \mathbb{R}^{n^2}$, defined as $\mathbf{x} = \frac{1}{T} \sum_{t=1}^T \mathbf{x}_t$. In [25,27], a reformulation of the model (4) was done in the covariance domain in order to exploit temporal information. In the following, we proceed similarly but consider a simplified modeling where only auto-covariance vectors are taken into account, thus neglecting cross-terms.

Considering the frames $\{\mathbf{y}_t\}_{t=1}^T$ as T realisations of a random variable \mathbf{y} , the sample auto-covariance (variance) vector $\tilde{\mathbf{r}}_{\mathbf{y}} \in \mathbb{R}^{n^2}$ of \mathbf{y} can be estimated by:

$$\tilde{\mathbf{r}}_{\mathbf{y}} \approx \frac{1}{T-1} \sum_{t=1}^T (\mathbf{y}_t - \bar{\mathbf{y}})^2, \quad (5)$$

where $\bar{\mathbf{y}} = \frac{1}{T} \sum_{t=1}^T \mathbf{y}_t$ denotes the empirical mean. From (4) and (5), the following model thus holds between the auto-covariance vectors:

$$\tilde{\mathbf{r}}_{\mathbf{y}} = \Psi^2 \mathbf{r}_{\mathbf{x}} + \tilde{\mathbf{r}}_{\mathbf{n}}, \quad (6)$$

where $\mathbf{r}_{\mathbf{x}} \in \mathbb{R}^{n^2}$ and $\tilde{\mathbf{r}}_{\mathbf{n}} \in \mathbb{R}^{n^2}$ are the auto-covariance vectors associated to the samples $\{\mathbf{x}_t\}_{t=1}^T$ and $\{\mathbf{n}_t\}_{t=1}^T$, respectively, and where by $\Psi^2 \in \mathbb{R}^{n^2 \times n^2}$ we denote the matrix $\Psi \odot \Psi$ where \odot denotes the point-wise Hadamard product. Finally, note that by assumption $\tilde{\mathbf{r}}_{\mathbf{n}} = s \mathbf{1}$, where $\mathbf{1} = (1, \dots, 1) \in \mathbb{R}^{n^2}$.

Algorithm 1 Model-based and PnP support estimation

Require: $\tilde{\mathbf{r}}_{\mathbf{y}}, \mathbf{r}_{\mathbf{x}}^0 \in \mathbb{R}^{n^2}$ and parameters $(\tau, \lambda > 0$ for model-based, σ for PnP)**repeat**

$$s^{k+1} = \mathbf{1}^T (\tilde{\mathbf{r}}_{\mathbf{y}} - \Psi^2 \mathbf{r}_{\mathbf{x}}^k)$$

$$\mathbf{z}^{k+1} = \mathbf{r}_{\mathbf{x}}^k - \tau \lambda (\Psi^2)^T (\tilde{\mathbf{r}}_{\mathbf{y}} - \Psi^2 \mathbf{r}_{\mathbf{x}}^k - s^{k+1} \mathbf{1})$$

$$\mathbf{r}_{\mathbf{x}}^{k+1} = \begin{cases} \text{prox}_{\tau \lambda R}(\mathbf{z}^{k+1}) & \% \text{ model-based (7)} \\ \mathbf{r}_{\mathbf{x}}^{k+1} = D_{\sigma}(\mathbf{z}^{k+1}) & \% \text{ PnP} \end{cases}$$

until convergence**return** $\Omega := \{i : (\hat{\mathbf{r}}_{\mathbf{x}})_i \neq 0\}, \hat{s}$

Remark 1. Note that, upon reshaping, $\tilde{\mathbf{r}}_{\mathbf{y}} \in \mathbb{R}^{n^2}$ is in fact the second-order SOFI image associated to the stack $\{\mathbf{y}_t\}_{t=1}^T$, which, thanks to the ‘squaring’ of the underlying point spread function, enjoys better spatial resolution in comparison, e.g., to $\bar{\mathbf{y}}$, see [5] for details.

Remark 2. In comparison to the covariance-based modelling in SPARCOM [25] and COLORME [27], model (6) is indeed a simplification. In those papers, the whole sample-covariance matrix $\mathbf{R}_{\mathbf{y}} \in \mathbb{R}^{n^2 \times n^2}$ with main diagonal $\tilde{\mathbf{r}}_{\mathbf{y}}$ was computed. Such a matrix is not diagonal, due to the correlation induced by the PSF. In order to deal with a simplified model and benefit from faster calculations, we consider in the following the relation (6) involving only auto-covariance terms and leave a complete modelling involving also cross terms for future work. The resulting observation model 6 is thus less rich but exact (not approximated).

Based on (6), we are now interested in finding the fluorescent molecule locations and estimate noise information. Namely, we are interested in finding the support of $\mathbf{r}_{\mathbf{x}}$, $\Omega := \{i : (\mathbf{r}_{\mathbf{x}})_i \neq 0\}$ and the unknown noise variance $s \geq 0$. We do so by considering the following minimisation problem:

$$(\hat{\mathbf{r}}_{\mathbf{x}}, \hat{s}) \in \arg \min_{\mathbf{r}_{\mathbf{x}} \geq 0, s \geq 0} \frac{\lambda}{2} \|\tilde{\mathbf{r}}_{\mathbf{y}} - \Psi^2 \mathbf{r}_{\mathbf{x}} - s \mathbf{1}\|_2^2 + \mathcal{R}(\mathbf{r}_{\mathbf{x}}), \quad (7)$$

where $\lambda > 0$ is a regularization parameter and $\mathcal{R}(\cdot)$ is a regularisation term to be defined to enforce desirable properties (sparsity, for instance) of the solution. Problem (7) can be solved by Algorithm 1, where, to improve convergence speed, a global minimisation on s is performed followed by a proximal-gradient step on $\mathbf{r}_{\mathbf{x}}$. An analogous (a priori, slower) algorithm benefiting from theoretical convergence guarantees is the Proximal Alternating Linearized Minimisation (PALM) algorithm whose convergence is studied in [1].

Once Ω and \hat{s} have been computed, following [29,27] a second algorithmic step can be performed to estimate image intensities only in correspondence with the support points in Ω , i.e. by solving:

$$(\hat{\mathbf{x}}, \hat{\mathbf{b}}) \in \arg \min_{\mathbf{x} \in \mathbb{R}_+^{|\Omega|}, \mathbf{b} \in \mathbb{R}_+^2} \frac{1}{2} \|\Psi_{\Omega} \mathbf{x} - (\bar{\mathbf{y}} - \mathbf{b})\|_2^2 + \frac{\mu}{2} \|\nabla_{\Omega} \mathbf{x}\|_2^2 + \frac{\beta}{2} \|\nabla \mathbf{b}\|_2^2, \quad (8)$$

where the data term models the presence of Gaussian noise, $\bar{\mathbf{y}} = \sum_{t=1}^T \mathbf{y}_t$ and $\mu, \beta > 0$ are regularisation parameters. Moreover, $\Psi_\Omega \in \mathbb{R}^{n^2 \times |\Omega|}$ is a matrix whose i -th column is extracted from Ψ for all indexes $i \in \Omega$ and ∇_Ω denotes the discrete gradient operator restricted to points in the support Ω .

A hand-crafted regularisation model \mathcal{R} in (7) introduces reconstruction biases. For example, using the ℓ_1 norm [25] or the continuous exact relaxation of the ℓ_0 pseudo-norm [26]) enforces sparsity by promoting point reconstruction. Solutions thus appear dotted as reconstructed points have a given inter-distance which cannot be decreased [7]. For reconstructing filaments, a solution is to use a regularising term promoting curves. Such method is proposed, e.g., in [18] in an off-the-grid setting, but the numerical aspects are difficult and still under development. To overcome this limitation, in the following section, we propose a Plug-and-Play extension of the approach above where the proximal step is replaced by a denoiser D_σ trained on an appropriate dataset of covariance images representing the geometrical structures of interest.

3.1 Plug-and-play extension

The proximal step naturally appearing when solving problem (7) by proximal gradient algorithms, can be replaced by an off-the-shelf denoiser. To do so, we make use of a proximal gradient step denoiser as proposed by Hurault *et al.* in [14]. In their paper, the authors showed that this choice corresponds indeed to the proximal operator associated to a non-convex smooth function which allows the authors to derive convergence guarantees of the resulting proximal gradient scheme [14, Theorem 4.1]. Note, that differently to the setting proposed in [14], our algorithm processes auto-covariance images due to the model (6) and, along with $\hat{\mathbf{r}}_{\mathbf{x}}$, it provides an estimate \hat{s} of the noise variance by alternate minimisation. We report the iterative scheme in Algorithm 1 and refer in the following to PnP-COLORME to the case when a PnP regulariser is employed.

In [14] the authors considered a denoiser D_σ in the form of a gradient step (3) of a functional $\mathcal{R}_\sigma : \mathbb{R}^{n^2} \rightarrow \mathbb{R}$ with specific properties, e.g., bounded from below, and parameterised by a deep neural network N_σ . Recalling the characterisation of proximity operators [10] introduced by Gribonval & Nikolova, the authors proved in fact that D_σ can be written as proximal operator of a function ϕ_σ defined by:

$$\phi_\sigma(\mathbf{w}) := \mathcal{R}_\sigma(D_\sigma^{-1}(\mathbf{w})) - \frac{1}{2} \|D_\sigma^{-1}(\mathbf{w}) - \mathbf{w}\|_2^2, \quad \mathbf{w} \in \mathbb{R}^{n^2}.$$

The function minimised when employing PnP COLORME reads: $F_\sigma(\mathbf{r}_{\mathbf{x}}, s) := \frac{1}{2} \|\mathbf{r}_{\mathbf{y}} - \Psi^2 \mathbf{r}_{\mathbf{x}} - s \mathbf{1}\|_2^2 + \phi_\sigma(\mathbf{r}_{\mathbf{x}})$, which, after recalling that $\mathbf{r}^k = D_\sigma(\mathbf{z}^k)$ at each k , can be written as:

$$F_\sigma(\mathbf{r}_{\mathbf{x}}^k, s^k) = \frac{1}{2} \|\mathbf{r}_{\mathbf{y}} - \Psi^2 \mathbf{r}_{\mathbf{x}}^k - s^k \mathbf{1}\|_2^2 + \mathcal{R}_\sigma(\mathbf{z}^k) - \frac{1}{2} \|\mathbf{z}^k - \mathbf{r}_{\mathbf{x}}^k\|_2^2. \quad (9)$$

In [14, Theorem 4.1] the authors show that thanks to the structure of F_σ , the PnP proximal gradient scheme converges indeed to a stationary point of F_σ , whose decay can be indeed assessed throughout the iterations.

Note that, the regularisation parameter $\lambda > 0$ appearing in (1) to regulate the strength of the regularisation term \mathcal{R} has been replaced by the hyperparameter σ in (9). Intuitively, the value of σ should correspond to the variance of AWGN appearing in the gradient steps of the proximal gradient algorithm 1, hence its tuning is not straightforward. As discussed in [32], a possible remedy for avoiding a time-consuming parameter tuning consists in introducing a rescaling parameter whose setting is easier than σ .

4 Numerical results

We now present some results obtained by using PnP-COLORME on temporal sequences of blurred and noisy data. A natural extension to the actual problem of super-resolution where $\Psi = \mathbf{M}_q \mathbf{H} \in \mathbb{R}^{m^2 \times n^2}$ with $\mathbf{H} \in \mathbb{R}^{n^2 \times n^2}$ is PSF convolution matrix and $\mathbf{M}_q \in \mathbb{R}^{m^2 \times n^2}$ is a downsampling operator with $n = qm, q > 1$, is left for future work.

To train the denoiser D_σ we created a dataset composed of clean and noisy image pairs. The geometrical features of the images in this dataset should be the same as the one of the images to restore. Differently from other methods, the proposed algorithm works with a model formulated in the covariance domain, so that the denoiser takes as an input noisy sample auto-covariance matrices of a fluctuating temporal sequence of images. Hence, to create the dataset we first started by creating different spatial patterns (thin filaments) shown in Figure 1 where the emitters have different positions in the continuous grid. Such patterns are the superposition, after rotations with different angles, of the ground truth spatial pattern provided in the MT0 microtubule training dataset uploaded for SMLM 2016 challenge³. Then, we used the fluctuation model discussed in [9] to simulate temporal fluctuations and create a temporal stack of $T = 500$ frames for each spatial pattern. Two exemplar frames of one temporal stack of images are reported in Figures 2a and 2b. For each temporal stack of images, we could therefore calculate the temporal auto-covariance image (see Figure 2c) corresponding to one instance of the clean images \mathbf{r}_x^{GT} in our dataset. To create now its noisy version we added Gaussian noise $\boldsymbol{\eta}$ with constant variance σ^2 , $\boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{Id})$, with σ following a uniform distribution, $\sigma \sim \mathcal{U}(\sigma_1, \sigma_2)$. We remark that since the noise in the covariance data comes from additive Gaussian noise on the individual frames, its actual distribution is indeed χ^2 . However, since the number of the degrees of freedom is high (as $T = 500$), the distribution can be approximated by a Gaussian distribution. In our experiments, after normalising \mathbf{r}_x^{GT} with maximum value equal to 1, we select $\sigma_1 = \epsilon \ll 1$ and $\sigma_2 = 50/255$.

Training was performed following the procedure in [14] and using the code available on the authors' GitHub repository⁴. For the neural network $N_\sigma(\cdot)$ used to parameterise the denoiser (see (3)), we used DRUNet, a CNN proposed in [33]. For training, we used 500 pairs of clean-noise auto-covariance images

³ <https://srm.epfl.ch/Challenge/ChallengeSimulatedData>

⁴ <https://github.com/samuro95/Prox-PnP>




Fig. 1: Simulated spatial patterns




Fig. 2: (a-b) Two different frames of a simulated fluctuating stack made from the first spatial pattern from Figure 1, (c) The auto-covariance image \mathbf{r}_x^{GT} estimated from the whole temporal sequence.

and 100 for validation. The network was trained using 1215 epochs via ADAM optimization and batch size equal to 16. In the following experiments, the choice $\tau = 1$ and $\lambda = 0.99$ in Algorithm 1 was performed to guarantee convergence, see [14, Section 4.1] for details.

4.1 Simulated data

We first apply PnP COLORME to simulated data presented in Figure 3. The PSF used to generate the data has a FWHM equal to 176.6 nm, the pixel size is equal to 25 nm and the images have a size of 256×256 pixels.

Thanks to its training, we observe that the proposed approach is able to capture the filaments' geometry fairly well. We observe that in comparison to the ground truth support in Fig. 3c, the reconstruction in Fig. 3g is rather accurate. For the evaluation of the localization precision the Jaccard Index (JI) has been used. It is a quantity in $[0, 1]$ computed as the ratio between correct detections (CD) and the total (correct, false negatives false positive) detections, i.e. $\text{JI} := \text{CD} / (\text{CD} + \text{FN} + \text{FP})$, up to a tolerance $\delta > 0$, measured in nm (see, e.g., [23]). For the reconstruction in Fig. 3g, the tolerance precision was chosen $\delta = 40$ nm. Moreover, by solving (8), intensities can also be estimated with high precision, see Fig. 3h. However, for the challenging dataset in Figure 3, the appearance of small artefacts (e.g. incorrect duplication of filaments) due to the




Fig. 3: (a) Mean of the acquired temporal sequence, (b) First frame (c) The ground truth support (d) The ground truth intensity image (e) Evolution of cost function F_σ in (9) (f) the evolution of $\min_{i \leq k} \|\mathbf{r}_x^{i+1} - \mathbf{r}_x^i\|^2 / \|\mathbf{r}_x^0\|^2$, in logarithmic scale, (g) Reconstructed support (h) Reconstructed intensity image.

training dataset we built are observed. They could be potentially removed by retraining the model with more heterogeneous data.

4.2 Real data

We then applied the proposed approach to high-density SMLM acquisitions using a publicly available dataset created for the 2013 SMLM challenge⁵, see Figure 4. Although in SMLM the molecules do not have a blinking behaviour, but rather an on-to-off transition, we can consider as blinking the temporal behaviour of one pixel in high-density videos due to the presence of many molecules per pixel. The dataset contains $T = 500$ images, the PSF of the microscope used to acquire these data has a FWHM of 351.8 nm and the pixel size is equal to 100 nm. The support $\hat{\Omega}$ computed by the model-based COLORME approach in [29,27] based on the use of a relaxation of the ℓ_0 pseudo-norm is compared to the one PnP-COLORME variant of Algorithm 1. Since no ground truth is available for these data, no quantitative assessment can be computed, however better continuation properties than COLORME [27] are observed.

⁵ <https://srm.epfl.ch/Challenge/Challenge2013>




Fig. 4: HD-SMLM data: (first row) The temporal mean and the first frame of the acquired temporal sequence (second row) Support (ℓ_0 -based [27] VS. PnP) and intensity reconstruction, $\sigma = 10/255$.

5 Conclusions

We presented a PnP model for support localisation for the deconvolution of imaging data in fluorescence microscopy. PnP approaches rely on the use of off-the-shelf denoisers to model implicit prior regularisation functionals. They can be effectively used to replace proximal steps in proximal gradient algorithms. Following [14], we choose a denoiser with a particular structure to benefit from convergence guarantees. Our results show that the geometry of specific structures (thin filaments) can be captured by suitable training. Future work should take into account the presence of a downsampling operator in the image formation model and a more accurate modelling making use of also cross terms in the covariance data.

6 Acknowledgements

VS and LBF are supported by the 3IA Côte d’Azur Investments (with reference number ANR-19-P3IA-0002). LC acknowledges the support received by the ANR project TASKABILE (ANR-22-CE48-0010) and the GdR ISIS project SPLIN. VS, LC and

LBF acknowledge the support received by the ANR project MICROBLIND (ANR-21-CE48-0008). All authors acknowledge the support received by the H2020 RISE projects NoMADS (GA. 777826).

References

1. Bolte, J., Sabach, S., Teboulle, M.: Proximal alternating linearized minimization for nonconvex and nonsmooth problems. *Mathematical Programming* **146**(1), 459–494 (2014)
2. Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J.: Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.* **3**(1), 1–122 (2011)
3. Cohen, R., Blau, Y., Freedman, D., Rivlin, E.: It has potential: Gradient-driven denoisers for convergent solutions to inverse problems. In: *Adv. Neural Inf. Process. Syst.* vol. 34, pp. 18152–18164. Curran Associates, Inc. (2021)
4. Dabov, K., Foi, A., Katkovnik, V., Egiazarian, K.: Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Trans. Image Process.* **16**(8), 2080–2095 (2007)
5. Dertinger, T., Colyer, R., Iyer, G., Weiss, S., Enderlein, J.: Fast, background-free, 3D super-resolution optical fluctuation imaging (SOFI). *PNAS* **106** (2009)
6. Douglas, J., Rachford, H.H.: On the numerical solution of heat conduction problems in two and three space variables. *Trans. Am. Math. Soc.* **82**(2), 421–439 (1956)
7. Duval, V., Peyre, G.: Exact support recovery for sparse spikes. *Found. Comput. Math.* **15**, 1315–1355 (2015)
8. Geman, D., Yang, C.: Nonlinear image recovery with half-quadratic regularization. *IEEE Trans. Image Process.* **4**(7), 932–946 (1995)
9. Girsault, A., Lukes, T., Sharipov, A., Geissbuehler, S., Leutenegger, M., Vandenberg, W., Dedecker, P., Hofkens, J., Lasser, T.: SOFI simulation tool: A software package for simulating and testing super-resolution optical fluctuation imaging. *PLOS ONE* **11**(9), 1–13 (2016)
10. Gribonval, R., Nikolova, M.: A characterization of proximity operators. *J. Math. Imaging Vis.* **62** (2020)
11. Gustafsson, N., Culley, S., Ashdown, G., Owen, D.M., Pereira, P.M., Henriques, R.: Fast live-cell conventional fluorophore nanoscopy with ImageJ through super-resolution radial fluctuations. *Nature comm.*, 7(1) pp. 12471–12471 (2016)
12. Hell, S.W., Wichmann, J.: Breaking the diffraction resolution limit by stimulated emission: stimulated-emission-depletion fluorescence microscopy. *Opt. Lett.* **19**, 780–782 (1994)
13. Hurault, S., Leclaire, A., Papadakis, N.: Gradient step denoiser for convergent plug-and-play. In: *International Conference on Learning Representations* (2022)
14. Hurault, S., Leclaire, A., Papadakis, N.: Proximal denoiser for convergent plug-and-play optimization with nonconvex regularization. *Proceedings of Machine Learning Research*, vol. 162, pp. 9483–9505. PMLR (2022)
15. Kamilov, U.S., Bouman, C.A., Buzzard, G.T., Wohlberg, B.: Plug-and-play methods for integrating physical and learned models in computational imaging: Theory, algorithms, and applications. *IEEE Signal Process. Mag.* **40**(1), 85–97 (2023)
16. Kamilov, U.S., Mansour, H., Wohlberg, B.: A plug-and-play priors approach for solving nonlinear imaging inverse problems. *IEEE Signal Process. Lett.* **24**(12), 1872–1876 (2017)

17. Laumont, R., Bortoli, V.D., Almansa, A., Delon, J., Durmus, A., Pereyra, M.: Bayesian imaging using plug & play priors: When Langevin meets tweedie. *SIAM J. Imaging Sci.* **15**(2), 701–737 (2022)
18. Laville, B., Blanc-Féraud, L., Aubert, G.: Off-the-grid curve reconstruction through divergence regularisation: an extreme point result (2022), HAL preprint
19. Parikh, N., Boyd, S.: Proximal algorithms. *Found. Trends Optim.* **1**(3) (2014)
20. Reehorst, E.T., Schniter, P.: Regularization by denoising: Clarifications and new interpretations. *IEEE Trans. Comput. Imaging* **5**(1), 52–67 (2019)
21. Romano, Y., Elad, M., Milanfar, P.: The little engine that could: Regularization by denoising (red). *SIAM J. Imaging Sci.* **10**(4), 1804–1844 (2017)
22. Ryu, E., Liu, J., Wang, S., Chen, X., Wang, Z., Yin, W.: Plug-and-play methods provably converge with properly trained denoisers. In: *Proceedings of the 36th International Conference on Machine Learning*. vol. 97, pp. 5546–5557. PMLR, Long Beach, California, USA (09–15 Jun 2019)
23. Sage, D., Kirshner, H., Pengo, T., Stuurman, N., Min, J., Manley, S., Unser, M.: Quantitative evaluation of software packages for single-molecule localization microscopy. *Nature methods*, 12 (06 2015). <https://doi.org/10.1038/nmeth.3442>
24. Sage, Daniel, et al.: Super-resolution fight club: Assessment of 2D & 3D single-molecule localization microscopy software. *Nat. Methods* **16** (05 2019)
25. Solomon, O., Eldar, Y.C., Mutzafi, M., Segev, M.: SPARCOM: Sparsity-based super-resolution correlation microscopy. *SIAM J. Imaging Sci.* **12**(1) (2019)
26. Soubies, E., Blanc-Féraud, L., Aubert, G.: A continuous exact ℓ^0 penalty (CEL0) for least squares regularized problem. *SIAM J. Imaging Sci.* **8**(3), 1607–1639 (2015)
27. Stergiopoulou, V., Calatroni, L., de Morais Goulart, H., Schaub, S., Blanc-Féraud, L.: COLORME: Super-resolution microscopy based on sparse blinking/fluctuating fluorophore localization and intensity estimation. *Biological Imaging* **2** (2022)
28. Stergiopoulou, V., Calatroni, L., Schaub, S., Blanc-Féraud, L.: 3D image super-resolution by fluorophore fluctuations and ma-tirf microscopy reconstruction (3D-COLORME). In: *2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI)*. pp. 1–4 (2022)
29. Stergiopoulou, V., de M. Goulart, J.H., Schaub, S., Calatroni, L., Blanc-Féraud, L.: COLORME: Covariance-based ℓ_0 super-resolution microscopy with intensity estimation. In: *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*. pp. 349–352 (2021)
30. Terris, M., Repetti, A., Pesquet, J.C., Wiaux, Y.: Building firmly nonexpansive convolutional neural networks. In: *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing*. pp. 8658–8662 (2020)
31. Venkatakrishnan, S.V., Bouman, C.A., Wohlberg, B.: Plug-and-play priors for model based reconstruction. In: *2013 IEEE Global Conference on Signal and Information Processing*. pp. 945–948 (2013)
32. Xu, X., Liu, J., Sun, Y., Wohlberg, B., Kamilov, U.S.: Boosting the performance of plug-and-play priors via denoiser scaling. In: *2020 54th Asilomar Conference on Signals, Systems, and Computers*. pp. 1305–1312. IEEE (2020)
33. Zhang, K., Li, Y., Zuo, W., Zhang, L., Van Gool, L., Timofte, R.: Plug-and-play image restoration with deep denoiser prior. *IEEE Trans. Pattern Anal. Mach. Intell.* (2021)
34. Zhang, K., Zuo, W., Chen, Y., Meng, D., Zhang, L.: Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising. *IEEE Trans. Image Process.* **26**(7), 3142–3155 (2017)