

Preconditioned Plug-and-Play ADMM with Locally Adjustable Denoiser for Image Restoration

Mikael Le Pendu and Christine Guillemot

Abstract—Plug-and-Play optimization recently emerged as a powerful technique for solving inverse problems by plugging a denoiser into a classical optimization algorithm. The denoiser accounts for the regularization and therefore implicitly determines the prior knowledge on the data, hence replacing typical handcrafted priors. In this paper, we extend the concept of plug-and-play optimization to use denoisers that can be parameterized for non-constant noise variance. In that aim, we introduce a preconditioning of the ADMM algorithm, which mathematically justifies the use of such an adjustable denoiser. We additionally propose a procedure for training a convolutional neural network for high quality non-blind image denoising that also allows for pixel-wise control of the noise standard deviation. We show that our pixel-wise adjustable denoiser, along with a suitable preconditioning strategy, can further improve the plug-and-play ADMM approach for several applications, including image completion, interpolation, demosaicing and Poisson denoising.

Index Terms—Plug-and-Play Prior, Inverse Problems, Preconditioning, ADMM, Denoising, Interpolation, Demosaicing, Image Completion.



1 INTRODUCTION

Inverse problems arising in image restoration require the use of prior knowledge on images in order to determine the most likely solutions among an infinity of possibilities. Traditional optimization methods rely on priors modeled as convex regularization functions such as the total variation, encouraging smoothness, or the l_1 norm for sparsity. In this context, convex optimization algorithms have played an important role and their convergence properties have been well-established. However, the methods based on “hand-crafted” convex priors are now significantly outperformed by deep neural networks that directly learn an inverse mapping from the degraded measurements to the solution space. Here, the prior knowledge on images and the degradation to recover from do not need a formal mathematical definition. Instead, they are implicitly taken into account when training the network from a large dataset of degraded images (i.e. network’s input) along with their original versions (i.e. network’s output). This approach has enabled a significant leap in performance for common image restoration problems including denoising, demosaicing, super-resolution, etc. However, different neural networks must be carefully designed and trained for each problem specifically. Furthermore, it is usually impossible to interpret what task is performed by the different layers or sub-modules of the trained network which acts as a black box.

In order to increase the interpretability and genericity of deep neural networks, the field is evolving towards methods that combine them with traditional optimization algorithms. A popular approach, seen for example in [1], [2], [3], [4], [5], [6], [7], consists in defining so-called “unrolled” neural

networks that simulate several iterations of an optimization algorithm. For instance, in unrolled gradient descent, the gradient of the data-fidelity term can be computed knowing its closed form expression, while a convolutional neural network (CNN) is used to simulate the gradient computation of a regularization term. End-to-end training of this regularizing network is then performed by “unrolling” a given number of iterations of the optimization algorithm. Here, the trained CNN is meant to represent only prior knowledge on images, which does not depend on a specific image restoration task. However, in practice these networks require re-training for each task, meaning that task-specific features are also learnt in the end-to-end unrolled training.

A more universal approach referred to as “Plug-and-Play” (PnP), relies on proximal algorithms such as the Alternating Direction Method of Multipliers (ADMM) [8] or Half Quadratic Splitting (HQS) [9], where the proximal operator of the regularization term is usually replaced by a denoiser as in [10], [11], [12], [13], [14]. While Venkatakrishnan et al. [10] introduced this approach using traditional denoising techniques including BM3D [15] or K-SVD [16], more recent methods attain higher performance thanks to neural networks trained for denoising. The main limitation, however, is that the objective function to minimize is typically non-convex when the proximal operator of the regularization term is defined as a high performance denoiser. Hence the proximal algorithms used may not provide the optimal solution.

In this paper, we propose a preconditioned ADMM in order to improve the algorithm’s performance in this challenging “Plug-and-Play” scenario. While the proposed formulation remains mathematically equivalent to the original problem, the preconditioning makes it possible to use a denoiser that takes into account spatially varying noise levels. This is particularly advantageous for applications

The authors are with Inria Rennes – Bretagne-Atlantique
263 Avenue Général Leclerc, 35042 Rennes Cedex, France (e-mail: first-name.lastname@inria.fr).

where the variance of the approximation error made at each iteration varies spatially in a predictable way. For instance, in image demosaicing, completion, or interpolation, only the unknown pixels that require interpolation are expected to have an approximation error. Hence, the pattern of known and unknown pixels can be used in our scheme to define a sensible preconditioning matrix. Likewise, when an image is corrupted by Poisson noise, the noise variance at each pixel is known to be proportional to the noiseless pixel intensity. In summary, the main contributions of our paper are as follows:

- We formulate a preconditioned ADMM, which enables the use of a denoiser that can take into account spatially varying standard deviation.
- We propose a procedure for training a denoising neural network that takes as input a map of pixel-wise standard deviation, along with the image to be denoised.
- We demonstrate the effectiveness of the approach and define suitable preconditioning schemes for several applications: image completion, interpolation, demosaicing and Poisson denoising. Note that for Poisson denoising, we present further derivations accounting for the negative log-likelihood of the Poisson distribution, which replaces the conventional least square data fidelity term (only suitable for Gaussian noise).

2 BACKGROUND: PLUG-AND-PLAY ADMM

2.1 Plug-and-Play Prior

Let us consider the linear inverse problem of recovering a clean image $\hat{x} \in \mathbb{R}^n$ from its degraded measurements $\mathbf{b} \in \mathbb{R}^m$ obtained with the degradation model $\mathbf{b} = \mathbf{A}\mathbf{x} + \nu$, where $\mathbf{A} \in \mathbb{R}^{m \times n}$ is the degradation matrix, $\nu \in \mathbb{R}^m$ is additive white Gaussian noise with standard deviation σ , and \mathbf{x} is the unknown ground truth image arranged as a vector. The problem is generally formulated as the Maximum a Posteriori (MAP) estimation of \mathbf{x} , given its prior probability density function p . The MAP is given by:

$$\hat{x} = \arg \max_{\mathbf{x}} p(\mathbf{x}|\mathbf{b}) = \arg \max_{\mathbf{x}} p(\mathbf{b}|\mathbf{x}) \cdot p(\mathbf{x}), \quad (1)$$

$$= \arg \min_{\mathbf{x}} -\ln(p(\mathbf{b}|\mathbf{x})) - \ln(p(\mathbf{x})), \quad (2)$$

$$= \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \sigma^2 \mathcal{R}(\mathbf{x}). \quad (3)$$

In practice, the prior distribution p is not used directly, but is represented by the regularizer $\mathcal{R}(\mathbf{x}) = -\ln(p(\mathbf{x}))$, which penalizes unlikely solutions.

The PnP approach introduced in [10] goes one step further by removing the need for an explicit definition of the regularizer. This is made possible by the use of proximal algorithms such as ADMM in which the regularizer only appears in the evaluation of the proximal operator of $\gamma\mathcal{R}$ for some scalar factor γ . This operator is defined as:

$$\text{prox}_{\gamma\mathcal{R}}(\mathbf{u}) = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{x} - \mathbf{u}\|_2^2 + \gamma \cdot \mathcal{R}(\mathbf{x}). \quad (4)$$

It can be noted that the expression of $\text{prox}_{\gamma\mathcal{R}}$ is a particular case of Eq. (3), where the degradation matrix \mathbf{A} is the identity matrix, i.e. the degradation model only consists in the

addition of white Gaussian noise. In other words $\text{prox}_{\gamma\mathcal{R}}$ is the MAP denoiser assuming additive white Gaussian noise of variance γ and given the underlying prior $p(\mathbf{x}) = e^{-\mathcal{R}(\mathbf{x})}$. Based on this observation, the authors of [10] replaced $\text{prox}_{\gamma\mathcal{R}}$ by a standard image denoiser which implicitly determines the prior distribution. In the following, we will note the denoiser as a function of the image and the noise standard deviation:

$$F_{\mathcal{R}}(\mathbf{u}, \sqrt{\gamma}) = \text{prox}_{\gamma\mathcal{R}}(\mathbf{u}) \quad (5)$$

More recent works use a trained CNN to represent the proximal operator. For instance, in [13], adversarial training is used to learn a projection operator. Here, the assumption is that the regularizer is the indicator function of the set of “natural images”. Hence, the corresponding proximal operator projects the input image to the closest “natural image”, which can be seen as a form of blind denoising. While this avoids the need for a parameter γ , it does not allow for controlling the denoising “strength” within the algorithm. In [11], a neural network called DRUNet, that combines Res-Net [17] and U-Net [18] architectures, is trained for the task of Gaussian denoising. The noise standard deviation is given as a constant input map in addition to the noisy image. Hence, compared to approaches based on denoising CNNs trained for a single noise level [19], or for blind denoising [13], [20], [21], the method in [11] can control the algorithm’s parameters more precisely at each iteration, while keeping the highest denoising performance. Since a single network is used, it also simplifies the method in [12], where 25 denoisers are trained for different standard deviations in order to cover a larger range of noise levels.

In this paper, we argue that a denoiser parameterized with an input map for standard deviation can even be trained to control the noise level at each pixel independently. Such a denoiser further improves the PnP scheme for several applications thanks to the proposed preconditioning.

2.2 ADMM Formulation

In this section, we present the classical ADMM formulation in the case of the least squares inverse problem in Eq. (3). The same notations will be used throughout the paper. In order to use the ADMM, the problem is first cast into a constrained optimization problem by splitting the two terms:

$$\begin{aligned} \hat{x} = \arg \min_{\mathbf{x}, \mathbf{y}} \quad & \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \sigma^2 \mathcal{R}(\mathbf{y}) \\ \text{subject to} \quad & \mathbf{x} = \mathbf{y}, \end{aligned} \quad (6)$$

The constraint $\mathbf{x} = \mathbf{y}$ is included in the optimization by defining the augmented Lagrangian function \mathcal{L} , which introduces a dual variable $\mathbf{l} \in \mathbb{R}^n$ and a penalty parameter ρ :

$$\begin{aligned} \mathcal{L}(\mathbf{x}, \mathbf{y}, \mathbf{l}) = \quad & \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \sigma^2 \mathcal{R}(\mathbf{y}) \\ & + \mathbf{l}^\top (\mathbf{x} - \mathbf{y}) + \frac{\rho}{2} \|\mathbf{x} - \mathbf{y}\|_2^2, \\ = \quad & \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \sigma^2 \mathcal{R}(\mathbf{y}) + \frac{\rho}{2} \left\| \mathbf{x} - \mathbf{y} + \frac{\mathbf{l}}{\rho} \right\|_2^2 - \frac{\|\mathbf{l}\|_2^2}{2\rho}. \end{aligned} \quad (7)$$

The ADMM method consists in alternatively minimizing the augmented Lagrangian \mathcal{L} for the variables \mathbf{x} and \mathbf{y} separately, and by updating the dual variable \mathbf{l} . In practice, the penalty parameter ρ may also be updated at each iteration to accelerate the convergence. The ADMM iteration thus reads as:

$$\mathbf{x}^{k+1} = \arg \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \rho^k \left\| \mathbf{x} - \left(\mathbf{y}^k - \frac{\mathbf{l}^k}{\rho^k} \right) \right\|_2^2, \quad (9)$$

$$\mathbf{y}^{k+1} = \arg \min_{\mathbf{y}} \frac{1}{2} \left\| \mathbf{y} - \left(\mathbf{x}^{k+1} + \frac{\mathbf{l}^k}{\rho^k} \right) \right\|_2^2 + \frac{\sigma^2}{\rho^k} \mathcal{R}(\mathbf{y}), \quad (10)$$

$$\mathbf{l}^{k+1} = \mathbf{l}^k + \rho^k (\mathbf{x}^{k+1} - \mathbf{y}^{k+1}), \quad (11)$$

$$\rho^{k+1} = \rho^k \cdot \alpha, \quad \text{with } \alpha > 1, \quad (12)$$

where the dual variable \mathbf{l} is typically zero-initialized. An initialization of \mathbf{x} is also required and is often performed as $\mathbf{x}_0 = \mathbf{A}^\top \mathbf{b}$, but a finer initialization strategy may be used depending on the problem.

The y-update in Eq. (10) can be equivalently written using a proximal operator:

$$\mathbf{y}^{k+1} = \text{prox}_{\gamma \mathcal{R}}(\mathbf{x}^{k+1} + \mathbf{l}^k / \rho^k), \quad (13)$$

where $\gamma = \sigma^2 / \rho^k$. As seen in Section 2.1, this step can be performed in the PnP approach by denoising the image $\mathbf{x}^{k+1} + \mathbf{l}^k / \rho^k$ using a Gaussian denoiser for a noise standard deviation $\sqrt{\gamma} = \sigma / \sqrt{\rho^k}$.

3 PRECONDITIONED PLUG-AND-PLAY ADMM

In several problems, we have additional knowledge about locally varying noise level. For instance, in image completion or interpolation problems, the pixels already sampled in \mathbf{b} are known without error. Thus, at a given iteration, only the pixels at unknown positions need to be denoised. However, in the original ADMM formulation the denoising step considers the same noise level $\sigma / \sqrt{\rho^k}$ at each pixel. Hence, in order to finely tune the denoising effect and improve the performance of the PnP-ADMM, we reformulate the problem with a preconditioner that takes into account the knowledge of a variable noise level.

3.1 ADMM Reformulation

Let us consider the following problem, mathematically equivalent to Eq. (3):

$$\tilde{\mathbf{x}} = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{A}\mathbf{P}\mathbf{x} - \mathbf{b}\|_2^2 + \sigma^2 \mathcal{R}(\mathbf{P}\mathbf{x}), \quad (14)$$

$$\hat{\mathbf{x}} = \mathbf{P}\tilde{\mathbf{x}}, \quad (15)$$

where \mathbf{P} is a diagonal preconditioning matrix.

Similarly to the Section. 2.2, we can express the ADMM algorithm for solving Eq. (14) by replacing the updates of \mathbf{x} and \mathbf{y} in Eqs. (9) and (10) with respectively:

$$\mathbf{x}^{k+1} = \arg \min_{\mathbf{x}} \|\mathbf{A}\mathbf{P}\mathbf{x} - \mathbf{b}\|_2^2 + \rho^k \left\| \mathbf{x} - \left(\mathbf{y}^k - \frac{\mathbf{l}^k}{\rho^k} \right) \right\|_2^2, \quad (16)$$

$$\mathbf{y}^{k+1} = \arg \min_{\mathbf{y}} \frac{1}{2} \left\| \mathbf{y} - \left(\mathbf{x}^{k+1} + \frac{\mathbf{l}^k}{\rho^k} \right) \right\|_2^2 + \frac{\sigma^2}{\rho^k} \mathcal{R}(\mathbf{P}\mathbf{y}). \quad (17)$$

The variable \mathbf{x} can be updated directly using the well-known closed form solution of Eq. (16):

$$\mathbf{x}^{k+1} = (\mathbf{P}^\top \mathbf{A}^\top \mathbf{A} \mathbf{P} + \rho^k \mathbf{I})^{-1} (\mathbf{P}^\top \mathbf{A}^\top \mathbf{b} + \rho^k \mathbf{y}^k - \mathbf{l}^k). \quad (18)$$

However, for updating \mathbf{y} , the denoiser $F_{\mathcal{R}}$ from Eq. (5), is no longer suitable because of the matrix \mathbf{P} within the regularization term. The next section presents how the y-update can still be performed using a more general denoiser that considers Gaussian noise with variable variance.

3.2 New Prior Term Sub-Problem

3.2.1 General Gaussian Denoiser

Let us first define the general expression of a denoiser for Gaussian noise with covariance matrix Σ . The MAP estimate from a noisy image \mathbf{u} can be derived similarly to Eqs. (1)-(3), where the likelihood $p(\mathbf{u}|\mathbf{x})$ is a multivariate Gaussian distribution, i.e. $p(\mathbf{u}|\mathbf{x}) \propto e^{-\frac{1}{2} \|\Sigma^{-1/2}(\mathbf{u}-\mathbf{x})\|_2^2}$. Including the prior and taking the negative logarithm gives the expression of the denoiser:

$$G_{\mathcal{R}}(\mathbf{u}, \Sigma^{1/2}) = \arg \min_{\mathbf{x}} \frac{1}{2} \left\| \Sigma^{-1/2}(\mathbf{u} - \mathbf{x}) \right\|_2^2 + \mathcal{R}(\mathbf{x}), \quad (19)$$

Here, we will only consider the case of a diagonal matrix Σ . Hence, we still assume independent noise at each pixel (but not identically distributed), and the diagonal of Σ corresponds to pixel-wise variance. Note that the denoiser $F_{\mathcal{R}}$ is equivalently defined as $F_{\mathcal{R}}(\mathbf{x}, \sigma) = G_{\mathcal{R}}(\mathbf{x}, \sigma \mathbf{I})$ for a constant standard deviation σ . A practical image denoiser $G_{\mathcal{R}}$ can be obtained by training a CNN that takes as input a map of the pixel-wise noise levels in addition to the noisy image. A training procedure for such a denoiser is presented in Section 4.

3.2.2 Sub-Problem Solution

Now, let us rewrite Eq. (17) using the change of variable $\tilde{\mathbf{y}} = \mathbf{P}\mathbf{y}$, which gives $\mathbf{y} = \mathbf{P}^{-1}\tilde{\mathbf{y}}$. Thus, we have $\mathbf{y}^{k+1} = \mathbf{P}^{-1}\tilde{\mathbf{y}}^{k+1}$ and $\tilde{\mathbf{y}}^{k+1} = \arg \min_{\tilde{\mathbf{y}}} f(\tilde{\mathbf{y}})$ with:

$$f(\tilde{\mathbf{y}}) = \frac{\rho^k}{2\sigma^2} \left\| \mathbf{P}^{-1}\tilde{\mathbf{y}} - \left(\mathbf{x}^{k+1} + \frac{\mathbf{l}^k}{\rho^k} \right) \right\|_2^2 + \mathcal{R}(\tilde{\mathbf{y}}), \quad (20)$$

$$= \frac{1}{2} \left\| \frac{\sqrt{\rho^k}}{\sigma} \mathbf{P}^{-1} \left(\tilde{\mathbf{y}} - \mathbf{P} \left(\mathbf{x}^{k+1} + \frac{\mathbf{l}^k}{\rho^k} \right) \right) \right\|_2^2 + \mathcal{R}(\tilde{\mathbf{y}}). \quad (21)$$

Given the denoiser $G_{\mathcal{R}}$ in Eq. (19) we can then rewrite the y-update step as:

$$\mathbf{y}^{k+1} = \mathbf{P}^{-1} G_{\mathcal{R}} \left(\mathbf{P} \left(\mathbf{x}^{k+1} + \frac{\mathbf{l}^k}{\rho^k} \right), \frac{\sigma}{\sqrt{\rho^k}} \mathbf{P} \right) \quad (22)$$

Therefore, the preconditioned ADMM still solves the original problem but introduces a denoising step that assumes noise with spatially varying standard deviation. The standard deviation can be adjusted per pixel via the diagonal matrix \mathbf{P} and the global parameters σ and ρ^k . A suitable choice of the preconditioning matrix can be made depending on the problem in order to improve the performance of the PnP-ADMM.

3.3 Variables Interpretation and Initialization

It should be noted that in the preconditioned problem, an image \mathbf{x}^k only estimates the intermediate variable $\hat{\mathbf{x}}$, but the final result is $\hat{\mathbf{x}} = \mathbf{P}\hat{\mathbf{x}}$ (see Eq. (15)). Hence, \mathbf{x}^k (similarly \mathbf{y}^k) may not be a plausible image. However, the input of the denoiser is $\mathbf{P}(\mathbf{x}^{k-1} + \mathbf{l}^k/\rho^k)$, where the multiplication by \mathbf{P} rescales the image pixels consistently with the “natural image” $\hat{\mathbf{x}}$, as expected by the denoiser.

Inversely, for the initialization, the inverse scaling must be performed. Given an initial estimate $\hat{\mathbf{x}}^0$ of the true image $\hat{\mathbf{x}}$, a good initialization for our problem is thus $\mathbf{x}^0 = \mathbf{P}^{-1}\hat{\mathbf{x}}^0$.

4 DEEP LOCALLY ADJUSTABLE DENOISER

4.1 Denoising Network for Known Noise Level

Our algorithm relies on a deep neural network for solving the prior term sub-problem which corresponds to a Gaussian denoising problem where the noise level (i.e. standard deviation) is known and can be adjusted for each pixel. Here, we use the DRUNet architecture proposed in [11]. The overall structure consists of a U-Net where each level contains sequences of 4 residual-blocks and either a down-sampling layer (first branch of the U) or an upsampling layer (second branch).

The DRUNet network structure in [11] conforms to the definition of the denoisers $F_{\mathcal{R}}$ (Eq. (5)) or $G_{\mathcal{R}}$ (Eq. (19)) by taking as input the concatenation (in the channel dimension) of the noisy image and a noise level map whose pixels’ values are equal to the noise standard deviation. The advantage of using an input noise level map is that a single generic model can be trained to perform optimally for a large range of noise levels. However, in [11] the model is only trained considering constant noise level maps since their algorithm only uses the constant denoiser $F_{\mathcal{R}}$. In the rest of the paper, we refer to this network as *DRUNet-cst*.

In order to use our preconditioning approach, a locally adjustable denoiser $G_{\mathcal{R}}$ is required. Therefore, the training process must be adapted so that each input sample consists of an arbitrary noise level map along with an image corrupted with the corresponding Gaussian noise level for each pixel. We describe in the next section how to randomly generate suitable patterns so that the trained model generalizes well for any arbitrary noise level map.

4.2 Noise Level Map Generation

Let us first consider the case of a constant noise level map as in [11]. Here, all the pixels are equal to the same random variable S that can be simply defined as:

$$S = 2\mu \cdot X, \quad (23)$$

where X is a uniformly distributed random variable in the range $[0, 1]$. The parameter μ is equal to the expectation of the random variable S . It can be selected to train a denoiser that is sufficiently generic for all the noise levels σ in the range $[0, 2\mu]$.

Now, in order to train a locally adjustable denoiser $G_{\mathcal{R}}$, we generate a random map for each training image using the following random process:

$$S_i = 2\mu \cdot (X_i \cdot (1 - W) + O \cdot W), \quad (24)$$

where i is the pixel index, O , W and X_i (for each pixel i) are independent random variables with uniform distributions in the range $[0, 1]$. One can verify that μ is the expectation of S_i , and the range of possible values of S_i is $[0, 2\mu]$ similarly to the previous case with constant noise level.

The random weight W allows us to adjust the variance of the noise level map (i.e. lower weight corresponding to higher variance). The random offset O is also necessary to reduce the correlation between the variance and the mean of the generated maps. In particular, the offset enables the generation of maps with a low variance but a high average level, so that the trained denoiser generalizes well even in the constant noise level case. Hence, the method makes it possible to train the network with noise level maps covering a wide range of first and second order statistics, which prevents overfitting for a specific type of pattern. In the paper, we refer to the network trained with this method as *DRUNet-var*.

For some applications, the noise standard deviation may also need to be adjusted depending on the color component. For these applications, we train another network called *DRUNet-var-RBG* that takes an input noise level map for each color component. For the training, the noise level maps are generated for each component separately using the random process in Eq. (24).

4.3 Training Details

The training is performed by generating a noise level map randomly for each input image, as described in Section 4.2. Considering pixel data in the range $[0, 1]$, we use the parameter $\mu = 25/255$ in Eq. (23) (for the *DRUNet-cst* denoiser) or Eq. (24) (for the *DRUNet-var* and *DRUNet-var-RBG* denoisers). The maximum noise level is thus $2\mu = 50/255$. Gaussian noise with standard deviation defined by the noise level map is then added to the input image.

The remaining details of the training procedure are the same as described in [11]. We use the same large dataset of 8694 images, including the Berkeley Segmentation Dataset BSDS500 (without the validation images) [22], the Waterloo Exploration Database [23], the DIV2K dataset [24], and the Flickr2K dataset [25]. We minimize the l_1 loss between the denoiser’s output and the ground truth image using the ADAM optimizer [26] with a learning rate initialized at $1e - 4$ and decreased by half every 100000 iterations. The

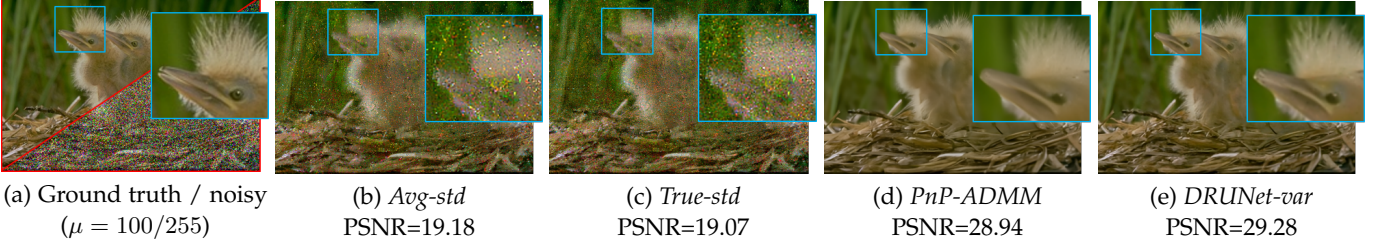


Fig. 1. Denoising results with pixel-wise standard deviation randomly selected in $[0, 2\mu]$. The average noise standard deviation is $\mu = 100/255$. For the *Avg-std*, *True-std* and *PnP-ADMM* schemes in (b-d), the network *DRUNet-cst* is used. In (e), the network *DRUNet-var* is used with the *True-std* scheme.

TABLE 1

Denoising performance (average PSNR over each dataset) for Gaussian noise with constant noise level (standard deviation $\sigma = 10, 30$ or 50).

Dataset noise level σ	CBSD68			Set5		
	10	30	50	10	30	50
CBM3D [15]	35.90	29.91	27.51	36.02	30.97	28.75
FFDNet [27]	36.14	30.32	27.97	36.16	31.35	29.24
RNAN [28]	36.43	30.65	28.30	36.62	31.77	29.61
<i>DRUNet-cst</i> [11]	36.51	30.79	28.48	36.67	31.90	29.80
<i>DRUNet-var</i>	36.51	30.80	28.48	36.67	31.91	29.79
<i>DRUNet-var-RBG</i>	36.47	30.78	28.46	36.61	31.90	29.77

training is stopped when the learning rate is smaller than $5e - 7$. Finally, each iteration of the training uses a batch of 16 patches of size 128×128 randomly selected from the images of the training dataset.

4.4 Denoising Performance

Here, we evaluate the performance of our denoisers *DRUNet-var* and *DRUNet-var-RBG* generalized for variable standard deviation, in comparison with *DRUNet-cst* [11] as well as other state-of-the-art denoisers that assume a constant noise level. These include the BM3D denoiser (noted CBM3D for color images) [15] and the two recent CNN-based methods FFDNet [27] and RNAN [28]. Similarly to *DRUNet-cst*, the FFDNet denoiser can be parameterized at inference time knowing the noise standard deviation. On the other hand, RNAN requires a separate training for each noise level, but has shown higher denoising performance thanks to a non-local attention module.

We perform our evaluations with the widely used Set5 [29] and CBSD68 [30] test datasets.

4.4.1 Constant Noise Level

First, let us consider input images corrupted with constant noise level. The results in Table 1 show that our locally adjustable denoisers *DRUNet-var* and *DRUNet-var-RBG* do not have significantly degraded performance compared to the reference network *DRUNet-cst*, despite being trained for more generic noise level maps. Only a very moderate loss is observed for the *DRUNet-var-RBG* denoiser compared to *DRUNet-cst*, mostly for the lowest and highest noise levels ($\sigma = 10$ or 50). Nevertheless, *DRUNet-var-RBG* still outperforms the other state-of-the-art methods, even the RNAN denoiser trained specifically for each noise level.

4.4.2 Spatially Varying Noise Level

Now, let us show the advantage of the proposed training procedure by comparing the performances of our *DRUNet-var* denoiser with the *DRUNet-cst* version when the input

TABLE 2

Denoising performance (average PSNR over each dataset) for Gaussian noise with pixel-wise variable standard deviation in the range $[0, 2\mu]$.

		<i>DRUNet-cst</i> [11]			<i>DRUNet-var</i>
Dataset		<i>Avg-std</i>	<i>True-std</i>	<i>PnP-ADMM</i>	<i>True-std</i>
$\mu = \frac{25}{255}$	Set5	29.01	29.03	34.35	34.48
	BSD68	28.14	28.22	32.75	32.90
$\mu = \frac{50}{255}$	Set5	24.45	24.37	32.03	32.24
	BSD68	23.59	23.60	30.09	30.27
$\mu = \frac{100}{255}$	Set5	19.82	19.58	29.69	30.06
	BSD68	19.03	18.88	27.73	28.01

noise standard deviation varies spatially. For these tests, we randomly select a standard deviation in the range $[0, 2\mu]$ for each pixel. For a complete comparison with *DRUNet-cst*, we have tested this network in three different ways:

- *Avg-std*: using the average noise level μ as input.
- *True-std*: using the true noise level map as input.
- *PnP-ADMM*: solving the problem of denoising with variable noise level (Eq.(19)) using PnP-ADMM without preconditioning based on the *DRUNet-cst* network.

For the *PnP-ADMM*, we set the parameters using Eq. (33) so that the noise of highest standard deviation $\sigma_{den}^0 = 2\mu$ is removed at the first iteration, while successive iterations refine the result by decreasing the noise level of the denoiser down to σ_{den}^N . Here, we use $\sigma_{den}^N = \frac{2\mu}{3}$ for $N = 25$ iterations, which we found to give the best results for this application.

Table 2 and Fig. 1 present the results of *DRUNet-cst* in each of the three schemes, and *DRUNet-var* in the *True-std* scheme. As expected, the results using only the average standard deviation are unsatisfying, since the pixels with above average noise level are not sufficiently denoised, while details are not optimally preserved in less noisy regions. The same behavior is observed when using the true noise level map as input to *DRUNet-cst*. This confirms that a denoiser trained for constant noise level does not generalize when the noise standard deviation strongly varies between pixels. More satisfying results are obtained when using *DRUNet-cst* in the PnP-ADMM scheme. However, our *DRUNet-var* denoiser trained directly for this task can still retrieve more details, which indicates that the PnP-ADMM without preconditioning is sub-optimal. More realistic applications are presented in the rest of the paper to demonstrate that our locally adjustable denoisers also allow for improved performances in practical scenarios thanks to the preconditioned PnP-ADMM.

5 APPLICATIONS

In this section we present several examples of applications in image restoration. Although the proposed denoiser allows us to adjust of the variance locally, it still assumes an independent noise distribution for each pixel. Hence, in order to demonstrate the advantage of our preconditioning for the PnP-ADMM scheme, we restrict our study to applications where the error variance is expected to vary only in the pixel domain. While this excludes deblurring, or super-resolution from an anti-aliased downsampled image, many applications of high practical interest are still concerned. These include image completion, interpolation, demosaicing and Poisson denoising.

5.1 Image Completion and Interpolation

In the problems of image completion and interpolation, a subset Ω of the original image pixels are sampled in the input vector \mathbf{b} , while the remaining ones are unknown and must be determined. In both problems, the sampling can be expressed in the matrix form $\mathbf{A}\mathbf{x} = \mathbf{b}$ with a sampling matrix \mathbf{A} such that each row contains only zeros, except one element equal to 1 at the index of a sampled pixel.

When solving the problem with ADMM, all the unknown pixels are expected to have the highest estimation error at the initialization stage. However, over iterations, more reliable estimates should be obtained especially for pixels close to a sampled pixel. Hence, the denoiser within our preconditioned PnP-ADMM should be adjusted locally depending on the proximity to a sampled pixel. Following this intuition, we define a preconditioning matrix \mathbf{P} that varies over iterations according to the following formula for each iteration k :

$$[\mathbf{P}^k]_{i,i} = \frac{\max(\mathbf{m}^k) + \epsilon}{[\mathbf{m}^k]_i + \epsilon} \quad (25)$$

where ϵ is a scalar parameter and \mathbf{m}^k is defined recursively by blurring \mathbf{m}^{k-1} , and where \mathbf{m}^0 is the mask indicating the known pixels:

$$\mathbf{m}^k = \mathbf{m}^{k-1} * g(\sigma_f), \text{ with } [\mathbf{m}^0]_i = \begin{cases} 1 & \text{if } i \in \Omega \\ 0 & \text{otherwise} \end{cases} \quad \forall i \quad (26)$$

Here, $*$ is the 2D convolution operation and g is a 2D Gaussian filter of parameter σ_f . Note that it is equivalent to have $\mathbf{m}^k = \mathbf{m}^0 * g(\sigma_f \sqrt{k})$. We can thus determine σ_f so that the preconditioning at the last iteration N does not depend on N , but only on a fixed parameter σ_f^{last} at the last iteration by taking $\sigma_f = \sigma_f^{last} / \sqrt{N}$.

Using this definition, the preconditioning values are always equal to 1 for the known pixels (i.e. $[\mathbf{P}^k]_{i,i} = 1 \quad \forall i \in \Omega$), and higher than 1 for the other pixels, which results in a stronger denoising for the unknown pixels in Eq. (22). The parameter ϵ prevents too high preconditioning values that would make the denoising impractical. The highest preconditioning value is $p_{max} = (1 + \epsilon)/\epsilon$. Additionally, thanks to the Gaussian blur, the preconditioning values of unknown pixels will depend on their distance with sampled pixels. Therefore, the least reliable pixels (i.e. that are far from sampled pixels) will be denoised more.

Note that we only use non-zero preconditioning values (and thus non-zero standard deviation in the denoiser) so

that the matrix \mathbf{P} remains invertible. However, ideally, no denoising should be performed at the sampled pixel positions to prevent unnecessary loss of information. Therefore, we force the denoiser's output to be equal to the input for these pixels.

5.2 Demosaicing

Digital cameras typically capture colors using a sensor with a color filter array (CFA). Thanks to the mosaic formed by the CFA, neighbor pixels on the sensor record a different color information. Knowing the CFA pattern, the full color information can then be retrieved by an inverse problem called demosaicing.

Red, green and blue filters are generally used so that each pixel directly records one of the RGB components. Demosaicing can then be seen as an interpolation problem where a pixel R, G or B value is either known or unknown. Therefore, we use the same preconditioning strategy described in Section 5.1. However, in a realistic scenario, sensor data may also contain noise which is preferably removed jointly with the demosaicing step [31]. In the case of noisy data, our method applies similarly, but we do not force the denoiser to keep the sampled pixels unchanged as explained in Section 5.1.

Furthermore, unlike the problems of completion and interpolation, the demosaicing uses different masks indicating the sampled pixels for the R, G and B components. The preconditioning matrix thus takes different values for each component of the same pixel, which requires using a denoiser parameterized with a RGB noise level map. Our *DRUNet-var-RBG* network trained for arbitrary noise level patterns is thus suitable for this application.

However, in order to illustrate the tradeoff between the genericity of the denoiser and the optimal performance, we also trained a more specialized denoiser for demosaicing that we call *DRUNet-dem*. For training this network, instead of generating noise level maps with the generic random process in Eq. (24), we generate the patterns from the CFA mask, that is also used in the preconditioned ADMM. These patterns are defined by Eqs. (25) and (26), where a mask \mathbf{m}^0 is defined for each color component by the CFA pattern. Experimental results using either the generic network *DRUNet-var-RBG* or the specialized one *DRUNet-dem*, are given in Section 6.4.

5.3 Poisson Denoising

In many practical scenarios, the assumption that images are corrupted by additive white Gaussian noise is inaccurate. Camera noise is typically better modelled by a Poisson process. In order to formulate a MAP optimization for the problem of denoising an image corrupted with Poisson noise, the least squares data term in Eq. (3) must be modified. Here, we re-derive our preconditioned ADMM algorithm for Poisson denoising.

5.3.1 Preconditioned ADMM for Poisson denoising

Applying Poisson noise to an original image $\mathbf{x} \in \mathbb{R}^n$ gives a noisy image $\mathbf{b} \in \mathbb{N}^n$ such that the likelihood distribution is:

$$p(\mathbf{b}|\mathbf{x}) = \prod_{i=1}^n \frac{\mathbf{x}_i^{\mathbf{b}_i} e^{-\mathbf{x}_i}}{\mathbf{b}_i!}. \quad (27)$$

As seen in Eqs. (1)-(3), the data term is the negative log-likelihood which can be derived as

$$-\ln(p(\mathbf{b}|\mathbf{x})) = -\mathbf{b}^\top \ln(\mathbf{x}) + \mathbf{1}^\top \mathbf{x} + c, \quad (28)$$

where c is a constant term that can be ignored for the minimization. The preconditioned Poisson denoising problem is then formulated by substituting \mathbf{x} with $\mathbf{P}\mathbf{x}$ and by adding the regularization term as in Eq. (14):

$$\tilde{\mathbf{x}} = \arg \min_{\mathbf{x}} -\mathbf{b}^\top \ln(\mathbf{P}\mathbf{x}) + \mathbf{1}^\top \mathbf{P}\mathbf{x} + \mathcal{R}(\mathbf{P}\mathbf{x}), \quad (29)$$

The problem is solved using the ADMM, by splitting the data and regularization terms following the methodology in Section 2.2. Since the regularization term remains the same as in Section 3, the y-update in Eq. (22) still applies, and is performed using the same Gaussian denoiser $G_{\mathcal{R}}$. However, given the new data term, the x-update becomes:

$$\mathbf{x}^{k+1} = \arg \min_{\mathbf{x}} -\mathbf{b}^\top \ln(\mathbf{P}\mathbf{x}) + \mathbf{1}^\top \mathbf{P}\mathbf{x} + \frac{\rho^k}{2} \|\mathbf{x} - \mathbf{u}^k\|_2^2, \quad (30)$$

where $\mathbf{u}^k = \mathbf{y}^k - \mathbf{l}^k/\rho^k$.

Knowing that \mathbf{P} is a diagonal matrix, this problem can be solved independently for each pixel. The following closed form solution is derived in Appendix A (similar derivations are also found in [14] for the case $\mathbf{P} = \mathbf{I}$):

$$[\mathbf{x}^{k+1}]_i = \frac{\rho^k \mathbf{u}_i - \mathbf{P}_{i,i} + \sqrt{(\rho^k \mathbf{u}_i - \mathbf{P}_{i,i})^2 + 4\rho^k \mathbf{b}_i}}{2\rho^k}. \quad (31)$$

It is worth noting that unlike the Gaussian case, the noise level only depends on the data and is not controlled by a parameter σ . However, in photography the amount of Poisson noise can be reduced by increasing the exposure. This effect can be simulated by applying Poisson noise to a ground truth image rescaled in the range $[0, \lambda]$. A high peak value λ thus simulates a high exposure and less noise. However, for an input image \mathbf{b} with values in the range $[0, \lambda]$, the algorithm must be modified since the denoiser $G_{\mathcal{R}}$ in Eq. (22) assumes data in the range $[0, 1]$. The y-update should then be performed by replacing $G_{\mathcal{R}}$ with the denoiser $G_{\mathcal{R}}^\lambda$ adapted to the range $[0, \lambda]$ by rescaling the input and output as:

$$G_{\mathcal{R}}^\lambda(\mathbf{u}, \Sigma^{1/2}) = \lambda G_{\mathcal{R}}(\mathbf{u}/\lambda, \Sigma^{1/2}/\lambda) \quad (32)$$

Similarly, the final ADMM result should be divided by λ to recover an image in the range $[0, 1]$.

5.3.2 Choice of Preconditioning Matrix

It is well known that the variance and the expected value of a random variable with the Poisson distribution are both equal to the distribution's parameter, i.e. the noiseless image pixel. This means that for each pixel of the image, the noise standard deviation is equal to the square root of the pixel value in the ground truth image. Although this image is unknown, the noisy image provides a first estimate of the noise variance, which can be refined at the next iterations, as we obtain a better estimate of the noiseless image.

Therefore, in this problem, we use $\rho_k = 1$ for each iteration, so that the denoiser is only controlled using the preconditioning matrix which is initialized as $\mathbf{P}^0 = \text{diag}(\sqrt{\mathbf{b}})$.

Here, the square root is applied element-wise and the notation diag forms a diagonal matrix from the vector. Following the initialization strategy in Section 3.3, the variable \mathbf{x} is then initialized such that $\hat{\mathbf{x}}^0 = \mathbf{P}^0 \mathbf{x}^0 = \mathbf{b}$, and thus $\mathbf{x}^0 = [\mathbf{P}^0]^{-1} \mathbf{b} = \sqrt{\mathbf{b}}$.

Then, at each iteration, a more accurate estimate of the noiseless image is given by $\mathbf{P}^k \mathbf{y}^k$ obtained after the denoising step in Eq. (22). Note that $\mathbf{P}^k \mathbf{y}^k$ is directly the output of the denoiser (before applying $[\mathbf{P}^k]^{-1}$) and is thus obtained without further matrix multiplication. Hence, the preconditioning matrix can be updated to better estimate the noise standard deviation using $\mathbf{P}^{k+1} = \text{diag}(\sqrt{\mathbf{P}^k \mathbf{y}^k})$.

6 EXPERIMENTAL RESULTS

6.1 ADMM Parameters Setting

The main hyper-parameters of the ADMM are the number of iterations N , the initial penalty parameter ρ^0 , and its increase factor α in Eq. (12). The additional parameter σ in Eq. (14) depends on the problem and corresponds to the Gaussian noise standard deviation in the measurement vector \mathbf{b} . For noise-free applications (i.e. completion, interpolation, noise-free demosaicing), using the theoretical value $\sigma = 0$ would remove the regularization term. Therefore, in order to keep the benefit of the regularization in these applications, we use a small value $\sigma = 1/255$. Also note that this parameter does not appear in the case of Poisson denoising in Eq. (29) since the noise standard deviation only depends on the ground truth image.

6.1.1 Completion, Interpolation and Demosaicing

For the completion, interpolation and demosaicing (either with or without noise) the parameters setting is inspired from [11]: since the noise level of the denoiser is globally controlled at each iteration by $\sigma_{den}^k = \sigma/\sqrt{\rho^k}$, we choose ρ^0 such that σ_{den}^0 is sufficiently large to remove initialization noise or artifacts at the first iteration (regardless of the loss in details). For example, in the completion problem, we use zero-initialisation and a large value of $\sigma_{den}^0 = 1$. For the interpolation and demosaicing problems, more accurate initialisation can be performed, thus we use a smaller value of $\sigma_{den}^0 = 50/255$. The parameter α is then determined so that σ_{den}^k decreases down to $\sigma_{den}^N = \sigma$, in order to best preserve the details in the final image. Therefore, we have $\sigma_{den}^k = \sigma_{den}^N/\sqrt{\rho^k}$. Using this equation for $k = 0$ and $k = N$, we can compute the parameters values:

$$\rho^0 = \left(\frac{\sigma_{den}^N}{\sigma_{den}^0} \right)^2 \quad \text{and} \quad \alpha = \left(\frac{1}{\rho^0} \right)^{1/N} \quad (33)$$

Note that the noise level assumed by the denoiser is also controlled locally by the preconditioning matrix \mathbf{P} in Eq. (22). However, our preconditioning strategy for these problems is such that $\mathbf{P}_{i,i} = 1$ if i is the index of a sampled pixel (see Section 5.1). Hence at the last iteration, the image is denoised assuming the correct noise standard deviation $\sigma_{den}^N = \sigma$ for the sampled pixels. In our experiments, we can thus compare our preconditioned PnP-ADMM to the non-preconditioned version using the same parameterization of σ_{den}^0 and σ_{den}^N .

TABLE 3

Summary of the parameters setting. For the demosaicing with preconditioning, we use either $\sigma_f^{last} = 0$ with the *DRUNet-dem* network or $\sigma_f^{last} = 0.3$ with the *DRUNet-var-RBG* network.

		without preco.		with preco.			
		σ_{den}^0	σ_{den}^N	σ_{den}^0	σ_{den}^N	σ_f^{last}	p_{max}
completion		1	$\frac{1}{255}$	1	$\frac{1}{255}$	0	10
interpolation		$\frac{50}{255}$	$\frac{1}{255}$	$\frac{50}{255}$	$\frac{1}{255}$	0.4	10
demosaicing	noise-free	$\frac{50}{255}$	$\frac{1}{255}$	$\frac{50}{255}$	$\frac{1}{255}$	0 or 0.3	10
	noise σ	$\frac{50}{255}$	σ	$\frac{50}{255}$	σ	0 or 0.3	10

		without preco.		with preco.	
Poisson denoising		ρ^0	α	ρ^0	α
(peak λ)		$\frac{1}{\lambda}$	$4^{1/N}$	1	1

The preconditioning matrix definition in Eqs. (25) and (26) introduces the additional parameters σ_f , controlling the blurring of the mask at each iteration, and ϵ , controlling the maximum preconditioning value $p_{max} = (1 + \epsilon)/\epsilon$. In all the experiments, we use $p_{max} = 10$ (i.e. $\epsilon = 1/9$). The parameter σ_f is set according to the number of iterations N using $\sigma_f = \sigma_f^{last}/\sqrt{N}$ as explained in Section 5.1, where σ_f^{last} is the blurring parameter at the last iteration that we set empirically for each application (see details in the respective sections and parameters summary in Table 3).

6.1.2 Poisson Denoising

For the problem of Poisson denoising, the noise level at each pixel and each iteration is fully controlled by the preconditioning matrix as explained in section 5.3, which removes the need for the penalty parameter. Therefore we simply use $\rho^0 = 1$ and $\alpha = 1$.

Nevertheless in order to evaluate the advantage of the preconditioning, we also implement the non-preconditioned version by taking $\mathbf{P} = \mathbf{I}$ and using the *DRUNet-cst* network. In this case, the noise standard deviation assumed by the denoiser in Eq. (22) is $1/\sqrt{\rho^k}$ at each iteration k . Given an input image data in the range $[0, \lambda]$ and with Poisson noise, the highest possible noise standard deviation is $\sqrt{\lambda}$. Hence, we initialize ρ_0 so that $1/\sqrt{\rho^0} = \sqrt{\lambda}$. We found empirically that the best results are always obtained by decreasing the denoiser's standard deviation down to $1/\sqrt{\rho^N} = \sqrt{\lambda}/2$ at the last iteration N . Therefore in our experiments without preconditioning, we always use $\rho^0 = 1/\lambda$ and $\alpha = 4^{1/N}$.

A summary of the parameters setting can be found in Table 3. It does not include the number of iterations N which is studied more in detail for each application in the following sections.

6.2 Interpolation

The interpolation problem can be seen as a particular case of super-resolution, where the low resolution input image was generated without applying an antialiasing filter before sub-sampling the ground truth image pixels. As a result, the main difficulty of the interpolation problem is to remove aliasing effects. First, we show in Fig. 2 that our method better removes aliasing when we update the preconditioning matrix \mathbf{P} by blurring the mask of sampled pixels, as described in Section 5.1 (see Eqs. (25)-(26)). In this

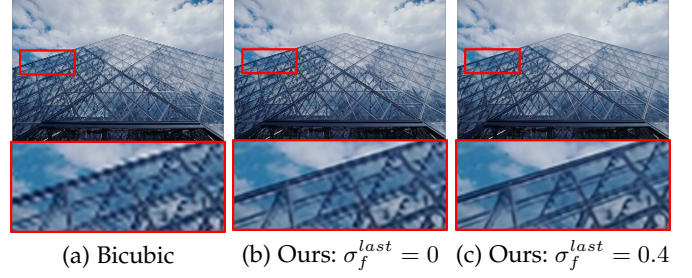


Fig. 2. Interpolation for x2 upsampling using either: (a) bicubic interpolation, (b) our method without \mathbf{P} update (i.e. $\sigma_f^{last} = 0$), (c) our method with $\sigma_f^{last} = 0.4$. For (b) and (c), we used $N = 6$ iterations.

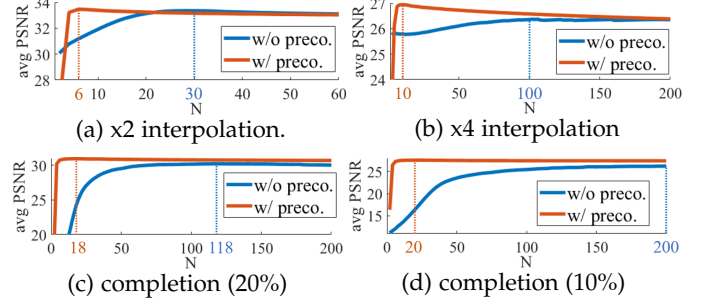


Fig. 3. Results of the PnP-ADMM either with or without preconditioning with respect to the number of iterations N for: (a) x2 interpolation, (b) x4 interpolation, (c) completion from 20% pixels, (d) completion from 10% pixels. The plots show the average PSNR over the Set5 Dataset.

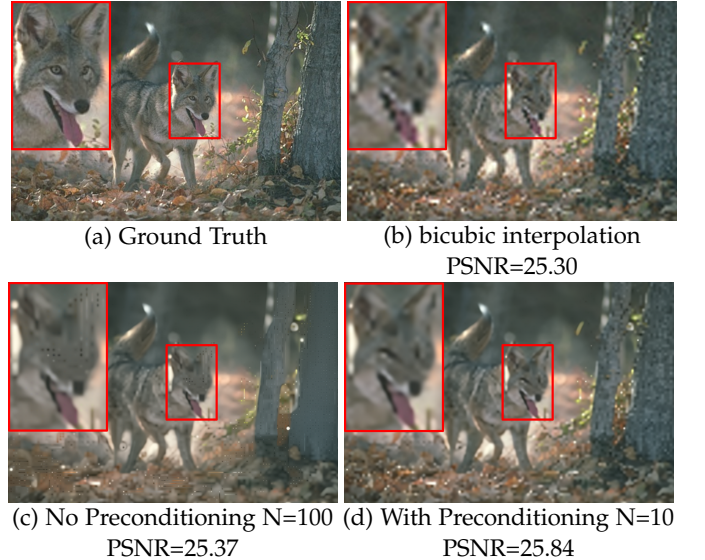


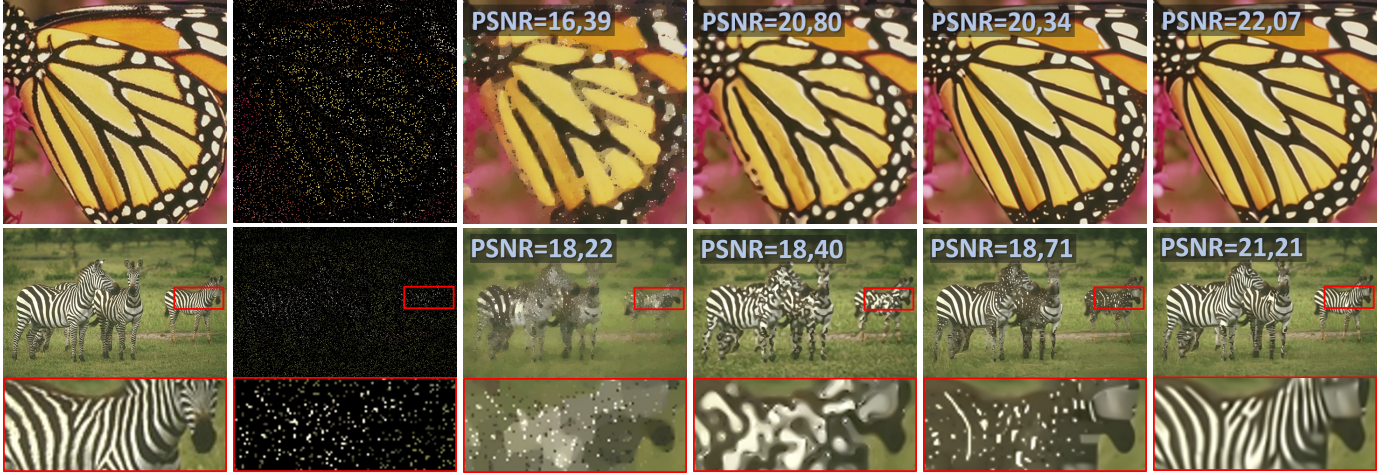
Fig. 4. Example of x4 interpolation results.

TABLE 4

Interpolation results for x2 and x4 factors (average PSNR over each dataset). In each case, the number of iterations N giving the best results is used for the PnP-ADMM.

		Set5	CBS68
x2	bicubic interpolation	31.63	27.85
	PnP-ADMM w/o preco. ($N = 30$)	33.34	28.85
	PnP-ADMM w/ preco. ($N = 6$)	33.47	29.00
x4	bicubic interpolation	25.66	23.20
	PnP-ADMM w/o preco. ($N = 100$)	26.36	23.65
	PnP-ADMM w/ preco. ($N = 10$)	26.94	23.90

example, the results of the bicubic interpolation in Fig. 2(a) displays strong aliasing artifacts. In Fig. 2(b), using our preconditioned PnP-ADMM with $\sigma_f^{last} = 0$ (i.e. no update of \mathbf{P} and no blurring of the mask) partially removes the aliasing. However, better results are obtained in Fig. 2(c),



(a) Ground Truth (b) Input (10% pixels) (c) Total Variation (d) Diffusion (EED) (e) PnP-ADMM (f) PnP-ADMM (preco.)

Fig. 5. Results of completion from 10% pixels for the butterfly image using (c) Total Variation regularization, (d) Edge-Enhancing Diffusion, (e) PnP-ADMM without preconditioning (using $N = 200$ iterations), (f) PnP-ADMM with our preconditioning (using $N = 20$ iterations).

TABLE 5

Completion results for 20% and 10% rates (average PSNR over each dataset). In each case, the number of iterations N giving the best results is used for the PnP-ADMM.

		Set5	CBS68
20%	Total Variation regularization	26.11	24.80
	Edge Enhancing Diffusion	28.61	25.55
	PnP-ADMM w/o preco. ($N = 118$)	30.20	26.75
	PnP-ADMM w/ preco. ($N = 18$)	30.94	27.43
10%	Total Variation regularization	22.86	22.66
	Edge Enhancing Diffusion	25.85	23.43
	PnP-ADMM w/o preco. ($N = 200$)	26.20	24.06
	PnP-ADMM w/ preco. ($N = 20$)	27.52	24.95

when updating \mathbf{P} using $\sigma_f^{last} = 0.4$.

Although many super-resolution algorithms exist and excellent performances have been obtained by deep learning techniques, these methods are generally trained assuming that the degradation already includes a given antialiasing filter. Therefore these methods are not suitable for the interpolation. Hence, we only compare our approach to the non-preconditioned PnP-ADMM, as well as the bicubic interpolation which we also use as an initialisation in the ADMM. The results of the PnP-ADMM either with or without preconditioning depends on the chosen number of iterations N . Fig. 3(a)-(b) shows that without preconditioning, the best results are obtained using respectively $N = 30$ and $N = 100$ iterations for x2 and x4 interpolation, while with our preconditioning, we only need to use $N = 6$ and $N = 10$ respectively. Using these values for the parameter N in each case, the results in Table 4 and Fig. 4 show that in addition to the faster convergence, our preconditioning improves the final results of the PnP-ADMM both in x2 and x4 interpolation.

6.3 Completion

For the completion problem, we evaluate the results using a rate of either 20% or 10% of known pixels. Unlike the interpolation problem, the known pixels do not form a regular grid and are instead selected randomly, which prevents the aliasing artifacts. For this reason, contrary to what we observed for the interpolation in Fig. 2, updating the preconditioning matrix \mathbf{P} with a blurred version of the

mask does not improve the completion results. Hence, we use $\sigma_f^{last} = 0$ in the following experiments.

Based on the results in Fig. 3(c)-(d), we use $N = 18$ and $N = 20$ iterations respectively for the 20% and 10% completion problems when using our preconditioning. Similarly to the interpolation problem, the convergence of the non-preconditioned PnP-ADMM is significantly slower: the best results for the Set5 dataset are obtained using respectively $N = 118$ and $N = 200$ for 20% and 10% completion.

For further comparisons, we also provide completion results obtained with the conventional total variation (TV) regularization, using the implementation in [34]. We also compare our results to a state of the art diffusion based method. A review of diffusion methods in [35] have shown that the best completion performance was obtained with a non-linear anisotropic diffusion scheme referred to as edge enhancing diffusion (EED). The EED was also exploited within compression schemes in more recent works (e.g. [36], [37]), where only a small percentage of pixels are encoded and the remaining ones must be recovered by the decoder. Therefore, we use the EED scheme for our comparisons.

The results in Fig. 5 show that, while the EED diffusion better recovers large image structures than solving the inverse problem with TV regularization, the PnP-ADMM scheme retrieves finer details and sharper edges thanks to the advanced regularization provided by the trained denoiser. However, without our preconditioning, some important structures of the image may be lost as shown in Fig. 5(e) (e.g. white spots on the butterfly, zebras' stripes). On the other hand, our preconditioned version recovers all the structures that could be inferred from the input data, hence resulting in a large gain in PSNR. Note that this observation is consistent with the interpolation results in Fig. 4, where important details of the image are lost in the non-preconditioned PnP-ADMM but are preserved with our preconditioning. The average PSNR results in Table 5 confirm the superiority of the proposed preconditioned PnP-ADMM over the other approaches for image completion from either 10% or 20% pixels.

TABLE 6

Demosaicing results (average PSNR over each dataset) for the Bayer pattern in the noise-free scenario ($\sigma = 0$). For our preconditioned PnP-ADMM method, we show the results using either the generic denoiser *DRUNet-var-RBG* or the specialized denoiser for demosaicing *DRUNet-dem*. The best and second best results are in red and blue respectively.

	Malvar et al. (Matlab) [32]	Gharbi et al. (DeepJoint) [31]	Henz et al. [21]	Kokkinos et al. (MMNet) [33]	PnP-ADMM		
					no preconditioning (<i>DRUNet-cst</i>)	with preconditioning	
						<i>DRUNet-var-RBG</i>	<i>DRUNet-dem</i>
Kodak	34.68	41.85	42.05	40.26	41.45	42.37	42.32
McMaster	34.46	39.13	39.50	36.84	39.24	39.27	39.31



Fig. 6. Demosaicing results in the noise-free scenario. The “PnP-ADMM (*DRUNet-cst*)” result corresponds to the non-preconditioned version, while our method is preconditioned and uses either the generic denoiser *DRUNet-var-RBG* or the specialized one *DRUNet-dem*.

TABLE 7

Demosaicing results (average PSNR over each dataset) for the Bayer pattern in the noisy scenario with noise of standard deviation $\sigma = 10/255$ or $\sigma = 20/255$. For our preconditioned PnP-ADMM method, we show the results using either the generic denoiser *DRUNet-var-RBG* or the specialized denoiser for demosaicing *DRUNet-dem*. The best and second best results are in red and blue respectively.

		Malvar et al. (Matlab) [32]	Gharbi et al. (DeepJoint) [31]	Kokkinos et al. (MMNet) [33]	PnP-ADMM		
					no preconditioning (<i>DRUNet-cst</i>)	with preconditioning	
						<i>DRUNet-var-RBG</i>	<i>DRUNet-dem</i>
$\sigma = 10/255$	Kodak	27.54	33.27	30.96	33.24	33.90	34.18
	McMaster	27.69	33.18	30.19	33.35	33.89	34.06
$\sigma = 20/255$	Kodak	22.38	30.04	23.81	29.92	30.84	31.18
	McMaster	22.75	30.18	23.65	30.09	31.09	31.36

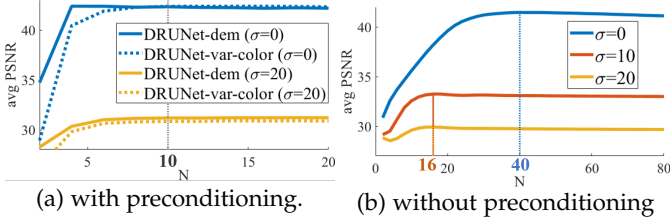


Fig. 7. Results of the PnP-ADMM for demosaicing with respect to the number of iterations N : (a) with preconditioning, (b) without preconditioning. The plots show the average PSNR over the Kodak Dataset.

6.4 Demosaicing

For the demosaicing, we evaluate the results both in the noise-free scenario and in the presence of Gaussian noise of standard deviation $\sigma = 10$ or $\sigma = 20$. In each case, the input images are filtered by the traditional Bayer CFA, i.e. only one of the red green or blue components is known at each pixel. The Bayer CFA forms a repetitive pattern of 2×2 blocks, each containing one red, one blue and two green pixels. For our preconditioned PnP-ADMM method, we evaluate the results using either the generic denoising network *DRUNet-var-RBG* or the specialized network *DRUNet-dem* that was trained for noise level patterns generated from the Bayer CFA, as explained in Section 5.2.

Similarly to the interpolation problem, the repetition of the 2×2 block pattern in the Bayer CFA may cause artifacts comparable to aliasing. When using the generic network *DRUNet-var-RBG* in our preconditioning scheme, these artifacts are better removed by updating the matrix \mathbf{P} with a blurred version of the mask. In this case, we use $\sigma_f^{last} = 0.3$ to control the blur at each iteration. However, when using

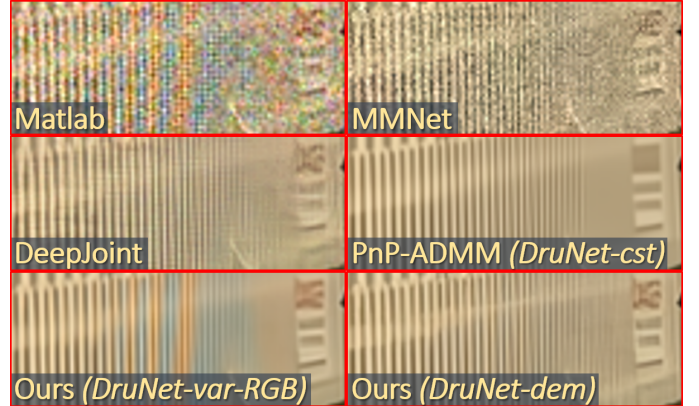


Fig. 8. Demosaicing results for the same detail as in Fig. 6, but in the noisy scenario with Gaussian noise of standard deviation $\sigma = 20$.

the denoiser *DRUNet-dem* specialized for such patterns, the aliasing artifacts from the Bayer CFA are naturally removed without the need for altering the preconditioning matrix. Thus we use $\sigma_f^{last} = 0$ in this case.

Note that for all the experiments, a fast initialisation is performed with Matlab’s demosaicing method [32] that uses bilinear interpolation and gradient-based corrections. Given this initialisation, we fix the number of iterations N of the PnP-ADMM according to the preliminary experiment in Fig. 7. Using our preconditioning (Fig. 7(a)), close to optimal results are obtained using $N = 10$, either with or without noise, and for both the *DRUNet-var-RBG* and *DRUNet-dem* networks. More iterations are needed for the non-preconditioned PnP-ADMM which reaches its best performance using $N = 40$ in the noise-free scenario ($\sigma = 0$), and using $N = 16$ in both noisy settings.

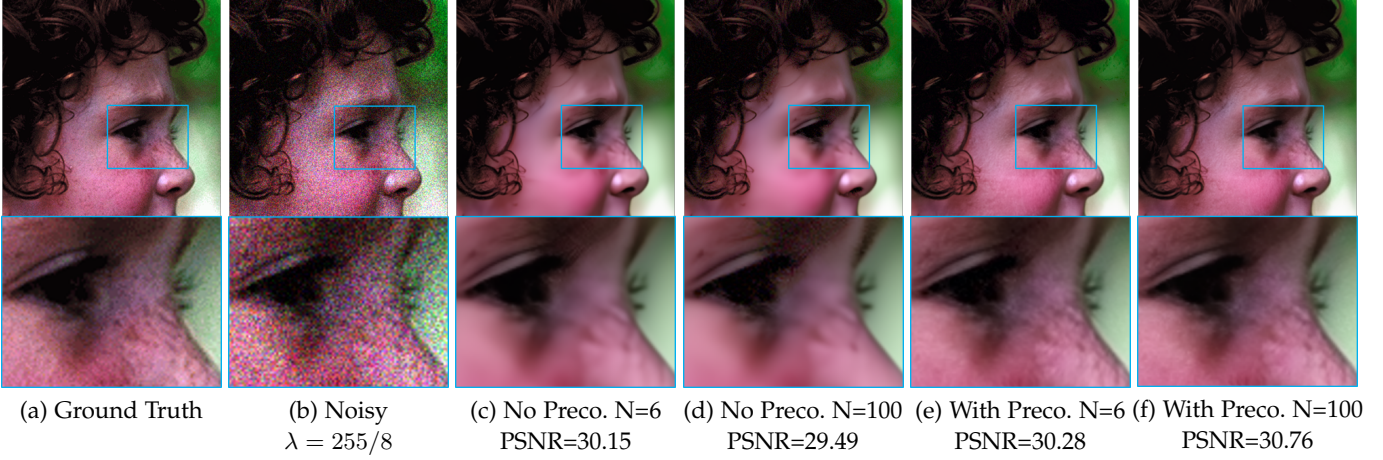
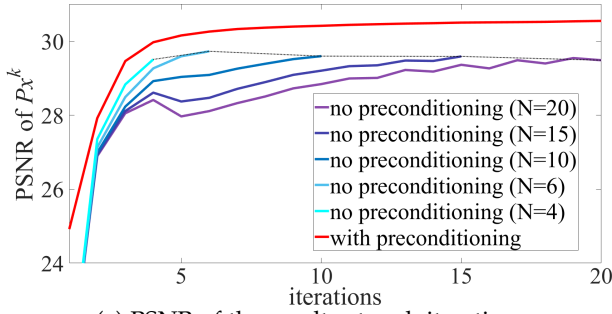
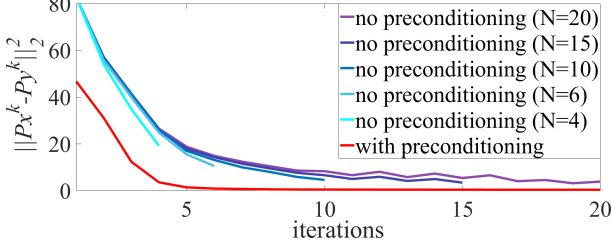


Fig. 9. Poisson Denoising given a peak value $\lambda = 255/8$. Results are shown for the PnP-ADMM either with preconditioning (e,f) or without (c,d) and using a number of iterations $N = 6$ (c,e) or $N = 100$ (d,f).



(a) PSNR of the results at each iteration.



(b) Mean Square Error between the two ADMM variables

Fig. 10. Convergence plots of the PnP-ADMM for Poisson denoising with $\lambda = 255/8$ and for the image “head” (from Fig. 9). For, the non-preconditioned case, several parameters depend on the fixed number of iterations N , thus, different curves are obtained with respect to N . With our preconditioning, the curves for different values of N are superimposed on the red curve because no parameter depend on N .

For the evaluation, we use the Kodak [38] and McMaster [39] test datasets that are widely used for demosaicing benchmarks as they contain natural images which include challenging features for this application. We compare both versions of our preconditioned PnP-ADMM with the non-preconditioned version (based on the *DRUNet-cst* denoiser), as well as other reference methods including Matlab’s demosaicing method by Malvar et al. [32] and the deep learning based methods of Gharbi et al. (DeepJoint) [31], Henz et al. [21] and Kokkinos et al. (MMNet) [33].

Table 6 and Fig. 6 present the results in the noise-free scenario. Our method performs among the best in average PSNR, indicating that it successfully recovers the fine details overall, and with greater accuracy than the non-preconditioned PnP-ADMM. However, in very challenging areas with high frequency patterns (see Fig. 6), our precon-

ditioning based on the generic network *DRUNet-var-RBG* leaves some zipper and color artifacts. Note that, on closer inspection, similar artifacts also remain in the Henz et al. and DeepJoint methods. However, using the specialized denoiser *DRUNet-dem* in our scheme completely solves this issue. This shows that it was not a limitation of the preconditioning or the PnP-ADMM scheme itself, but rather a limitation of the *DRUNet-var-RBG* denoiser that does not perform optimally when the noise level map is a structured pattern constructed from the Bayer CFA.

Similar conclusions can be drawn for the noisy scenario in Table 7 and Fig. 8. Our preconditioned PnP-ADMM best recovers the image details, and using the specialized *DRUNet-dem* network prevents the color issues that may appear in challenging areas when using the generic *DRUNet-var-RBG* network.

6.5 Poisson Denoising

For Poisson denoising, we compare the results of the PnP-ADMM either with or without our preconditioning. Note that the non-preconditioned version was previously described in [14] for this application. However, the BM3D denoiser was used in [14]. For a fair evaluation of our approach, we use instead the *DRUNet-cst* network which gives the best performance when no preconditioning is required. For our preconditioned scheme, we use the *DRUNet-var-RBG* denoiser since the Poisson noise variance depends not only on pixels’ positions, but also on the colour component.

First, we analyse the results for moderate to high levels of Poisson noise (using $\lambda \in \{255, 255/4, 255/8\}$). Table 8 and Fig. 10(a) show that the the proposed preconditioning improves the results compared to the original PnP-ADMM, regardless of the fixed number of iterations N . Furthermore, a better convergence is obtained with our approach as shown in Fig. 10. While the results of the non-preconditioned PnP-ADMM start degrading when using more than $N = 6$ iterations, our method successfully converges to the best result. The improved convergence is confirmed in Fig. 10(b) showing that the equality constraint between the two ADMM variables is better satisfied with our preconditioning. The visual results in Fig. 9 also show that our approach can recover finer details. Using a large

TABLE 8
Results of PnP-ADMM for Poisson denoising either with or without preconditioning (average PSNR over each dataset).

	Dataset	Without Preconditioning		With Preconditioning	
		$N = 6$	$N = 20$	$N = 6$	$N = 20$
$\lambda = 255$	Set5	36.54	36.33	36.82	36.81
	CBSD68	36.35	36.25	36.58	36.57
$\lambda = \frac{255}{4}$	Set5	33.34	33.13	33.65	33.69
	CBSD68	32.36	32.19	32.72	32.74
$\lambda = \frac{255}{8}$	Set5	31.73	31.41	31.94	32.06
	CBSD68	30.42	30.20	30.81	30.89

TABLE 9
Denoising Results for extreme Poisson noise with $\lambda = 1$ (average PSNR over each dataset).

	Anscombe initialization	PnP-ADMM (N=100) no preconditioning	PnP-ADMM (N=100) with preconditioning
Set5	17.50	22.04	23.82
CBSD68	17.30	20.50	23.05

number of iterations further improves the PSNR in this case, although the visual results are similar (see Fig. 9(e,f)). On the other hand, in the non-preconditioned PnP-ADMM, using a too large number of iterations visually degrades the results by removing details in the bright regions, while leaving more noise in the dark regions (see Fig. 9(c,d)).

Let us now evaluate our method for an extreme noise level, i.e. using $\lambda = 1$. In this situation, the noisy image is a very inaccurate estimate of the noise variance, and it can't be used to initialize the preconditioning matrix. Instead, we compute an initial estimation using variance stabilization and Gaussian denoising: the Anscombe transform [40] is first applied to obtain an image with approximately constant noise variance; our Gaussian denoiser is then used assuming a constant standard deviation; finally the inverse Anscombe transform is applied after denoising. The preconditioning matrix \mathbf{P} is thus computed as the square root of this initial estimate. We also disabled the update of \mathbf{P} which degraded our results in this case. The PSNR results using this initialization method are presented in Table 9, and visual comparisons are shown in Fig. 11. Here, the best results are obtained after convergence either with or without the preconditioning. Hence the results are given for a large number of iterations $N = 100$. Our approach significantly outperforms both the non-preconditioned PnP-ADMM and the Anscombe variance stabilization method used for the initialization.

7 CONCLUSION

In this paper, we have proposed a new preconditioning approach for Plug-and-Play optimization methods in the context of image restoration. Existing PnP schemes perform regularization of inverse problems using any conventional denoiser that assumes independent and identically distributed Gaussian noise. On the theoretical level, we have shown that our preconditioning removes this i.i.d. assumption by introducing in the algorithm a denoiser parameterized with a covariance matrix (directly related to the chosen preconditioning matrix) instead of a single standard deviation parameter. Greater control is thus granted for

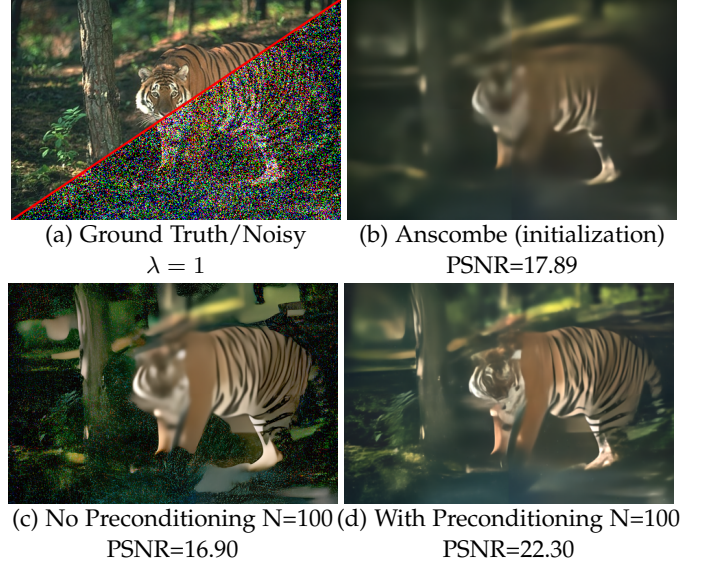


Fig. 11. Poisson Denoising for very high noise level (using peak value $\lambda = 1$).

better adapting the algorithm to each task by accounting for the specific error distribution of the current image estimate.

For our practical implementation, we have proposed a training procedure that generalizes a state-of-the-art CNN denoiser, enabling its parameterization with an arbitrary noise level map. While this restricts our study to diagonal covariance and preconditioning matrices, it is sufficient for many applications where the input image is degraded independently on each pixel. Such applications include image interpolation, completion, demosaicing and Poisson denoising. For each of these tasks, we have defined a suitable preconditioning scheme that significantly outperforms the non-preconditioned version both in convergence speed and image quality. The proposed method also preserves the genericity of the PnP approach since a denoiser is used to solve several inverse problems. Nevertheless, as we have shown with demosaicing, our method may be further specialized and improved for a given task by training the denoiser using the corresponding noise level maps.

A possible direction for future work would be to develop denoisers suitable for correlated noise (i.e. non-diagonal covariance matrix), hence extending the use of our preconditioned PnP scheme to a broader range of applications.

APPENDIX A

POISSON DATA TERM SUB-PROBLEM

Let us find the closed form solution of the minimization in Eq. (30) which we rewrite here without the iteration numbers for simplicity of notation:

$$\arg \min_x -\mathbf{b}^\top \ln(\mathbf{P}\mathbf{x}) + \mathbf{1}^\top \mathbf{P}\mathbf{x} + \frac{\rho}{2} \|\mathbf{x} - \mathbf{u}\|_2^2, \quad (34)$$

where \mathbf{P} is a diagonal matrix. Here, we additionally assume $\rho > 0$, $\mathbf{b}_i \geq 0$ and $\mathbf{P}_{i,i} > 0$, $\forall i$.

Since \mathbf{P} is diagonal, the problem is equivalently expressed for each element i independently. Using the scalar notations $b = \mathbf{b}_i$, $p = \mathbf{P}_{i,i}$ and $u = \mathbf{u}_i$, the problem is to find the value x that minimizes the function h defined as:

$$h(x) = -b \cdot \ln(x \cdot p) + x \cdot p + \frac{\rho}{2} (x - u)^2 \quad (35)$$

The function is defined for $x > 0$ and is convex. Therefore, finding $x > 0$ such that $\frac{dh}{dx}(x) = 0$ gives the global minimum. Differentiating h gives:

$$\frac{dh}{dx} = \frac{-b}{x} + p + \rho \cdot (x - u) = 0, \quad (36)$$

$$= \frac{\rho}{x} \cdot \left(x^2 + x \cdot \left(\frac{p - \rho u}{\rho} \right) - \frac{b}{\rho} \right). \quad (37)$$

Therefore, if we can find $x > 0$ such that x is a root of the second order polynomial $x^2 + x \cdot \left(\frac{p - \rho u}{\rho} \right) - \frac{b}{\rho}$, then x minimizes the function h . For $b_i > 0$, the solution always exists and is:

$$x = \frac{1}{2} \left(\frac{\rho u - p}{\rho} + \sqrt{\left(\frac{\rho u - p}{\rho} \right)^2 + \frac{4b}{\rho}} \right), \quad (38)$$

$$= \frac{1}{2\rho} \left(\rho u - p + \sqrt{(\rho u - p)^2 + 4\rho b} \right). \quad (39)$$

Note that if $b = 0$ and $u \leq p/\rho$, then the highest root is $x = 0$, which is not in the domain of h . However, when $b = 0$, we can simply redefine h using $h(x) = x \cdot p + \frac{\rho}{2}(x - u)^2$ which is defined for $x = 0$. So in practice, we can extend the domain of h to the value 0 (negative values are not valid pixel values and should remain excluded). With this definition, Eq. (39) gives the solution in every case.

Rewriting Eq.(39) using the matrix notations directly gives the solution to Eq. (34):

$$x_i = \frac{\rho u_i - P_{i,i} + \sqrt{(\rho u_i - P_{i,i})^2 + 4\rho b_i}}{2\rho}. \quad (40)$$

ACKNOWLEDGMENTS

This work was supported in part by the french ANR research agency in the context of the artificial intelligence project DeepCIM, and in part by the EU H2020 Research and Innovation Programme under grant agreement No 694122 (ERC advanced grant CLIM).

REFERENCES

- [1] S. Diamond, V. Sitzmann, F. Heide, and G. Wetzstein, "Unrolled optimization with deep priors," *arXiv preprint arXiv:1705.08041*, 2017.
- [2] M. Mardani, Q. Sun, S. Vasawanal, V. Pappayan, H. Monajemi, J. Pauly, and D. Donoho, "Neural proximal gradient descent for compressive imaging," *arXiv preprint arXiv:1806.03963*, 2018.
- [3] Y. Yang, J. Sun, H. Li, and Z. Xu, "Deep admm-net for compressive sensing mri," in *Proceedings of the 30th international conference on neural information processing systems*, 2016, pp. 10–18.
- [4] D. Gilton, G. Ongie, and R. Willett, "Neumann networks for linear inverse problems in imaging," *IEEE Transactions on Computational Imaging*, vol. 6, pp. 328–343, 2019.
- [5] Q. Ning, W. Dong, G. Shi, L. Li, and X. Li, "Accurate and lightweight image super-resolution with model-guided deep unfolding network," *IEEE Journal of Selected Topics in Signal Processing*, 2020.
- [6] S. Lefkimmiatis, "Universal denoising networks : A novel cnn architecture for image denoising," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2018, pp. 3204–3213.
- [7] W. Dong, P. Wang, W. Yin, G. Shi, F. Wu, and X. Lu, "Denoising prior driven deep neural network for image restoration," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 10, pp. 2305–2318, 2019.
- [8] S. Boyd, N. Parikh, and E. Chu, *Distributed optimization and statistical learning via the alternating direction method of multipliers*. Now Publishers Inc, 2011.
- [9] D. Geman and C. Yang, "Nonlinear image recovery with half-quadratic regularization," *IEEE transactions on Image Processing*, vol. 4, no. 7, pp. 932–946, 1995.
- [10] S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in *2013 IEEE Global Conference on Signal and Information Processing*, 2013, pp. 945–948.
- [11] K. Zhang, Y. Li, W. Zuo, L. Zhang, L. Van Gool, and R. Timofte, "Plug-and-play image restoration with deep denoiser prior," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [12] K. Zhang, W. Zuo, S. Gu, and L. Zhang, "Learning deep cnn denoiser prior for image restoration," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 3929–3938.
- [13] J. R. Chang, C.-L. Li, B. Póczos, B. Vijaya Kumar, and A. C. Sankaranarayanan, "One network to solve them all — solving linear inverse problems using deep projection models," in *2017 IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 5889–5898.
- [14] A. Rond, R. Giryes, and M. Elad, "Poisson inverse problems by the plug-and-play scheme," *Journal of Visual Communication and Image Representation*, vol. 41, pp. 96–108, 2016.
- [15] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-d transform-domain collaborative filtering," *IEEE Transactions on Image Processing*, vol. 16, no. 8, pp. 2080–2095, 2007.
- [16] M. Aharon, M. Elad, and A. Bruckstein, "K-svd: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Transactions on Signal Processing*, vol. 54, no. 11, pp. 4311–4322, 2006.
- [17] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [18] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention*, 2015, pp. 234–241.
- [19] T. Meinhardt, M. Moeller, C. Hazirbas, and D. Cremers, "Learning proximal operators: Using denoising networks for regularizing inverse imaging problems," in *International Conference on Computer Vision (ICCV)*, Oct. 2017, pp. 1799–1808.
- [20] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising," *IEEE Transactions on Image Processing*, vol. 26, no. 7, pp. 3142–3155, Jul. 2017.
- [21] B. Henz, E. S. Gastal, and M. M. Oliveira, "Deep joint design of color filter arrays and demosaicing," in *Computer Graphics Forum*, vol. 37, 2018, pp. 389–399.
- [22] P. Arbelaez, M. Maire, C. Fowlkes, and J. Malik, "Contour detection and hierarchical image segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 5, pp. 898–916, May 2011.
- [23] K. Ma, Z. Duanmu, Q. Wu, Z. Wang, H. Yong, H. Li, and L. Zhang, "Waterloo exploration database: New challenges for image quality assessment models," *IEEE Transactions on Image Processing*, vol. 26, no. 2, pp. 1004–1016, 2017.
- [24] E. Agustsson and R. Timofte, "Ntire 2017 challenge on single image super-resolution: Dataset and study," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, Jul. 2017.
- [25] B. Lim, S. Son, H. Kim, S. Nah, and K. M. Lee, "Enhanced deep residual networks for single image super-resolution," in *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2017, pp. 1132–1140.
- [26] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *ICLR (Poster)*, 2015.
- [27] K. Zhang, W. Zuo, and L. Zhang, "FFDNet: Toward a fast and flexible solution for CNN-based image denoising," *IEEE Transactions on Image Processing*, vol. 27, no. 9, pp. 4608–4622, 2018.
- [28] Y. Zhang, K. Li, K. Li, B. Zhong, and Y. Fu, "Residual non-local attention networks for image restoration," in *International Conference on Learning Representations (ICLR)*, May 2019.
- [29] M. Bevilacqua, A. Roumy, C. Guillemot, and M. line Alberi Morel, "Low-complexity single-image super-resolution based on nonnegative neighbor embedding," in *Proceedings of the British Machine Vision Conference*, 2012, pp. 135.1–135.10.
- [30] D. Martin, C. Fowlkes, D. Tal, and J. Malik, "A database of human segmented natural images and its application to evaluating

- segmentation algorithms and measuring ecological statistics," in *Proc. 8th Int'l Conf. Computer Vision*, vol. 2, Jul. 2001, pp. 416–423.
- [31] M. Gharbi, G. Chaurasia, S. Paris, and F. Durand, "Deep joint demosaicking and denoising," *ACM Transactions on Graphics (TOG)*, vol. 35, no. 6, pp. 1–12, 2016.
 - [32] H. S. Malvar, L.-W. He, and R. Cutler, "High-quality linear interpolation for demosaicing of bayer-patterned color images," in *Proc. IEEE ICASSP*, 2004.
 - [33] F. Kokkinos and S. Lefkimmiatis, "Iterative joint image demosaicking and denoising using a residual denoising network," *IEEE Transactions on Image Processing*, vol. 28, no. 8, pp. 4177–4188, 2019.
 - [34] J. Dahl, P. C. Hansen, S. H. Jensen, and T. L. Jensen, "Algorithms and software for total variation image reconstruction via first-order methods," in *Numerical Algorithms*, vol. 53, no. 1, Jul. 2009, pp. 67–92.
 - [35] I. Galić, J. Weickert, M. Welk, A. Bruhn, A. Belyaev, and H.-P. Seidel, "Image compression with anisotropic diffusion," *Journal of Mathematical Imaging and Vision*, vol. 31, no. 2, pp. 255–269, Apr. 2008.
 - [36] S. Andris, P. Peter, and J. Weickert, "A proof-of-concept framework for pde-based video compression," in *Picture Coding Symposium (PCS)*, 2016, pp. 1–5.
 - [37] C. Schmaltz, P. Peter, M. Mainberger, F. Ebel, J. Weickert, and A. Bruhn, "Understanding, optimising, and extending data compression with anisotropic diffusion," *Int. J. Comput. Vision*, vol. 108, no. 3, p. 222–240, Jul. 2014.
 - [38] R. Franzen, "Kodak lossless true color image suites," 1999, source: <http://r0k.us/graphics/kodak>.
 - [39] L. Zhang, X. Wu, A. Buades, and X. Li, "Color demosaicking by local directional interpolation and nonlocal adaptive thresholding," *Journal of Electronic Imaging*, vol. 20, no. 2, pp. 1 – 17, 2011.
 - [40] F. Anscombe, "The transformation of poisson, binomial and negative-binomial data," in *Biometrika*, vol. 35, no. 3-4, Dec. 1948, pp. 246–254.