Joint self-supervised blind denoising and noise estimation

Jean Ollion[*], Charles Ollion[II], Élisabeth Gassiat[†], Luc Lehéricy[⊤], and Sylvain Le Corff[‡]

[*]SABILab, Die. `sabilab.fr`
[II]CMAP, École Polytechnique, Institut Polytechnique de Paris, Palaiseau.
[†]Université Paris-Saclay, CNRS, Laboratoire de mathématiques d'Orsay, 91405, Orsay.
[⊤]Laboratoire J. A. Dieudonné, Université Côte d'Azur, CNRS, 06108, Nice.
[‡]Samovar, Télécom SudParis, département CITI, TIPIC, Institut Polytechnique de Paris, Palaiseau.

Abstract

We propose a novel self-supervised image blind denoising approach in which two neural networks jointly predict the clean signal and infer the noise distribution. Assuming that the noisy observations are independent conditionally to the signal, the networks can be trained jointly without any clean data. This approach is particularly relevant for biomedical image denoising where the noise is difficult to model precisely and clean training data are usually unavailable. Our method significantly outperforms current state-of-the-art self-supervised blind denoising algorithms on six publicly available biomedical image datasets, as well as its supervised counterpart. We also show empirically that it is able to capture noise distributions efficiently, both on different synthetic noise models and real multimodal and skewed data. Finally, the described framework is simple, lightweight and computationally efficient, making it useful in practical cases.

1   Introduction

Image denoising is a well-known Computer Vision task designed to restore pictures taken in poor imaging conditions. In scientific imagery (microscopy, astronomy, etc.) for instance, the optical setting may produce very noisy images, which limits their interpretability or their automatic processing. Formally, image denoising is the process of recovering a clean signal X_s given an observation Y_s at location $s \in \mathbb{N}^2$ corrupted by an additive noise ε_s. Classical denoising approaches are model-driven in the sense that they rely on strong assumptions on the noise distribution or on the structure of the signal but are often limited by the relevance of these assumptions. Recently, efficient data-driven methods have emerged. Most of them assume that pairs made of noisy data Y_s associated with a clean signal X_s are available in a supervised learning framework, see for instance [Weigert et al., 2017]. In [Lehtinen et al., 2018], the authors have demonstrated that it is possible to train an efficient denoising method using only pairs of independent noisy measurements (Y_s^1, Y_s^2) of the same signal X_s. Such assumptions have also been used to solve deconvolution problems with repeated measurements as in [Delaigle et al., 2008]. However, obtaining independent observations of the same signal is often unrealistic in practice.

Recent self-supervised methods have overcome this limitation [Batson and Royer, 2019, Krull et al., 2018, Quan et al., 2020] for instance by training a neural network to predict the value of a (corrupted) pixel

Y_s only using the noisy observations of the surrounding pixels. In such frameworks, the trained network extracts some local structure in the signal and therefore can be used as a denoiser. Such approaches are referred to as *blind-denoising* as they only assume that the noises associated with different observations are independent and centered. This is well suited to typical microscopy settings, in which the clean image is unavailable and the noise process is complex and not known.

These methods rely on training a function f_θ depending on an unknown parameter θ, usually implemented as a convolutional neural network. Using only the noisy observations and a binary mask M, the objective is to minimize a self-supervision loss of the form $\theta \mapsto \sum_{i=1}^{N} \|f_\theta(Y^{masked})_{s_i} - Y_{s_i}\|_2^2$, where Y^{masked} is an image in which pixels have been masked using M. This masking step is crucial to foster learning of local structure in the signal to predict the masked values.

While these methods are appealing in practice and result in efficient denoising functions, they suffer from several drawbacks: 1) it is not well understood why they are so effective in practice, i.e. what type of noise they are able to remove, and how sensible they are to the masking scheme for instance; 2) they often suffer from high frequency denoising artifacts known as *checkerboard pattern*.

To address these issues, we introduce a novel self-supervised method, based on the joint training of two neural networks referred to as D-net (denoiser net) and N-net (noise net). Similar to previous works, the denoiser is a convolutional neural network and receives a masked input during training. The flexible N-net recovers precisely the noise distribution during training, even for complex asymmetric noises. We derive this method from a different mathematical modeling of the denoising problem, which may open avenues for better understanding of why self-supervised methods achieve remarkable results.

The contribution of this work can be summarized as follows:

- we introduce a novel self-supervised blind-denoising method, modeling both the signal and the noise distributions;

- we show that the N-net recovers the noise distribution efficiently in varying experiments with synthetic and real noises;

- the proposed architectures outperform state-of-the-art algorithms over 6 standard microscopy datasets, without introducing denoising artifacts.

2   Related work

Masking and J-invariance.   The most typical class of denoising functions is chosen to be comprised of convolutional neural networks (CNNs), which are heavily parameterized functions and are not restricted to solve denoising problems. As an important consequence, a naive self-supervised loss without any masking would result in learning the identity function (i.e. the function outputing the noisy observation Y_s if it is not masked in the input data), as the considered CNNs can typically implement it. Starting from this intuition, many of the related works can be viewed as different masking schemes. This has been described in the *J-invariant* framework introduced by [Batson and Royer, 2019]: a *J-invariant* function does not depend on a few selected dimensions J of its input variables; typically this translates into a convolutional function which does not depend on the central pixel of the convolutional receptive field, but rather on the observations of neighboring pixels[1].

The first masked self-supervised denoising methods were introduced by Noise2Void (N2V) [Krull et al., 2018], in which $\{(Y^{masked}_{s_i}, Y_{s_i})\}_{1 \leqslant i \leqslant N}$ are sampled randomly in the picture, and masking consists in

---

[1]Those functions excluding the central pixel are sometimes also called *blind spot*, not to be confused with *blind denoising* in which the noise process is not known.

replacing Y_{s_i} by a random observed value in its neighborhood, with a positive probability of replacing Y_{s_i} by itself meaning that leaks in masking are introduced.

Noise2Self (N2S) [Batson and Royer, 2019] masking procedure differs from N2V in the sense that $\{(Y_{s_i}^{masked}, Y_{s_i})\}_{1 \leqslant i \leqslant N}$ are obtained with a fixed grid, and Y_i is replaced by the average of the 4 direct neighboring observations. In practice, the masking procedure has a strong impact on training: (i) improving masking schemes can improve denoising performance and (ii) as only masked pixels are used for training, typically representing a few percent of the image, this affects greatly the training efficiency.

The underlying CNN architecture implemented by these works is the U-net [Ronneberger et al., 2015], a typical convolutional autoencoder architecture, involving skip connections, which can reproduce fine grained details, while making use of higher-level spatially coarse information. While showing strong denoising performance in N2S and N2V, they however can produce *checkerboard patterns*, which are high frequency artifacts that arise in the denoised results. These works have been extended in DecoNoising [Goncharova et al., 2020] in which a Gaussian convolution is added after the neural network output to simulate microscope Point Spread Function. This technique improves performances, however the deconvolved image (predicted image before the Gaussian convolution) displays even stronger checkerboard pattern.

It is noteworthy that [Broaddus et al., 2020] showed that when the noise has local correlations (for instance a directional noise), masking can be adapted to remove them - by masking adjacent pixels in the same direction as the noise spatial correlation for instance.

Finally, a few methods with different have been developed different training schemes, such as Neighbor2Neighbor [Huang et al., 2021], in which a pair of sub-sampled images is generated from the noisy image. The author then use this pair in a way similar to [Lehtinen et al., 2018], in order to train their denoising model with one being the input and the other the output. One of the main benefits is that the whole downsampled image is used for training. On the other hand, it relies heavily on the downsampling procedure and does not enable to learn information about the noise distribution.

J-invariance without masking.   Instead of masking specific pixels, it is possible to design specific convolutional operators to limit the receptive field, ensuring that the resulting function is *J-invariant* by design. This was achieved in [Laine et al., 2019] by introducing directional convolution kernels, each kernel has its receptive field restricted to a half-plane that does not contain the central pixel. The associated function then takes values which only depend on pixels in specific directions, ensuring that it does not depend on pixels in the opposite direction. One drawback is that the inference has to be performed four times, one in each direction.

More recently, [Lee and Jeong, 2020] introduced a combination of specific convolution operators[2] with dilation and strides, guaranteeing that the function is independent of the central pixel by design, therefore *J-invariant*.

The benefit of these architectures compared to the masking-based training is that all output pixels can contribute to the loss function as in conventional training, rather than just the masked pixels; and they do not require a carefully tuned masking procedure. However, they strongly constrain the network architecture, which can hinder the denoising performance or result in more expensive inference schemes. Moreover, recent work show that fully *J-invariant* models lead to suboptimal performance [Xie et al., 2020].

Contribution of a noise model.   The denoising literature includes few works which explicitly model the noise distribution, either by choosing *a priori* a family of distributions (e.g. a Gaussian noise), or by selecting a more flexible class of distributions.

---

[2]It is interesting to note that with standard convolutions, a two layered network already cannot be made independent of the central pixel, which is why the authors had to rely on very specific convolutions.

The former is illustrated in [Laine et al., 2019], in which three types of corrupting noise are considered: Gaussian noise independent of the signal; Poisson-Gaussian noise, i.e. a Gaussian noise with variance scaling linearly with the signal value; finally impulse noise, i.e. a uniform noise. In each case, the noise parameters are either known or estimated with an auxiliary neural network. As the signal distribution and the noise distribution belong to a known parametric family, the noisy central pixel can be included at test time to improve performances. However, as the noise type has to be chosen *a priori*, the method is restricted to known and synthetic noise types and therefore is a *non-blind denoising* method.

In [Krull et al., 2019, Prakash et al., 2020b], the authors make use of a more flexible noise model, which is a generic modelisation of the conditional distribution of the noise given the signal intensity and thus can better model real noises. In these works, noise models are approximated using 2D histograms of denoised and noisy observed values, either using additional calibration data (in that case the method is not fully self-supervised) or using a previously trained denoising function [Prakash et al., 2020b]. This increases the complexity of the method, as it requires several training procedures and calibration. In the latter variant, the noise distribution is parametrized with a centered Gaussian mixture model with empirically designed constraints.

In [Prakash et al., 2020a], the authors use the Variational AutoEncoder formalism, adding a noise model in their architecture. This provides new interesting possibilities, as it can generate a diversity of denoised results, but sometimes resulting in visual artifacts or blurry results. When co-learning the noise model, the noise variance is an affine function of the signal value, which corresponds to assuming a Poisson-Gaussian noise distribution.

Perhaps closest to our method, in FBI [Byun et al., 2021], the authors focus on Poisson-Gaussian noise, enabling to both denoise efficiently and recover the Poisson-Gaussian parameter. However, they rely heavily on their Poisson-Gaussian assumption to derive their image transformation process and loss function, and therefore unadapted to other types of noise. While this assumption is reasonable for some types of real-world noises, it falls short on other cases, typically with Speckle noise, or at extreme values of signal where the noise distribution can be ill-defined.

It is worth noting that supervised *blind denoising* methods have used parametrized noise models, such as [Zhang et al., 2017, Yue et al., 2019], which explicitly used large neural networks to model a complex noise, with even less assumptions (it can be slightly structured). Even though the algorithm proposed in [Yue et al., 2019] is able to train jointly a noise network and a denoiser, their modeling only works in a supervised setting.

Chosen approach.   Given a short overview of the related denoising, the following summarizes the key differences with our approach:

- *J-invariance*: The proposed approach is not strictly *J-invariant*, which enables for more flexibility in the network architectures, as well as the use of the central pixel at test time. This also alleviates the need to have several inferences steps such as in [Laine et al., 2019].

- *Modeling the noise*: Our method enables to model and recover the noise distribution. While some other methods also enable this, they rely on a type of noise assumption (typically: Gaussian of Poisson-Gaussian), while our parametrised noise distributions are much more flexible. This is critical in real noise settings where the noise is not standard nor known.

- *Joint training*: We propose a single model in which all the parameters of our models are trained joinlty (both the denoising network and the noise model) by optimizing a single criterion. This results in a very simple, practical and stable procedure, which does not require preprocessing or calibration steps.

3   Model

Estimating a signal corrupted by additive noise is a challenging statistical problem. In such frameworks, the received observation at location $s \in \mathbb{N}^2$, Y_s, is given by $Y_s = X_s + \varepsilon_s$, where X_s is the signal and ε_s is the noise. A lot of works have been devoted to deconvolution where the aim is to recover the distribution of the signal based on the observations. It has been for instance applied in a large variety of disciplines and has stimulated a great research interest in signal processing [Moulines et al., 1997, Attias and Schreiner, 1998] and in image reconstruction [Kundur and Hatzinakos, 1996, Campisi and Egiazarian, 2017], see also [Meister, 2009]. Recently, [Gassiat et al., 2021] proved that it is possible to recover the signal distribution when X has at least two dimensions and may be decomposed into two subsets of random variables which satisfy some weak dependency assumption. This identifiability result does not require any assumption on the noise distribution but illustrates that the components of the signal must be dependent to allow its identification.

Let $\mathsf{S} \subset \mathbb{N}^2$ be a finite set corresponding to the pixel location of an image. In this work, for all $s \in \mathsf{S}$, X_s denotes the signal value at location s and Y_s its noisy observation. For all $s \in \mathsf{S}$, we write

$$Y_s = X_s + \varepsilon_s \,, \tag{1}$$

where $(\varepsilon_s)_{s \in \mathsf{S}}$ are signal-dependent random variables. In this paper, we consider the following assumptions.

- The random variables $(\varepsilon_s)_{s \in \mathsf{S}}$ are independent conditionally to $(X_s)_{s \in \mathsf{S}}$ and the conditional distribution of ε_s depends on X_s only.

- For all $s \in \mathsf{S}$, ε_s is centered.

- For all $s \in \mathsf{S}$, conditionally to X_s, ε_s follows a Gaussian mixture model (GMM) with signal-dependent weights $(\alpha_{k,\theta_n}(X_s))_{1 \leqslant k \leqslant N}$, means and variances $\{(\mu_{k,\theta_n}(X_s), \sigma^2_{k,\theta_n}(X_s))\}_{1 \leqslant k \leqslant N}$.

The nonnegative mixture weights $(\alpha_{k,\theta_n}(X_s))_{1 \leqslant k \leqslant N}$ are such that $\sum_{k=1}^{N} \alpha_{k,\theta_n}(X_s) = 1$, and the means and variances $\{(\mu_{k,\theta_n}(X_s), \sigma^2_{k,\theta_n}(X_s))\}_{1 \leqslant k \leqslant N}$ of the mixture model are parameterized by a convolutional neural network, called N-net and with unknown weights θ_n, see Section 4.1.

The aim of this paper is to propose an estimator of the signal value without introducing a prior model on $(X_s)_{s \in \mathsf{S}}$. Our approach is motivated by recent data-driven solutions which allow to avoid using too restrictive prior models with poor predictive performance. For all $s \in \mathsf{S}$, let Ω_s be a neighborhood of s, i.e. a set of pixel locations around s. In the following, for any finite set $A \subset \mathsf{S}$, $X_A = \{X_s \,;\, s \in A\}$ and $Y_A = \{Y_s \,;\, s \in A\}$. We propose to estimate the conditional mean of X_s given (Y_s, Y_{Ω_s}) by a parametric function denoted by μ_{θ_d} so that $\mathbb{E}[X_s | Y_s, Y_{\Omega_s}]$ is estimated by $\mu_{\theta_d}(Y_s, Y_{\Omega_s})$. The function μ_{θ_d} is parameterized by a convolutional neural network, called D-net and with unknown weights θ_d, see Section 4.1.

For all $s \in \mathsf{S}$, the conditional likelihood of a the central pixel given its neighborhood is given by

$$p_\theta(Y_s | Y_{\Omega_s}) = \int p_\theta(x_s, Y_s | Y_{\Omega_s}) \mathrm{d}x_s = \int p_\theta(x_s | Y_{\Omega_s}) p_\theta(Y_s | x_s) \mathrm{d}x_s \,,$$

where the last equality comes from the fact that Y_s is independent of Y_{Ω_s} given X_s. Assuming that $p_\theta(\mathrm{d}x_s | Y_{\Omega_s})$ is a peaked distribution, we propose to approximate it by the Dirac mass at $\mu_{\theta_d}(g(Y_{\Omega_s}), Y_{\Omega_s})$, i.e.

$$p_\theta(\mathrm{d}x_s | Y_{\Omega_s}) \simeq \delta_{\mu_{\theta_d}(g(Y_{\Omega_s}), Y_{\Omega_s})}(\mathrm{d}x_s) \,,$$

where g is a known function. In the experiments below, we chose to set $g(Y_{\Omega_s})$ as the empirical mean of the noisy pixels in Y_{Ω_s}.

Considering now a set of pixel locations $(s_i)_{1 \leqslant i \leqslant p}$ in S with non-overlapping set of neighborhoods $(\Omega_{s_i})_{1 \leqslant i \leqslant p}$, we therefore propose the following pseudo-loglikelihood function to be maximized during training to estimate θ_n and θ_d:

$$\ell_p(\theta_n, \theta_d) = \frac{1}{p} \sum_{i=1}^{p} \ell_\theta(Y_{s_i}|Y_{\Omega_{s_i}})\,,$$

with

$$\ell_\theta(Y_s|Y_{\Omega_s}) = \log \left(\sum_{k=1}^{N} \alpha_{k,\theta_n}(\tilde{\mu}_{\theta_d}^s) \varphi_{\check{\mu}_{\theta_d}^s, \sigma_{k,\theta_n}(\tilde{\mu}_{\theta_d}^s)}(Y_s) \right)\,,$$

where $\tilde{\mu}_{\theta_d}^s = \mu_{\theta_d}(g(Y_{\Omega_s}), Y_{\Omega_s})$, $\check{\mu}_{\theta_d}^s = \tilde{\mu}_{\theta_d}^s + \mu_{k,\theta_n}(\tilde{\mu}_{\theta_d}^s)$ and $\varphi_{\mu,\sigma}$ is the Gaussian probability density function with mean μ and standard deviation σ.

Note that our model encompasses the case of additive Gaussian noise when $N = 1$ and in this case, we denote by $\sigma_{\theta_n} : \mathbb{R} \to \mathbb{R}$ the function describing the local standard deviation of the noise distribution. This model contrasts with most common denoising algorithms where ε is assumed to be a centered Gaussian random variable and with a variance which is either known and constant or has a Poisson-Gaussian shape i.e., scales with $\alpha x + \eta^2$. As illustrated in Section 5, these assumptions do not usually hold, in particular when considering biomedical images, and they may have a severe impact on denoising performances. We display in Section 5 the benefit of our mixture model to account for positive skewness which cannot be modeled with a single Gaussian distribution. The results provided in Section 5 illustrate how such models improve denoising performance for asymmetrical noise distributions.

In [Gassiat et al., 2021], the variance of the noise is assumed to be constant and the target signal is assumed to be weakly dependent to obtain identifiability of the noise and the signal distributions. In (1), we extend the model proposed by [Gassiat et al., 2021] by considering mixture models with state-dependent standard deviations and identifiability remains an open problem. However, we assume in this work that X_s is dependent with the signal in the neighbooring pixels X_{Ω_s} so that heteroscedasticity is the main challenge to obtain identifiability of (1).

4 Experiments

4.1 Model Architecture

D-net The function μ_{θ_d} is parametrized by a U-net, with slight architecture modifications. Architecture and training information can be found in AppendixA. The receptive field of this network is 35x35 pixels, which means that the network may use pixels from the neighborhood that are masked. At test-time, we averaged the prediction of the image with the predictions of its transposed and flipped versions on each axis, which improves performances.

N-net In the Gaussian Mixture Model (GMM) case, the network has several outputs: for a mixture of N Gaussian distributions, there are N variances, $N - 1$ means (the last mean is computed to ensure that the resulting distribution is centered) and N mixture weights parametrized by the N-net. In the case where $N = 1$, the function $\sigma_{\theta_n} : \mathbb{R} \to \mathbb{R}$ describing the local variance of the noise distribution is a fully-connected deep neural network with several hidden layers. This choice is motivated by the large expressivity of such a network, necessary to approximate complex noise distributions. In practice, it is applied to each pixel, so it is implemented efficiently as a fully convolutional network using only 1x1 convolutional layers. The full architecture details for both models are available in AppendixA.
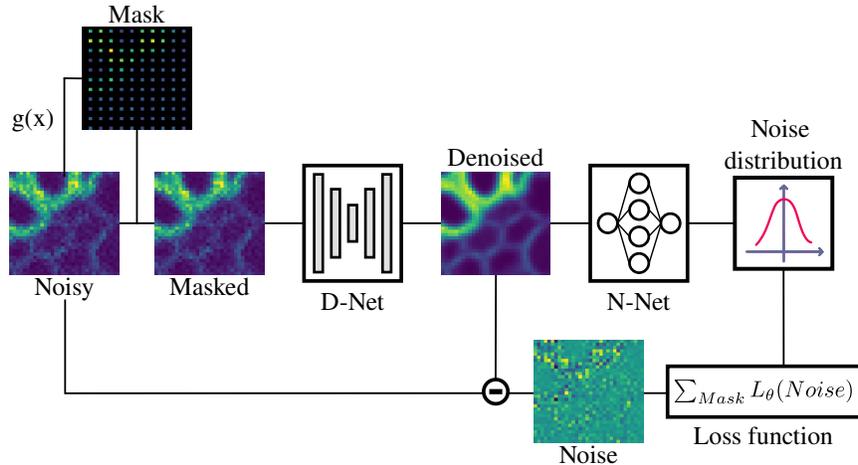
Figure 1: Training setup. For each mini-batch, a random grid is drawn. The masking function $x \mapsto g(x)$ is applied on each element of the grid, replacing the original pixels in the masked image. The denoised image predicted by the D-net is fed to the N-net that predicts a noise distribution for each pixel. The loss function is then computed on each element of the grid.

4.2   Datasets

We train and evaluate our method on 6 publicly available datasets of microscopy images. In those datasets, ground truth (X) is estimated by averaging several observations (Y) of the same field-of-view (FOV). This allows to have access to an estimation of the noise $Y - X$, which we refer to as *real noise* in this article.

The 3 first datasets (*PN2V-C*, *PN2V-MN*, *PN2V-MA*) have been published along with the PN2V method [Krull et al., 2019], each is composed of several observations of one single FOV. For a fair comparison, we use the same training and evaluation sets as the authors: for each sample type the whole dataset is used for training, and only a subset of the FOV is used for evaluation (see AppendixC.1 for details). The 3 last datasets are the 3 channels of the W2S dataset [Zhou et al., 2020] referred to as *W2S-1*, *W2S-2* and *W2S-3*. The dataset is composed of 120 FOV, the first 80 are used for training and the last 40 for evaluation (see AppendixC.1 for more details). Following the authors, for each FOV, only one observation is used for training and for evaluation, which better corresponds to a real setting where only one observation per FOV is available.

4.3   Masking procedure

Following [Batson and Royer, 2019], we mask pixels along a grid and compute the loss only on masked pixels. We obtained the best results by replacing the central value by the weighted average of the 8 direct neighbors with Gaussian weights ($\sigma = 1$). The drawback of masking along a grid is that pixels are masked at fixed relative positions with regards to the central pixel. If grid spacing is too small, then too many masked pixels are present in the receptive field and perturb the performances, because the available information is reduced. On the other hand, the larger the spacing, the less pixels are used for training, which reduces dramatically training efficiency. In order to push the limits of this trade-off, we use a random dynamic spacing

between 3 and 5 pixels, which allows to have relative positions of masked pixels that change randomly. On average, 6.8% of the image is masked.

Furthermore, we observed that datasets *PN2V-C* and *PN2V-MA* display axial correlation in the noise, for those datasets we adapted the masking procedure introduced in [Broaddus et al., 2020]: the replacement value was computed on a neighborhood excluding the neighbors along the correlation axis, and neighbors were masked along this axis, within an extent of 3 pixels. This can be determined easily in a self-supervised setup because the neural network tends to amplify the noise correlation.

4.4 Training

Networks are trained using Adam optimizer with a learning rate of $4 \cdot 10^{-4}$, decreased by $1/2$ on plateau of 30 epochs until 10^{-6}. We train networks for 400 epochs of 200 steps. Training time is about 2 min per epoch on a NVIDIA Tesla P4. We obtain better and more reproducible results using the weights of the trained model at the last epoch instead of the weights of the model with the best validation loss, possibly because the loss is a bad proxy for the denoising performances. For that reason, we do not use a validation step. Batch size is set to 1, and each batch is split into 100 (overlapping) tiles of 96x96 pixels. Tiles are augmented with random horizontal and/or vertical flip and/or a random rotation with an angle chosen within $(90°, 180°, 270°)$, except for datasets with axial noise correlation where axes transpositions are avoided.

4.5 Evaluation

We compared denoised image to ground truth with the classical Peak Signal-to-Noise Ratio (PSNR) metric. However, PSNR is not highly indicative of perceived similarity, in particular it does not reflect similarity of high frequency information such as textures and local contrasts [Wang et al., 2004], that denoising methods tend to reduce. It is thus essential to have other metrics that take them into account. To address this short-coming, we used Structural Similarity (SSIM) that take textures and edges into account [Wang et al., 2004], computed as in the original work.

5 Results

5.1 Noise estimation

Estimation on synthetic noise   To evaluate the capacity of the N-net to capture blindly different noise distributions, we generated 3 datasets by adding synthetic noise to the ground truth of dataset *W2S-1*, and we chose the parameters of the noise models so that PSNR of noisy images match the one of the original dataset (in a range of ±0.1dB). We used 3 classical noise models: additive Gaussian, Poisson-Gaussian (which is a good model for shot noise) and speckle (see AppendixC.2 for details). Empirical and predicted distributions of the noise standard deviation are illustrated in Fig. 2. One of the most striking result of this experiment is that for the 3 cases of synthetic noise, the predicted standard deviation provided by the N-net is a very sharp approximation of the known theoretical standard deviation. It shows in particular that our method is able to capture the different noise distributions even in areas where signal is rare.

Improving estimation on real noise   We observed that contrary to the classical noise models considered in the denoising literature, real noise often displays a certain amount of skewness, as illustrated in Fig. 3. In order to be able to capture this aspect, we predict a Gaussian mixture model (GMM) instead of a simple Gaussian model as described in Section 3. Fig. 3 shows that noise skewness is well described by the predicted
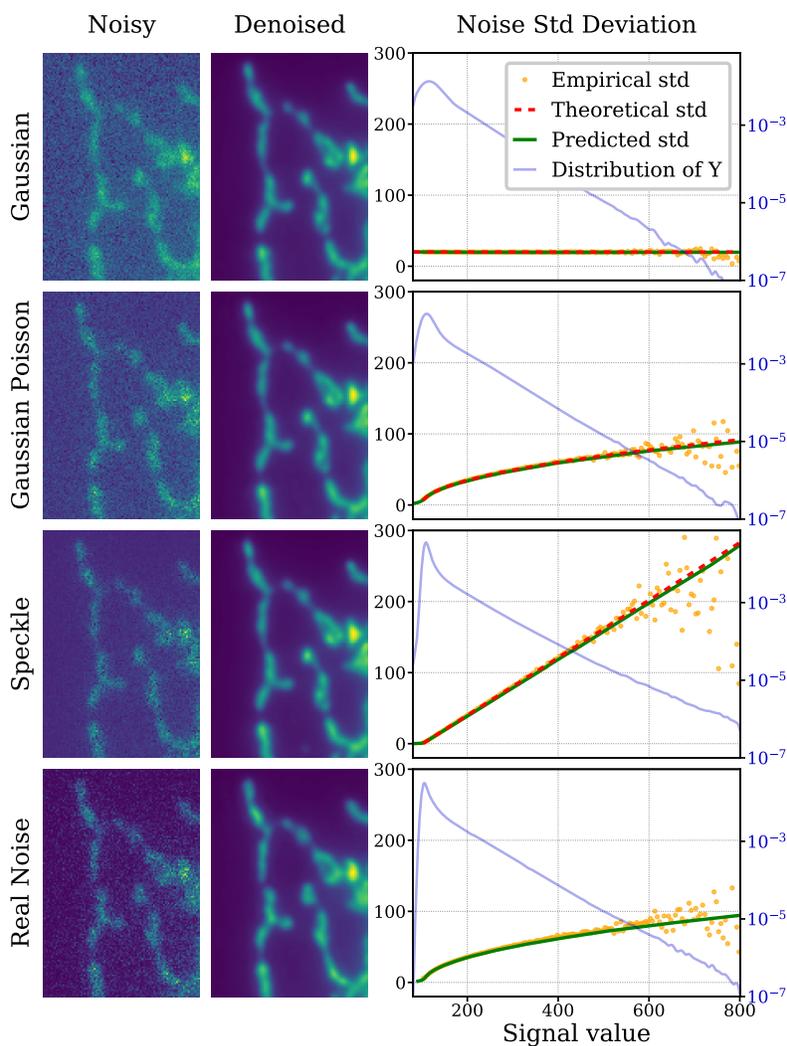
Figure 2: Noise estimation. For 3 models of synthetic noise as well as the real noise, the plots display the empirical standard deviation of the noise $Y - X$, as well as the predicted standard deviation of the noise by the N-net as a function of X (note that display range was shrunk in Y-axis for visualization purposes.). Theoretical standard deviation of the noise is displayed for the 3 models of synthetic noise. The empirical distribution of Y is displayed in blue, in logarithmic scale. Examples of noisy images corrupted with the corresponding noise model and the predicted denoised images are displayed above each graph.

model, and the noise distribution is better described by a GMM than by a single Gaussian. This applies for all datasets and the equivalent figures can be found in AppendixE. In this example, it is interesting to note that the Kullback–Leibler divergence between the empirical noise distribution and the predicted distribution (as a function of the signal value) is improved by considering a GMM instead of a unimodal distribution. This supports the use of our flexible N-net to capture a large variety of noise distributions (with multimodality and/or skewness) which can be observed in experimental datasets. This comment paves the way to several perspectives for our work such as the design of statistically consistent model selection procedures to choose automatically the number of mixing components. Such approaches have been proposed in more simple cases using for instance penalized maximum likelihood based algorithms. This remains an open problem in our framework and we leave this topic for future research.
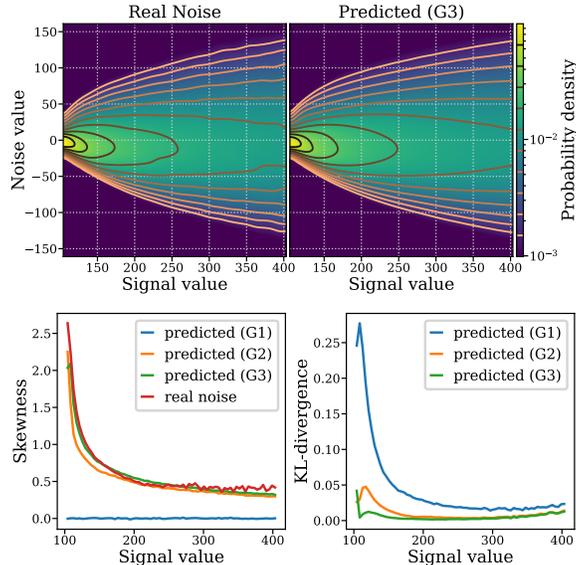

Figure 3: Real noise estimation for dataset *W2S-1*. The two left graphs represent the empirical distribution of the noise $Y - X$ as a function of X and the corresponding predicted noise distribution for a 3-component-GMM. The probability density is normalized for each signal value bin. Skewness of real and predicted noise distribution as a function of X, estimated with Pearson's moment coefficient of skewness. Kullback–Leibler divergence between real noise distribution and predicted distribution generated by each model, as a function of X. *G1* stands for Gaussian model, *G2* for a 2-component-GMM and *G3* a 3-component-GMM.

5.2   Denoising performances

We compared our method to 6 baselines: D-net trained in a supervised setup with L2 objective (with all other training and prediction hyperparameters unchanged), N2V, DecoNoising, which is the self-supervised blind denoising method that has shown best results on the datasets we considered, Neighbor2Neighbor and FBI-Denoiser, as well as one of the most simple denoising method: convolution by a Gaussian, whose standard deviation is chosen to maximize the PSNR on the evaluation dataset. We believe the latter makes a good

Table 1: Evaluation of our method on 6 datasets with PSNR/SSIM metrics. SSIM estimates structural similarity (sharpness). Metrics computed on noisy images are displayed in the *Noisy* column. The supervised version of our method was not trained for PN2V datasets as they cannot be split into independent train/evaluation sets. For DecoNoising and N2V, PSNR are taken from [Goncharova et al., 2020] and SSIM are computed on prediction made by networks we trained (see Section B). *Gaussian* corresponds to the optimal Gaussian baseline defined in Section 5.2. *Neigh2Neigh* stands for Neighbor2Neighbor. Results for our method predicting a Gaussian model are shown in column *Ours (G1)* (see Appendix Table 3 for 2/3-component GMM). Best PSNR $\pm 0.1 dB$ / SSIM $\pm 1\%_{oo}$ scores are underlined.

METHOD/DATASET	PN2V-C	PN2V-MN	PN2V-MA	W2S-1	W2S-2	W2S-3
NOISY	28.98 / 0.7713	28.10 / 0.6836	23.71 / 0.3731	21.85 / 0.3490	19.33 / 0.2256	20.39 / 0.2232
GAUSSIAN	34.92 / 0.9409	35.53 / 0.9392	34.07 / 0.8739	33.87 / 0.9326	32.27 / 0.8531	34.66 / 0.9013
SUPERVISED	NA / NA	NA / NA	NA / NA	35.22 / 0.9608	33.24 / 0.8828	36.31 / 0.9252
N2V	35.85 / 0.9404	35.86 / 0.9419	33.35 / 0.8384	34.30 / 0.9026	31.80 / 0.8311	34.65 / 0.8637
DECONOISING	36.39 / 0.9483	36.34 / 0.9489	34.04 / 0.8633	34.90 / 0.9169	32.31 / 0.8524	35.09 / 0.9051
NEIGH2NEIGH	34.42 / 0.9519	34.01 / 0.9070	32.70 / 0.8521	34.74 / 0.9552	32.96 / 0.8739	36.14 / 0.9216
FBI-DENOISER	36.23 / 0.9556	36.85 / 0.9609	32.89 / 0.8344	34.40 / 0.9500	32.26 / 0.8452	34.94 / 0.8910
OURS (G1)	<u>38.33</u> / <u>0.9754</u>	<u>39.08</u> / <u>0.9776</u>	<u>34.79</u> / <u>0.8905</u>	<u>35.33</u> / <u>0.9619</u>	<u>33.46</u> / <u>0.8867</u>	<u>36.57</u> / <u>0.9263</u>

reference, as it is one of the simplest denoising methods, and it removes noise efficiently but also other high-frequency information such as local contrasts. Training procedure is detailed in AppendixB.

The considered metrics are summarized in Table 1. Our method significantly outperforms the baselines both in terms of PSNR and SSIM on all datasets. Note that datasets PN2V-C and PN2V-MA have horizontally-correlated noise (that we take into account: see Section 4.2), which significantly lowers the performances of methods that are sensitive to it, such as FBI-Denoiser or N2V. For the version predicting a simple Gaussian distribution, the average PSNR gain over DecoNoisng is +1.42dB. This is also confirmed by the visual aspect, displayed in Fig. 4: our method produces images closer to the ground truth, smoother, sharper, more detailed and without visual artifacts. Remarkably, our method performs significantly better than the supervised method CARE [Weigert et al., 2017], with an average PSNR gain of +1.49dB (compared to PSNR values reported in [Goncharova et al., 2020]). Moreover, our method also performs better than its supervised counterpart with an average gain of +0.21dB (see Table 1). This could be explained by the fact that training with masking induces a lower dependency to central pixel compared to supervised training, and thus pushes the network to make better use of the neighborhood.

5.3   Ablation experiments

To better understand the contribution of the N-net and other hyperparameters, we ran ablation experiments in Table 2. The N-net has a very significant impact on performances, in particular when training is short (e.g. as in N2V, N2S [Krull et al., 2018, Goncharova et al., 2020]). We observe that the N-net greatly stabilizes the training[3] and improves convergence speed (see Fig. 5). We also observe that in all our hyperparameter settings that adding the N-net always improves performances over L2 objective, although not always significantly.

From the perspective of our framework, the L2 objective (as used in N2V, N2S) can be understood as a particular case of the N-net predicting a constant unit Gaussian noise. When the actual noise is significantly

---

[3]As an illustration of the excellent training stability, we measured a SEM of $6 \cdot 10^{-3}$ for the PSNR on dataset W2S-1 over 5 runs.
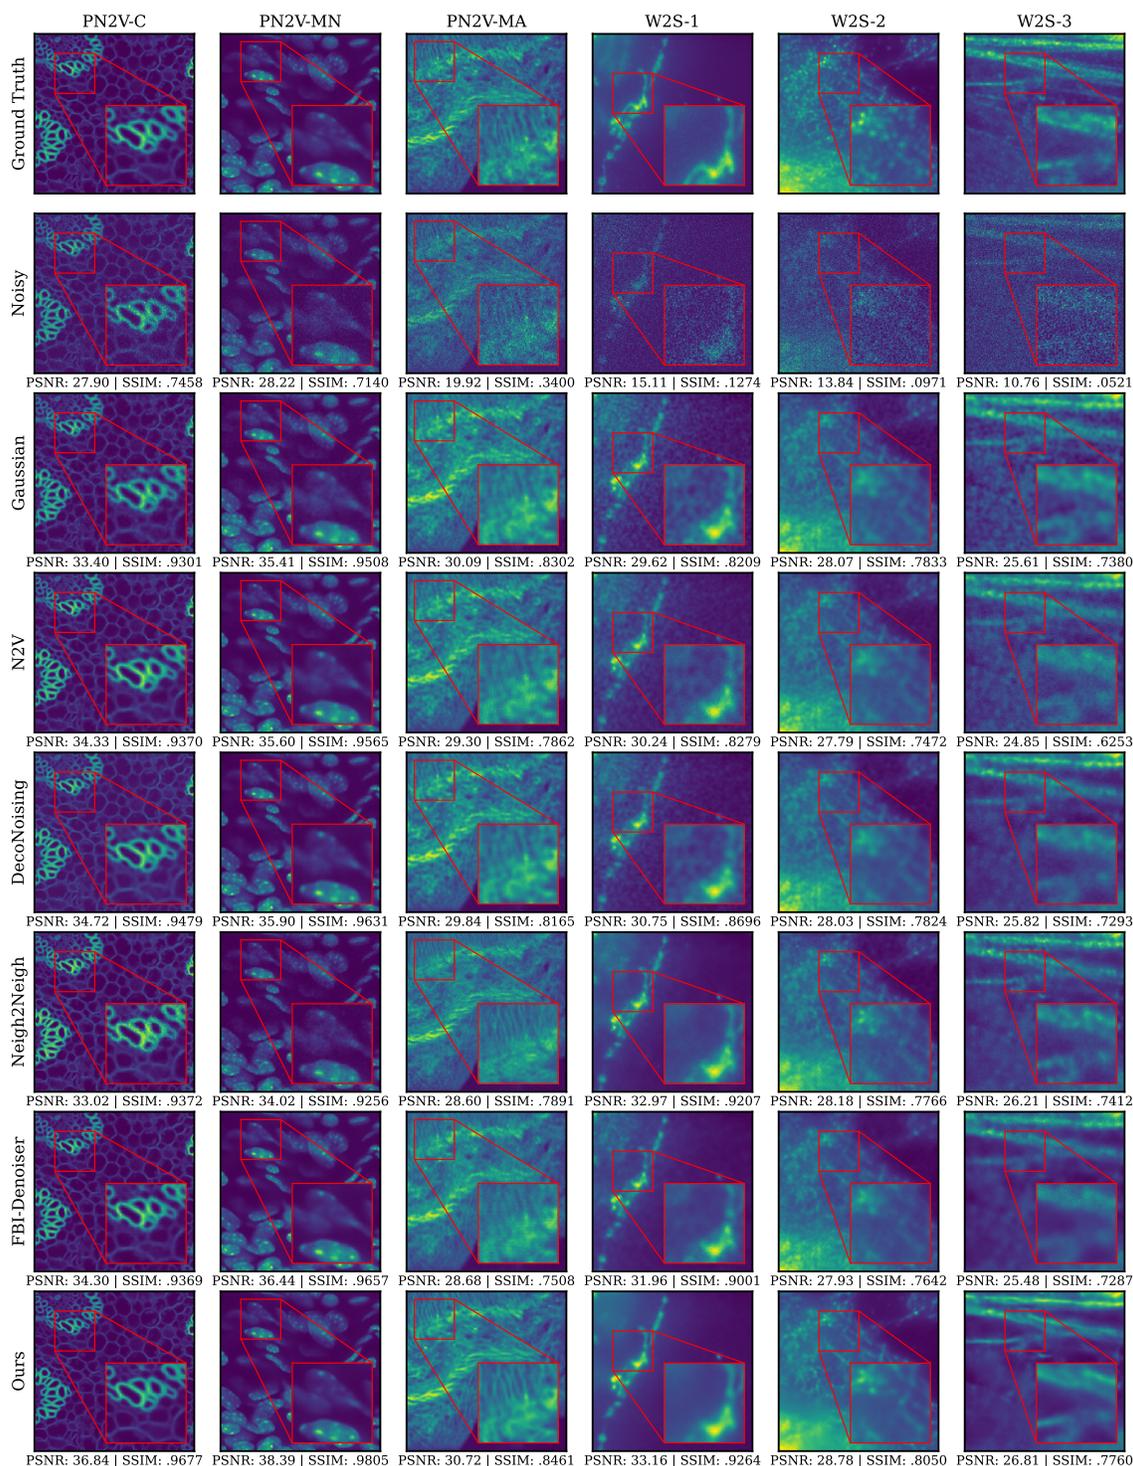
Figure 4: Visual comparison of denoising on the considered datasets. For each dataset a 256x256 portion of an evaluation image is displayed, on which metrics are computed and displayed below. Training of baselines is described in section B. *Gaussian* corresponds to the optimal Gaussian baseline defined in section 5.2.

different from it, the N-net has a stronger impact (e.g. *Speckle* in Fig. 5B). From a statistical perspective, it is known that considering a sufficiently rich noise model is crucial in deconvolution problem for signal estimation (misspecifying the noise density can lead to very poor deconvolution estimators for instance). This was a motivation to introduce the N-net. It is worth noting that considering mixture models improves the PSNR in two datasets out of six (Table 3). As mentioned in Section 5.1, an optimal and data-driven choice of the number of components remains an open (and challenging) statistical problem but we believe that such experiments support future research in this direction.

Table 2: Ablation experiment on dataset W2S-1. *N2V*: Noise2Void (L2). + *N-net (G1)*: N-net with a single Gaussian component. + *LT*: Longer training, see 4.4. *Ours* as described in 4.

CONDITION	PSNR	SSIM
N2V	34.30	0.9026
+ N-NET (G1)	34.99	0.9572
+ N-NET (G1) + LT	35.22	0.9608
OURS WITH N-NET (G1)	<u>35.33</u>	<u>0.9619</u>

6   Discussion and Limitations

We introduced a novel self-supervised blind-denoising method modeling both the signal and the noise distributions. We believe its simplicity, performances and the interpretability of the noise distribution will be useful both in practical applications, and as a basis for future research.

First, future works could consider more complex families of noise distributions such as structured or non-centered noises, that can also arise in real-life setups. In particular, [Lehtinen et al., 2018] managed to remove very structured non-centered noises such as overlaid text. With stronger assumptions and architecture changes, it might be possible to capture such noises.

Second, more theoretical works could explore the model proposed in this work (i) to obtain identifiability of model (1) and extend [Gassiat et al., 2021] to state-dependent standard deviations and (ii) to establish rates of convergence for the proposed estimators.

Finally, it would also be interesting to understand the role of the central pixel at test time, as it has a significative impact on performance: it depends on the masking procedure and the convolutional architecture, but the network is not trained explicitly to use it. Our mathematical modeling could be a good basis to study this specific dependency on the central pixel.

References

[Attias and Schreiner, 1998] Attias, H. and Schreiner, C. E. (1998). Blind source separation and deconvolution: the dynamic component analysis algorithm. *Neural computation*, 10(6):1373–1424.

[Batson and Royer, 2019] Batson, J. and Royer, L. (2019). Noise2self: Blind denoising by self-supervision. *arXiv preprint arXiv:1901.11365*.

[Broaddus et al., 2020] Broaddus, C., Krull, A., Weigert, M., Schmidt, U., and Myers, G. (2020). Removing structured noise with self-supervised blind-spot networks. In *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, pages 159–163. IEEE.

[Byun et al., 2021] Byun, J., Cha, S., and Moon, T. (2021). Fbi-denoiser: Fast blind image denoiser for poisson-gaussian noise. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5768–5777.

[Campisi and Egiazarian, 2017] Campisi, P. and Egiazarian, K. (2017). *Blind image deconvolution: theory and applications*. CRC press.

[Delaigle et al., 2008] Delaigle, A., Hall, P., Meister, A., et al. (2008). On deconvolution with repeated measurements. *The Annals of Statistics*, 36(2):665–685.

[Gassiat et al., 2021] Gassiat, E., Le Corff, S., and Lehéricy, L. (2021). Deconvolution with unknown noise distribution is possible for multivariate signals. *To appear in The Annals of Statistics, ArXiv:2006.14226*.

[Goncharova et al., 2020] Goncharova, A. S., Honigmann, A., Jug, F., and Krull, A. (2020). Improving blind spot denoising for microscopy. In *Proceedings of the European Conference on Computer Vision, Workshops*.

[Huang et al., 2021] Huang, T., Li, S., Jia, X., Lu, H., and Liu, J. (2021). Neighbor2neighbor: Self-supervised denoising from single noisy images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14781–14790.

[Krull et al., 2018] Krull, A., Buchholz, T.-O., and Jug, F. (2018). Noise2void-learning denoising from single noisy images. 2019 ieee. In *CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2124–2132.

[Krull et al., 2019] Krull, A., Vicar, T., and Jug, F. (2019). Probabilistic noise2void: Unsupervised content-aware denoising. *arXiv preprint arXiv:1906.00651*.

[Kundur and Hatzinakos, 1996] Kundur, D. and Hatzinakos, D. (1996). Blind image deconvolution. *IEEE signal processing magazine*, 13(3):43–64.

[Laine et al., 2019] Laine, S., Karras, T., Lehtinen, J., and Aila, T. (2019). High-quality self-supervised deep image denoising. In *Advances in Neural Information Processing Systems*, pages 6970–6980.

[Lee and Jeong, 2020] Lee, K. and Jeong, W.-K. (2020). Noise2kernel: Adaptive self-supervised blind denoising using a dilated convolutional kernel architecture. *arXiv preprint arXiv:2012.03623*.

[Lehtinen et al., 2018] Lehtinen, J., Munkberg, J., Hasselgren, J., Laine, S., Karras, T., Aittala, M., and Aila, T. (2018). Noise2Noise: Learning image restoration without clean data. In *Proceedings of the 35th International Conference on Machine Learning*, pages 2965–2974.

[Meister, 2009] Meister, A. (2009). *Deconvolution problems in nonparametric statistics*. Springer.

[Moulines et al., 1997] Moulines, E., Cardoso, J.-F., and Gassiat, E. (1997). Maximum likelihood for blind separation and deconvolution of noisy signals using mixture models. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 5, pages 3617–3620. IEEE.

14

[Prakash et al., 2020a] Prakash, M., Krull, A., and Jug, F. (2020a). Divnoising: Diversity denoising with fully convolutional variational autoencoders. *arXiv preprint arXiv:2006.06072*.

[Prakash et al., 2020b] Prakash, M., Lalit, M., Tomancak, P., Krul, A., and Jug, F. (2020b). Fully unsupervised probabilistic noise2void. In *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, pages 154–158. IEEE.

[Quan et al., 2020] Quan, Y., Chen, M., Pang, T., and Ji, H. (2020). Self2self with dropout: Learning self-supervised denoising from single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

[Ronneberger et al., 2015] Ronneberger, O., Fischer, P., and Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer.

[Wang et al., 2004] Wang, Z., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P. (2004). Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612.

[Weigert et al., 2017] Weigert, M., Schmidt, U., Boothe, T., Andreas, M., Dibrov, A., Jain, A., Wilhelm, B., Schmidt, D., Broaddus, C., Culley, S., et al. (2017). Content-aware image restoration: Pushing the limits of fluorescence microscopy. biorxiv.

[Xie et al., 2020] Xie, Y., Wang, Z., and Ji, S. (2020). Noise2same: Optimizing a self-supervised bound for image denoising. *Advances in Neural Information Processing Systems (NeurIPS)*.

[Yue et al., 2019] Yue, Z., Yong, H., Zhao, Q., Meng, D., and Zhang, L. (2019). Variational denoising network: Toward blind noise modeling and removal. In *Advances in Neural Information Processing Systems*.

[Zhang et al., 2017] Zhang, K., Zuo, W., Chen, Y., Meng, D., and Zhang, L. (2017). Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing*, 26(7):3142–3155.

[Zhou et al., 2020] Zhou, R., El Helou, M., Sage, D., Laroche, T., Seitz, A., and Süsstrunk, S. (2020). W2S: Microscopy data with joint denoising and super-resolution for widefield to SIM mapping. In *ECCV*.

A   Additional Implementation details

A.1   Networks and training

D-net architecture details The architecture is based on U-net [Ronneberger et al., 2015]. We propose several changes from the original version: we do not crop the image and use zero-padding instead, we use 2 levels of contractions/expansions with 64 filters, expansions are performed by an upsampling layer with nearest-neighbor approximation directly followed by 2x2 convolution. We also add two layers of 1x1 convolution with 64 filters and ReLU activation at the end of the network, and set no activation function at the output layer.

N-net architecture details In the case of a Gaussian noise, the N-net is composed 3 successive blocks, each block being composed of two 1x1 convolutions layers of 64 filters, each followed by a non-linear activation layer (alternatively tanh and leaky ReLU with alpha parameter set to 0.1). A convolution 1x1 with a single channel followed by an exponential activation function is placed after the last block (to ensure that the predicted σ is positive).

In the case where the N-net predicts a GMM with N components with weights $(\alpha_{k,\theta_n})_{1 \leqslant k \leqslant N}$, means $(\mu_k)_{1 \leqslant k \leqslant N}$ and variances $(\sigma^2_{k,\theta_n})_{1 \leqslant k \leqslant N}$, the second block is connected to three distinct blocks, each connected to a convolution 1x1 with:

- N channels, followed by an exponential activation function to predict σ_{k,θ_n}.
- N channels, followed by a softmax activation to predict α_{k,θ_n}.[4]
- N-1 channels to predict the distribution means μ_{k,θ_n}.

To ensure that the distribution is centered, the center of the last distribution is computed as

$$\mu_{N,\theta_n} = -\frac{1}{\alpha_{N,\theta_n}} \sum_{k=1}^{N-1} \alpha_{k,\theta_n} \cdot \mu_{k,\theta_n} \,.$$

B   Baselines

- DecoNoisng was trained using the source code provided by the authors, using no positivity constraint: `https://github.com/juglab/DecoNoising`. For N2V we used the same code without the convolution.

- FBI-Denoiser was trained using the code provided by authors: `https://github.com/csm9493/FBI-Denoiser/tree/de86420934a2416d4052dfa1298334af0d2ca49f`. We used same training options as for FMD datasets that are also fluorescence microscopy datasets. We tried to double the number of epochs without improvement.

- Neighbor2Neighbor was trained using the code provided by the authors: `https://github.com/TaoHuang2018/Neighbor2Neighbor/tree/e66389ebef1e64d306d8fcb95f096cf427243452`. As this method was not trained on fluorescence microscopy datasets, we tried to optimise hyperparameters as best as we could on dataset W2S-1 and used the same hyperparameters for all datasets. More

---

[4]When $N = 2$, only one channel is used and followed by a sigmoid activation function.

precisely we tested $\gamma = 1$ and $\gamma = 2$, an initial learning rate of $3 \cdot 10^{-4}$ or $1 \cdot 10^{-4}$ as suggested by the authors. Moreover, we also increased the number of epochs from 100 to 1000, and tried to decrease the learning rate to $1 \cdot 10^{-6}$ instead of $2 \cdot 10^{-5}$ (which surprisingly led to much lower performances). In the end, we used $\gamma = 1$, an initial learning rate of $3 \cdot 10^{-4}$, a minimal learning rate of $2 \cdot 10^{-5}$ and 1000 epochs.

C Datasets

C.1 Experimental Datasets

Datasets published along with the PN2V [Krull et al., 2019]

- *Convallaria* dataset, referred to as *PN2V-C* is composed of 100 images of size 1024x1024. Evaluation subset is: $Y \in [0\,;512], X \in [0\,;512]$.

- *Mouse skull nuclei* referred to as *PN2V-MN* is composed 200 images of size 512x512. Evaluation subset is: $Y \in [0\,;512], X \in [0\,;256]$.

- *Mouse Actin* referred to as *PN2V-MA* is composed of 100 images of size 1024x1024. Evaluation subset is: $Y \in [0\,;1024], X \in [0\,;512]$.

The *PN2V-C* and *PN2V-MA* datasets are acquired on a spinning disc confocal microscope and *PN2V-MN* dataset is acquired with a point scanning confocal microscope. Datasets can be respectively downloaded at: `https://doi.org/10.5281/zenodo.5156913`, `https://doi.org/10.5281/zenodo.5156960` and `https://doi.org/10.5281/zenodo.5156937`

We observed that datasets acquired with a spinning disc confocal microscope display axial noise correlation (see 4.2).

Datasets published in [Zhou et al., 2020]   We used the 16-bit raw images kindly provided by the authors. The dataset is composed of 120 FOV of 400 observations of size 512x512 pixels. The first 80 are used for training and the last 40 for evaluation. Following the authors, for each FOV, only the observation of index 249 is used for training and evaluation images are acquired with a electron-multiplying charge-coupled device camera on a wide-field microscope. It can be downloaded at: `https://datasets.epfl.ch/w2s/W2S_raw.zip`

Normalization   Images were normalized using the modal value as center and the difference between modal value and 95% percentile as scale factor, computed on the whole dataset. This is relevant in fluorescence microscopy data where signal is often less abundant than background with proportion that vary among images and signal distribution often has a heavy tail towards high values.

Metrics   For the 6 chosen datasets, images are encoded in 16-bit. PSNR is defined as

$$\text{PSNR} = 10 \log_{10}(d/\text{MSE})\,,$$

with d the maximum possible pixel value range of the image and MSE the mean squared error. For 8-bit encoded images d is simply 255, and for 16-bit images it would be 65635 but this does not correspond to the actual possible range of microscopy data, thus the actual range of values of each ground truth image is used. This is also what is done in [Goncharova et al., 2020] as we obtain the same PSNR values for raw images. The same applies for SSIM computation.

C.2 Synthetic noise datasets

- Additive Gaussian: $Y = X + \varepsilon$ with $\varepsilon \sim \mathcal{N}(0, \sigma^2)$, $\sigma = 20$.

- Poisson-Gaussian: $Y = X + (\alpha * (X - \underline{X}) + \eta^2)^{1/2}\varepsilon$ with $\varepsilon \sim \mathcal{N}(0, 1)$, $\alpha = 5, \eta = 12$ and $\underline{X}$ being the minimal value of the ground truth on the whole dataset.

- Speckle: $X = X + (X - \underline{X})\varepsilon$ with $\varepsilon \sim \mathcal{N}(0, \sigma^2)$, $\sigma = 0.405$ and $\underline{X}$ being the minimal value of the ground truth on the whole dataset.

D Impact of the N-net

Figure 5: PSNR at each epoch on the test set for N-net (G1) or D-net only (L2). A: training on real noise. B, C: training on different models/levels of synthetic noise added to ground truth image. B: *Gaussian* and *Speckle* correspond to noises presented in 5.1, that have no and high signal dependency respectively. C: *Speckle* with variance multiplied by 0.5, 2 and 4. D: Mean and SEM over 5 runs.

E Noise estimation

This section contains the figures corresponding to Fig 3 for each dataset. Note that they are computed using regular signal value bins and excluding signal values greater to the 99.5% percentile of the dataset so that there are enough observed samples in each bin to compute statistically significant metrics.

Table 3: Evaluation of the impact of N-net on performances with PSNR/SSIM metrics. Comparison between our method predicting a GMM with 1, 2, or 3 components (respectively G1, G2, G3) and D-net only trained with L2 loss (with training and prediction scheme unchanged). Best PSNR $\pm 0.1 dB$ / SSIM $\pm 1\text{‰}$ scores are underlined.

DATASET	NOISY	OURS (G1)	OURS (G2)	OURS (G3)
PN2V-C	28.98 / 0.7713	38.33 / <u>0.9754</u>	<u>38.47</u> / 0.9738	38.28 / <u>0.9756</u>
PN2V-MN	28.10 / 0.6836	39.08 / <u>0.9776</u>	<u>39.22</u> / <u>0.9779</u>	39.18 / <u>0.9780</u>
PN2V-MA	23.71 / 0.3731	<u>34.79</u> / <u>0.8905</u>	34.68 / 0.8880	<u>34.70</u> / 0.8877
W2S-1	21.85 / 0.3490	<u>35.33</u> / <u>0.9619</u>	35.27 / <u>0.9623</u>	35.27 / <u>0.9624</u>
W2S-2	19.33 / 0.2256	<u>33.46</u> / <u>0.8867</u>	<u>33.48</u> / <u>0.8871</u>	<u>33.47</u> / <u>0.8871</u>
W2S-3	20.39 / 0.2232	<u>36.57</u> / <u>0.9263</u>	<u>36.60</u> / <u>0.9269</u>	<u>36.59</u> / <u>0.9269</u>

F   Code

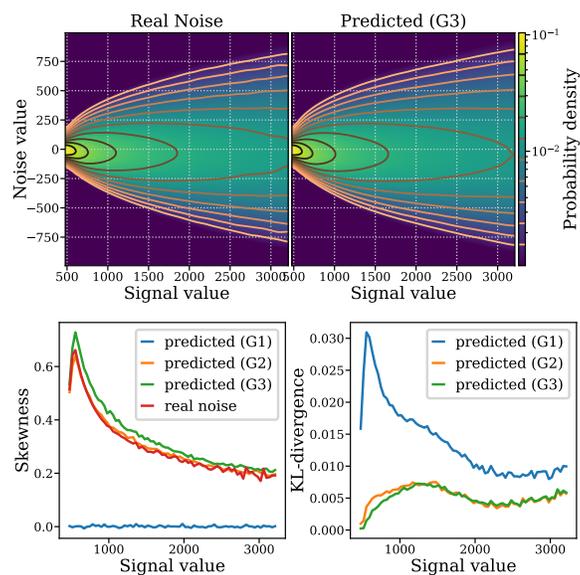The code will be available after the peer review process.

Figure 6: Real noise estimation for dataset *PN2V-C*. See main text Fig 3.
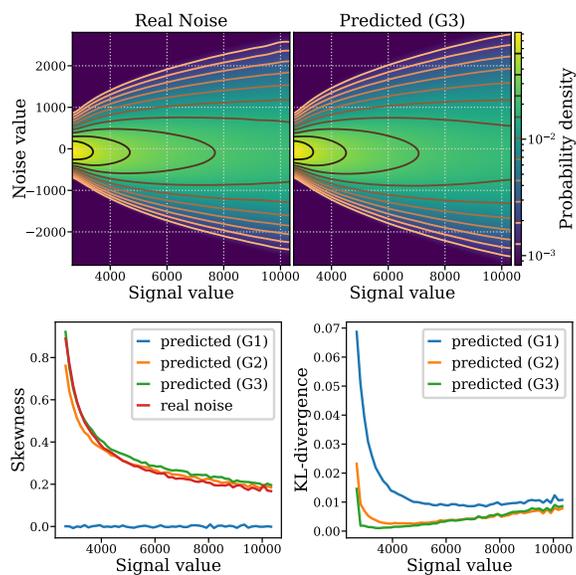

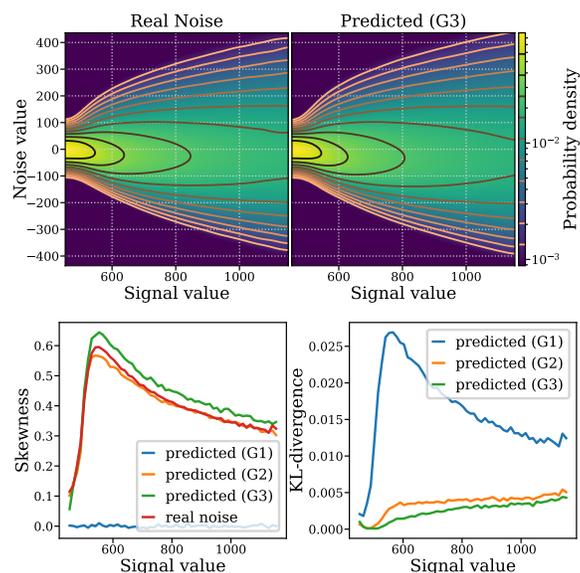Figure 7: Real noise estimation for dataset *PN2V-MN*. See main text Fig 3.

Figure 8: Real noise estimation for dataset *PN2V-MA*. See main text Fig 3.
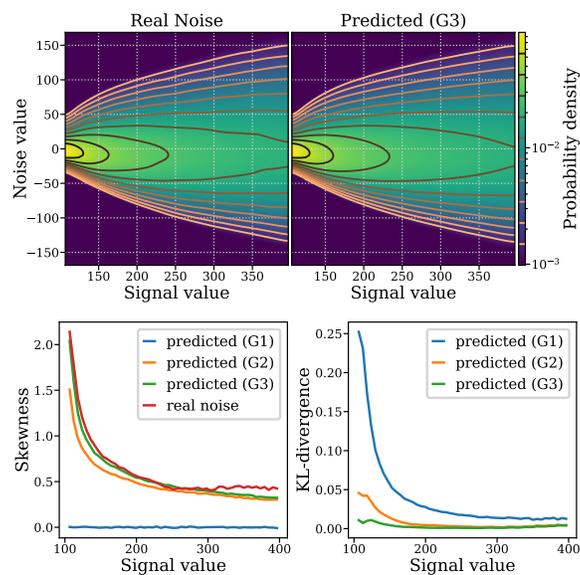

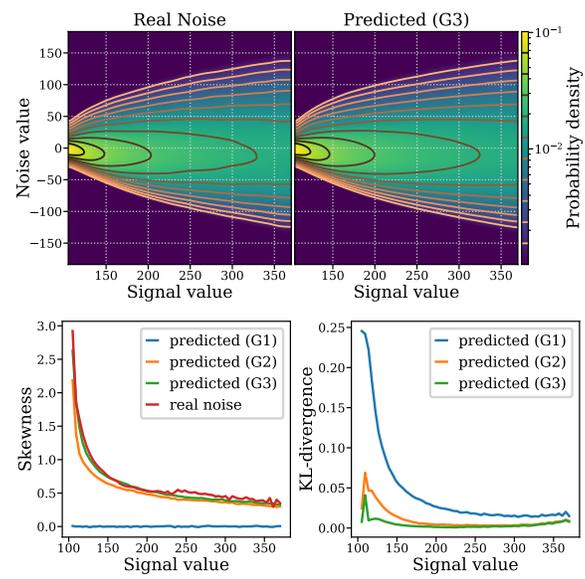Figure 9: Real noise estimation for dataset *W2S-2*. See main text Fig 3.

Figure 10: Real noise estimation for dataset *W2S-3*. See main text Fig 3.