

The Little Engine That Could: Regularization by Denoising (RED)*

Yaniv Romano[†], Michael Elad[‡], and Peyman Milanfar[‡]

Abstract. Removal of noise from an image is an extensively studied problem in image processing. Indeed, the recent advent of sophisticated and highly effective denoising algorithms has led some to believe that existing methods are touching the ceiling in terms of noise removal performance. Can we leverage this impressive achievement to treat other tasks in image processing? Recent work has answered this question positively, in the form of the Plug-and-Play Prior (P^3) method, showing that any inverse problem can be handled by sequentially applying image denoising steps. This relies heavily on the ADMM optimization technique in order to obtain this chained denoising interpretation. Is this the only way in which tasks in image processing can exploit the *image denoising engine*? In this paper we provide an alternative, more powerful, and more flexible framework for achieving the same goal. As opposed to the P^3 method, we offer Regularization by Denoising (RED): using the denoising engine in defining the regularization of the inverse problem. We propose an explicit image-adaptive Laplacian-based regularization functional, making the overall objective functional clearer and better defined. With a complete flexibility to choose the iterative optimization procedure for minimizing the above functional, RED is capable of incorporating any image denoising algorithm, can treat general inverse problems very effectively, and is guaranteed to converge to the globally optimal result. We test this approach and demonstrate state-of-the-art results in the image deblurring and super-resolution problems.

Key words. image denoising engine, Plug-and-Play Prior, Laplacian regularization, inverse problems

AMS subject classifications. 62H35, 68U10, 94A08

DOI. 10.1137/16M1102884

1. Introduction. We open this paper with a bold and possibly controversial statement: *To a large extent, removal of zero-mean white additive Gaussian noise from an image is a solved problem in image processing.*

Before justifying this statement, let us describe the basic building block that will be the star of this paper: *the image denoising engine*. From the humble linear Gaussian filter to the recently developed state-of-the-art methods using convolutional neural networks, there is no shortage of denoising approaches. In fact, these algorithms are so widely varied in their definition and underlying structure that a concise description will need to be made carefully. Our story begins with an image \mathbf{x} corrupted by zero-mean white additive Gaussian noise,

$$(1) \quad \mathbf{y} = \mathbf{x} + \mathbf{e}.$$

*Received by the editors November 9, 2016; accepted for publication (in revised form) August 29, 2017; published electronically October 19, 2017.

<http://www.siam.org/journals/siims/10-4/M110288.html>

Funding: The first author's research was supported by European Research Council grant 320640.

[†]The Electrical Engineering Department, The Technion - Israel Institute of Technology, Haifa, 32000, Israel (yaniv.romano@gmail.com).

[‡]Google Research, Mountain View, CA 94043 (elad@cs.technion.ac.il, peyman.milanfar@gmail.com).

In our notation, we consider an image as a vector of length n (after lexicographic ordering). In the above description, the noise vector is normally distributed: $\mathbf{e} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$. In the most general terms, the image denoising engine is a function $f : [0, 255]^n \rightarrow [0, 255]^n$ that maps an image \mathbf{y} to another image of the same size $\hat{\mathbf{x}} = f(\mathbf{y})$, with the hope of getting as close as possible to the original image \mathbf{x} . Ideally, such functions operate on the input image \mathbf{y} to remove the deleterious effect of the noise while maintaining edges and textures beneath.

The claim made above about the denoising problem being solved is based on the availability of algorithms proposed in the past decade that can treat this task extremely effectively and stably, getting very impressive results, which also tend to be quite concentrated (see, for example, the work reported in [3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 42]). Indeed, these documented results have led researchers to the educated guess that these methods are getting very close to the optimally possible denoising performance [21, 22, 23]. This aligns well with the unspoken observation in our community in recent years that investing more work to improve image denoising algorithms seems to lead to diminishing returns.

While the above may suggest that work on denoising algorithms is turning to a dead-end avenue, a new opportunity emerges from this trend: seeking ways to leverage the vast progress made on the image denoising front in order to treat other tasks in image processing, bringing their solutions to new heights. One natural path towards addressing this goal is to take an existing and well-performing denoising algorithm and generalize it to handle a new problem. This has been the logical path that has led to contributions such as [24, 25, 26, 27, 28, 29] and many others. These papers, and others like them, offer an exhaustive manual adaptation of existing denoising algorithms, carefully retailored to handle specific alternative problems. This line of work, while often successful, is quite limited, as it offers no flexibility and no general scheme for diverting image denoising engines to treat new image processing tasks.

Could one offer a more systematic way to exploit the abundance of high-performing image denoising algorithms to treat a much broader family of problems? The recent work by Venkatakrishnan, Bouman, and Wohlberg provides a positive and tantalizing answer to this question, in the form of the Plug-and-Play Prior (P^3) method [30, 31, 32, 33]. This technique builds on the use of an implicit prior for regularizing general inverse problems. When handling the obtained optimization task via the ADMM optimization scheme [57, 58, 59], the overall problem decomposes into a sequence of image denoising tasks, coupled with simpler L_2 -regularized inverse problems that are much easier to handle.

While the P^3 scheme may sound like the perfect answer to our prayers, reality is somewhat more complicated. First, this method is not always accompanied by a clear definition of the objective function, since the regularization being effectively used is only implicit, implied by the denoising algorithm. Indeed, it is not clear at all that there is an underlying objective function behind the P^3 scheme if arbitrary denoising engines are used [31]. Second, parameter tuning of the ADMM scheme is a delicate matter, and especially so under a nonprovable convergence regime, as is the case when using sophisticated denoising algorithms. Third, being intimately coupled with the ADMM, the P^3 scheme does not offer easy and flexible ways of replacing the iterative procedure. Because of these reasons, the P^3 scheme is neither a turn-key tool, nor is it free from emotional involvement. Nevertheless, the P^3 method has drawn much attention (see, e.g., [31, 32, 33, 34, 35, 36, 37, 38]), and rightfully so, as it offers a clear path towards harnessing a given image denoising engine for treating more general inverse

problems, just as described above.

Is there a more general alternative to the P^3 method that could be simpler and more stable? This paper puts forward such a framework, offering a systematic use of such denoising engines for regularization of inverse problems. We term the proposed method Regularization by Denoising (RED), relying on a general structured smoothness penalty term harnessed to regularize any desired inverse problem. More specifically, the regularization term we propose in this work is the following:

$$(2) \quad \rho(\mathbf{x}) = \frac{1}{2} \mathbf{x}^T [\mathbf{x} - f(\mathbf{x})],$$

in which the denoising engine itself is applied on the candidate image \mathbf{x} , and the penalty induced is proportional to the inner product between this image and its denoising residual. This defined smoothness regularization is effectively using an image-adaptive Laplacian, which in turn draws its definition from the arbitrary image denoising engine of choice, $f(\cdot)$. Surprisingly, under mild assumptions on $f(\cdot)$, it is shown that the gradient of the regularization term is manageable, given as the denoising residual, $\mathbf{x} - f(\mathbf{x})$. Therefore, armed with this regularization expression, we show that any inverse problem can be handled while calling the denoising engine iteratively.

RED, the newly proposed framework, is much more flexible in the choice of the optimization method to use, not being tightly coupled to one specific technique, as in the case of the P^3 scheme (relying on ADMM). Another key difference w.r.t. the P^3 method is that our adaptive Laplacian-based regularization functional is explicit, making the overall Bayesian objective function clearer and better defined. RED is capable of incorporating any image denoising algorithm, and can treat general inverse problems very effectively, while resulting in an overall algorithm with very simple structure.

An important advantage of RED over the P^3 scheme is the flexibility with which one can choose the denoising engine $f(\cdot)$ to plug in the regularization term. While most of the discussion in this paper keeps focusing on white Gaussian noise (WGN) removal, RED can actually deploy almost any denoising engine. Indeed, we define a set of two mild conditions that $f(\cdot)$ should satisfy and show that many known denoising methods obey these properties. As an example, in our experiments we show how the median filter can become an effective regularizer. Last but not least, we show that the defined regularization term is a convex function, implying that in most cases, in which the log-likelihood term is convex too, the proposed algorithms are guaranteed to converge to a global optimum solution. We demonstrate this scheme, showing state-of-the-art results in image deblurring and single image super-resolution.

This paper is organized as follows: In the next section we present the background material for this work, discussing the general form of inverse problems as optimization tasks, and presenting the P^3 scheme. Section 3 focuses on the image denoising engine, defining it and its properties clearly, so as to enable its use in the proposed Laplacian paradigm. Section 4 serves the main part of this work—introducing RED: a new way to use an image denoising engine to handle general structured inverse problems. In section 5 we analyze the proposed scheme, discussing convexity, an alternative formulation, and a qualitative comparison to the P^3 scheme. Results on the image deblurring and single-image super-resolution problems are brought in section 6, demonstrating the strength of the proposed scheme. We conclude the

paper in section 7 with a summary of the open questions that we identify for future work.

2. Preliminaries. In this section we provide background material that serves as the foundation to this work. We start by presenting the breed of optimization tasks we will work on throughout the paper for handling the inverse problems of interest. We then introduce the P^3 method and discuss its merits and weaknesses.

2.1. Inverse problems as optimization tasks. Bayesian estimation of an unknown image \mathbf{x} given its measured version \mathbf{y} uses the posterior conditional probability, $P(\mathbf{x}|\mathbf{y})$, in order to infer \mathbf{x} . The most popular estimator in this regime is the maximum a posteriori probability (MAP), which chooses the mode (\mathbf{x} for which the maximum probability is obtained) of the posterior. Using Bayes's rule, this implies that the estimation task is turned into an optimization problem of the form

$$(3) \quad \hat{\mathbf{x}}_{MAP} = \underset{\mathbf{x}}{\operatorname{Argmax}} P(\mathbf{x}|\mathbf{y}) = \underset{\mathbf{x}}{\operatorname{Argmax}} \frac{P(\mathbf{y}|\mathbf{x})P(\mathbf{x})}{P(\mathbf{y})} = \underset{\mathbf{x}}{\operatorname{Argmax}} P(\mathbf{y}|\mathbf{x})P(\mathbf{x}) \\ = \underset{\mathbf{x}}{\operatorname{Argmin}} \{-\log\{P(\mathbf{y}|\mathbf{x})\} - \log\{P(\mathbf{x})\}\}.$$

In the above derivations we exploited the fact that $P(\mathbf{y})$ is not a function of \mathbf{x} , and thus it can be omitted. We also used the fact that the $-\log$ function is monotonic decreasing, turning the maximization into a minimization problem.

The term $-\log\{P(\mathbf{y}|\mathbf{x})\}$ is known as the log-likelihood term, and it encapsulates the probabilistic relationship between the desired image \mathbf{x} and the measurements \mathbf{y} , under the assumption that the desired image is known. We shall rewrite this term as

$$(4) \quad \ell(\mathbf{y}, \mathbf{x}) = -\log\{P(\mathbf{y}|\mathbf{x})\}.$$

As a classic example for the log-likelihood that will accompany us throughout this paper, the expression $\ell(\mathbf{y}, \mathbf{x}) = \frac{1}{2\sigma^2} \|\mathbf{H}\mathbf{x} - \mathbf{y}\|_2^2$ refers to the case of $\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{e}$, where \mathbf{H} is a linear degradation operator and \mathbf{e} is WGN contamination of variance σ^2 . Naturally, if the noise distribution changes, we depart from the comfortable L_2 form.

The second term in (3), $-\log\{P(\mathbf{x})\}$, refers to the *prior*, bringing in the influence of the statistical nature of the unknown. This term is also referred to as the regularization, as it helps in better conditioning the overall optimization task in cases where the likelihood alone cannot lead to a unique or stable solution. We shall rewrite this term as

$$(5) \quad \lambda\rho(\mathbf{x}) = -\log\{P(\mathbf{x})\},$$

where λ is a scalar that encapsulates the confidence in this term.

What is $\rho(\mathbf{x})$ and how is it chosen? This is the holy grail of image processing, with a progressive advancement over the years of better modeling the image statistics and leveraging this for handling various tasks in image processing. Indeed, one could claim that almost everything done in our field surrounds this quest for choosing a proper prior, from the early smoothness prior $\rho(\mathbf{x}) = \lambda\mathbf{x}^T\mathbf{L}\mathbf{x}$ using the classic Laplacian [39], through total variation [40] and wavelet sparsity [41], all the way to recent proposals based on patch-based GMM [42, 43] and sparse-representation modeling [44]. Interestingly, the work we report here builds on the

surprising comeback of the Laplacian regularization in a much more sophisticated form, as reported in [47, 48, 49, 50, 51, 52, 53, 54, 55, 56].

Armed with a clear definition of the relation between the measurements and the unknown, and with a trusted prior, the MAP estimation boils down to the optimization problem of the form

$$(6) \quad \hat{\mathbf{x}}_{MAP} = \underset{\mathbf{x}}{\text{Argmin}} \ell(\mathbf{y}, \mathbf{x}) + \lambda\rho(\mathbf{x}).$$

This defines a wide family of inverse problems that we aim to address in this work, which includes tasks such as denoising, deblurring, super-resolution, demosaicing, tomographic reconstruction, optical-flow estimation, segmentation, and many other problems. The randomness in these problems is typically due to noise contamination of the measurements, and this could be Gaussian, Laplacian, Gamma-distributed, Poisson, and other noise models.

2.2. The Plug-and-Play Prior (P^3) approach. For completeness of this exposition, we briefly review the P^3 approach. Aiming to solve the problem posed in (6), the ADMM technique [57, 58, 59] suggests handling this by variable splitting, leading to the equivalent problem

$$(7) \quad \{\hat{\mathbf{x}}_{MAP}, \hat{\mathbf{v}}\} = \underset{\mathbf{x}, \mathbf{v}}{\text{Argmin}} \ell(\mathbf{y}, \mathbf{x}) + \lambda\rho(\mathbf{v}) \quad \text{s.t.} \quad \mathbf{x} = \mathbf{v}.$$

The constraint is turned into a penalty term, relying on the augmented Lagrangian method (in its scaled dual form [57]), leading to

$$(8) \quad \{\hat{\mathbf{x}}_{MAP}, \hat{\mathbf{v}}\} = \underset{\mathbf{x}, \mathbf{v}}{\text{Argmin}} \ell(\mathbf{y}, \mathbf{x}) + \lambda\rho(\mathbf{v}) + \frac{\beta}{2}\|\mathbf{x} - \mathbf{v} + \mathbf{u}\|_2^2,$$

where \mathbf{u} serves as the Lagrange multiplier vector for the set of constraints. ADMM addresses the resulting problem by updating \mathbf{x} , \mathbf{v} , and \mathbf{u} sequentially in a block-coordinate-descent fashion, leading to the following series of subproblems:

1. **Update of \mathbf{x} :** When considering \mathbf{v} (and \mathbf{u}) as fixed, the term $\rho(\mathbf{v})$ is omitted, and our task becomes

$$(9) \quad \hat{\mathbf{x}} = \underset{\mathbf{x}}{\text{Argmin}} \ell(\mathbf{y}, \mathbf{x}) + \frac{\beta}{2}\|\mathbf{x} - \mathbf{v} + \mathbf{u}\|_2^2,$$

which is a far simpler inverse problem, where the regularization is an L_2 proximity one, which is easy to solve in most cases.

2. **Update of \mathbf{v} :** In this stage we freeze \mathbf{x} (and \mathbf{u}), and thus the log-likelihood term drops, leading to

$$(10) \quad \hat{\mathbf{v}} = \underset{\mathbf{v}}{\text{Argmin}} \lambda\rho(\mathbf{v}) + \frac{\beta}{2}\|\mathbf{x} - \mathbf{v} + \mathbf{u}\|_2^2.$$

This stage is nothing but a denoising of the image $\mathbf{x} + \mathbf{u}$, assumed to be contaminated by a white additive Gaussian noise of power $\sigma^2 = 1/\beta$. This is easily verified by returning to (6) and plugging the log-likelihood term $\|\mathbf{v} - \mathbf{x} - \mathbf{u}\|_2^2/2\sigma^2$ referring to this case. Indeed, this is the prime observation in [30], as the authors suggest replacing

the direct solution of (10) by activating an image denoising engine of choice. This way, we do not need to define explicitly the regularization $\rho(\cdot)$ to be used, as it is implied by the engine chosen.

3. **Update of \mathbf{u} :** We complete the algorithm description by considering the update of the Lagrange multiplier vector \mathbf{u} , which is done by $\hat{\mathbf{u}} = \mathbf{u} + \mathbf{x} - \mathbf{v}$.

Although the above algorithm has a clear mathematical formulation and only two parameters, denoted by β and λ , it turns out that tuning these is not a trivial task. The source of complexity emerges from the fact that the input noise level to the denoiser is equal to $\sqrt{\lambda/\beta}$. The confidence in the prior is determined by λ , and the penalty on the distance between \mathbf{x} and \mathbf{v} is affected by β . Empirically, setting a fixed value of β does not seize the potential of this algorithm; following previous work (see, e.g. [32, 37]), a common practical strategy to achieve a high-quality estimation is to increase the value of β as a function of the iterations: starting from a relatively small value, i.e., allowing an aggressive regularization, then proceeding to a more conservative one that limits the smoothing effect, up to a point where β should be large enough to ensure convergence [32] and to avoid an undesired oversmoothed outcome. As one can imagine, it is cumbersome to choose the rate in which β should be increased, especially because the corrupted image $\mathbf{x} + \mathbf{u}$ is a function of the Lagrange multiplier, which varies through the iterations as well.

In terms of convergence, the P^3 scheme has been shown to be well-behaved under some conditions on the denoising algorithm used. While the work reported in [31] requires the denoiser to be a symmetric smoothing and nonexpansive filter, the later work in [32] relaxes this condition to much simpler boundedness of the denoising effect. However, both of these prove at best a convergence to a steady-state outcome, which is very far from the desired claim of getting to the global minimizer of the overall objective function. The work reported in [33] offers clear conditions for a global convergence of P^3 , requiring the denoiser to be nonexpansive, and emerging as the minimizer of a convex functional. A recently released paper extends the above by using a specific GMM-based denoiser, showing that these two conditions are met, thus guaranteeing global convergence of their ADMM scheme [38].

Indeed, in that respect, a delicate matter with the P^3 approach is the fact that given a choice of a denoising engine, it does not necessarily refer to a specific choice of a prior $\rho(\cdot)$, as not every such engine could have a MAP-oriented interpretation. This implies a fundamental difficulty in the P^3 scheme, as in this case we will be activating a denoising algorithm while departing from the original setting we have defined, and having no underlying cost function to serve. Indeed, the work reported in [31] addresses this very matter in a narrower setting, by studying the identity of the effective prior obtained from a chosen denoising engine. The author chooses to limit the answer to symmetric smoothing filters, showing that even in this special case, the outcome is far from being trivial. As we are about to see in the next section, this shortcoming can be overcome by adopting a different regularization strategy.

3. The image denoising engine. Image denoising is a special case of the inverse problem posed in (6), referring to the case $\mathbf{y} = \mathbf{x} + \mathbf{e}$, where \mathbf{e} is WGN contamination of variance σ^2 . In this case, the MAP problem becomes

$$(11) \quad \hat{\mathbf{x}}_{\text{Denoise}} = \underset{\mathbf{x}}{\text{Argmin}} \frac{1}{2\sigma^2} \|\mathbf{y} - \mathbf{x}\|_2^2 + \lambda\rho(\mathbf{x}).$$

The *image denoising engine*, which is the focal point of this work, is any candidate solver to the above problem, under a specific choice of a prior. In fact, in this work we choose to widen the definition of the image denoising engine to be any function $f : [0, 255]^n \rightarrow [0, 255]^n$ that maps an image \mathbf{y} to another image $f(\mathbf{y})$ of the same size, and which aims to treat the denoising problem by the operation $\hat{\mathbf{x}} = f(\mathbf{y})$, be it MAP-based, MMSE-based, or any other approach.

Below, we accompany the definition of a denoiser with a few basic conditions on the function f . Just before doing so, we make the following broad observation: Among the various degradations that inverse problems come to remedy, removal of noise is fundamentally different. Consider the set of all reasonable “natural” images living on a manifold \mathcal{M} . If we blur any given image or downscale it, it is still likely to live in \mathcal{M} . However, if the image is contaminated by an additive noise, it pops out of the manifold along the normal to \mathcal{M} with high probability. Denoising is therefore a fundamental mechanism for an orthogonal “projection” of an image back onto \mathcal{M} .¹ This may explain why denoising is such a central operation, which has been so heavily studied. In the context of this work, in any given step of our iterations, this projection would allow us to project the temporary result back onto \mathcal{M} , so as to increase chances of getting a good-quality restored version of our image.

3.1. Conditions and properties of $f(\mathbf{x})$. We pose the following two necessary conditions on $f(\mathbf{x})$ that will facilitate our later derivations. Both of these conditions rely on the differentiability² of the denoiser $f(\mathbf{x})$.

- **Condition 1: (Local) homogeneity.** A denoiser applied to a positively scaled image $f(c\mathbf{x})$ should result in a scaled version of the original image. More specifically, for any scalar $c \geq 0$ we must have $f(c\mathbf{x}) = cf(\mathbf{x})$. In this work we shall relax this condition and demand its satisfaction for $|c - 1| \leq \epsilon$ for a very small ϵ .

A direct implication of the above property refers to the behavior of the directional derivative of the denoiser $f(\mathbf{x})$ along the direction \mathbf{x} . This derivative can be evaluated as

$$(12) \quad \nabla_{\mathbf{x}} f(\mathbf{x}) \mathbf{x} = \frac{f(\mathbf{x} + \epsilon \mathbf{x}) - f(\mathbf{x})}{\epsilon}$$

for a very small ϵ . Invoking the homogeneity condition this leads to

$$(13) \quad \begin{aligned} \nabla_{\mathbf{x}} f(\mathbf{x}) \mathbf{x} &= \frac{(1 + \epsilon)f(\mathbf{x}) - f(\mathbf{x})}{\epsilon} \\ &= f(\mathbf{x}). \end{aligned}$$

Thus, the filter $f(\mathbf{x})$ can be written as³

$$(14) \quad f(\mathbf{x}) = \nabla_{\mathbf{x}} f(\mathbf{x}) \mathbf{x}.$$

¹In the context of optimization, a smaller class of the general denoising algorithms we define is characterized as “proximal operators” [60]. These operators are in fact direct generalizations of orthogonal projections.

²A discussion on this requirement and possible ways to relax it appear in Appendix D.

³This result is sometimes known as Euler’s homogeneous function theorem [65].

- **Condition 2: Strong passivity.** The Jacobian $\nabla_{\mathbf{x}}f(\mathbf{x})$ of the denoising algorithm is stable, satisfying the condition

$$(15) \quad \eta(\nabla_{\mathbf{x}}f(\mathbf{x})) \leq 1,$$

where $\eta(\mathbf{A})$ is the spectral radius of the matrix \mathbf{A} . We interpret this condition as the restriction of the denoiser not to magnify the norm of an input image, since

$$(16) \quad \|f(\mathbf{x})\| = \|\nabla_{\mathbf{x}}f(\mathbf{x})\mathbf{x}\| \leq \eta(\nabla_{\mathbf{x}}f(\mathbf{x})) \cdot \|\mathbf{x}\| \leq \|\mathbf{x}\|.$$

Here we have relied on the relation $f(\mathbf{x}) = \nabla_{\mathbf{x}}f(\mathbf{x})\mathbf{x}$ that has been established in (14).

Note that we have chosen this specific condition over the natural weaker alternative, $\|f(\mathbf{x})\| \leq \|\mathbf{x}\|$, since the strong condition implies the weak one.

A reformulation of the denoising engine that will be found useful throughout this work is the one suggested in [51], where we assume that the algorithm is built of two phases—a first in which highly nonlinear decisions are made, and a second in which these decisions are used to adapt a linear filter to the raw noisy image in order to perform the actual noise removal. Algorithms such as the NLM, kernel-regression, K-SVD, and many others admit this structure, and for them we can write

$$(17) \quad \hat{\mathbf{x}}_{\text{Denoise}} = f(\mathbf{y}) = \mathbf{W}(\mathbf{y})\mathbf{y}.$$

The matrix \mathbf{W} is an $n \times n$ matrix, representing the (pseudo-) linear filter that multiplies the $n \times 1$ noisy image vector \mathbf{y} . This matrix is image dependent, as it draws its content from the pixels in \mathbf{y} . Nevertheless, this pseudo-linear format provides a convenient description of the denoising engine for our later derivations. We should emphasize that while this notation is true for only some of the denoising algorithms, the proposed framework we outline in this paper is general and applies to *any* denoising filter that satisfies the two conditions posed above. Indeed, the careful reader will observe that this pseudo-linear form is closely related to the directional derivative relation shown above: $f(\mathbf{y}) = \nabla_{\mathbf{y}}f(\mathbf{y})\mathbf{y}$. In this form, the right-hand side is now reminding us of the pseudo-linear form where the matrix $\mathbf{W}(\mathbf{y})$ is replaced by the Jacobian matrix $\nabla_{\mathbf{y}}f(\mathbf{y})$.

As a side note we mention that yet another natural requirement on \mathbf{W} (or $\nabla_{\mathbf{y}}f(\mathbf{y})$ in a wider perspective) is that it is row-stochastic, implying that (i) this matrix is (entrywise) nonnegative, and that (ii) the vector $\mathbf{1}$ is an eigenvector of \mathbf{W} . This would imply a constancy behavior—a denoiser $f(\mathbf{y})$ does not alter a constant image. More specifically, defining $\mathbf{1}$ as an n -dimensional column vector of all ones, for any scalar $c \geq 0$ we have $f(c\mathbf{1}) = c\mathbf{1}$. Algorithms such as the NLM [1] and its many variants all lead to such a row-stochastic Jacobian. We note that this property, while nice to have, is not required for the derivations in this paper.

An interesting consequence of the homogeneity property is the following stability of the pseudo-linear operator \mathbf{W} . Starting with a first-order Taylor expansion of $f(\mathbf{x} + \mathbf{h})$, and invoking the directional derivative relation $f(\mathbf{y}) = \nabla_{\mathbf{y}}f(\mathbf{y})\mathbf{y}$, we get

$$(18) \quad \begin{aligned} f(\mathbf{y} + \mathbf{h}) &\approx f(\mathbf{y}) + \nabla_{\mathbf{y}}f(\mathbf{y}) \cdot \mathbf{h} \\ &\approx \nabla f(\mathbf{y}) \cdot \mathbf{y} + \nabla_{\mathbf{y}}f(\mathbf{y}) \cdot \mathbf{h} \\ &\approx \nabla_{\mathbf{y}}f(\mathbf{y}) \cdot (\mathbf{y} + \mathbf{h}). \end{aligned}$$

This result implies that while $\mathbf{W}(\mathbf{y}) = \nabla_{\mathbf{y}}f(\mathbf{y})$ may indeed depend on \mathbf{y} , its sensitivity to its perturbation is negligible, rendering it as an essentially constant linear operator on the perturbed image $\mathbf{y} + \mathbf{h}$.

3.2. Denoisers obeying the above conditions. We cannot conclude this section without answering the key question: Which are the denoising engines to which we are constantly referring? While these could include any of the thousands of denoising algorithms published over the years, we obviously focus on the best performing ones, such as the nonlocal means (NLM) and its advanced variants [1, 2, 3], the K-SVD denoising method that relies on sparse representation modeling of image patches [4] and its nonlocal extension [7], the kernel-regression method that exploits local orientation [6], the well-known BM3D that combines sparsity and self-similarity of patches [5], the EPLL scheme that suggests patch modeling based on the GMM model [42], CSR and NCSR, which cluster the patches and sparsify them jointly [8, 9], the group-Wiener filtering applied on patches [10], the multilayer perceptron method trained to clean an image or the more recent CNN-based alternative called the trainable nonlinear reaction-diffusion (TNRD) algorithm [11, 20], more recent work that proposed low-rank modeling of patches and the use of the weighted nuclear norm [16], the nonlocal sparsity with GSM model [17], and the list goes on and on. Each and every one of these options (and many others) is a candidate engine that could be fit into our scheme.

A fair and necessary question is whether the denoisers we work with obey the two conditions we have posed above (homogeneity and passivity) and whether the preliminary requirement of differentiability is met. We choose to defer the discussion on the differentiability to Appendix D, due to its relevance to several spread parts of this paper and focus here on the homogeneity and passivity.

Starting with the homogeneity property, we give experimental evidence, accompanied by theoretical analysis, to substantiate the fulfillment of this property by a series of well-known denoisers. Figure 1 shows $f((1 + \epsilon)\mathbf{x})$ versus $(1 + \epsilon)f(\mathbf{x})$ as a scatter-plot, tested for K-SVD, BM3D, NLM, EPLL, and TNRD. In all these experiments, the image \mathbf{x} is set to be Peppers, $\epsilon = 0.01$ and $\sigma = 5$ (level of noise assumed within f). As can be seen, a tendency to an equality $f((1 + \epsilon)\mathbf{x}) \approx (1 + \epsilon)f(\mathbf{x})$ is obtained, suggesting that all these are indeed satisfying the homogeneity property. The deviation from exact equality in each of these tests has been evaluated as the standard deviation of the difference $f((1 + \epsilon)\mathbf{x}) - (1 + \epsilon)f(\mathbf{x})$, leading to $2.95e - 4$, $3.38e - 4$, $1.38e - 4$, $1.46e - 4$, $9.51e - 5$, respectively. A further discussion on the homogeneity property from a theoretical perspective is given in Appendix C.

Turning to the passivity condition, a conceptual difficulty is the need to explicitly obtain the Jacobian of the denoiser in question. Assuming that we overcame this problem somehow and got $\nabla_{\mathbf{x}}f(\mathbf{x})$, its spectral radius would be evaluated using the Power-Method [68] that applies iterations of the form

$$(19) \quad \mathbf{h}_{k+1} = \frac{\nabla_{\mathbf{x}}f(\mathbf{x}) \cdot \mathbf{h}_k}{\|\nabla_{\mathbf{x}}f(\mathbf{x}) \cdot \mathbf{h}_k\|_2}.$$

The spectral radius itself is easily obtained as

$$(20) \quad \eta(\nabla_{\mathbf{x}}f(\mathbf{x})) = \frac{\mathbf{h}_{k+1}^T \mathbf{h}_k}{\mathbf{h}_k^T \mathbf{h}_k}.$$

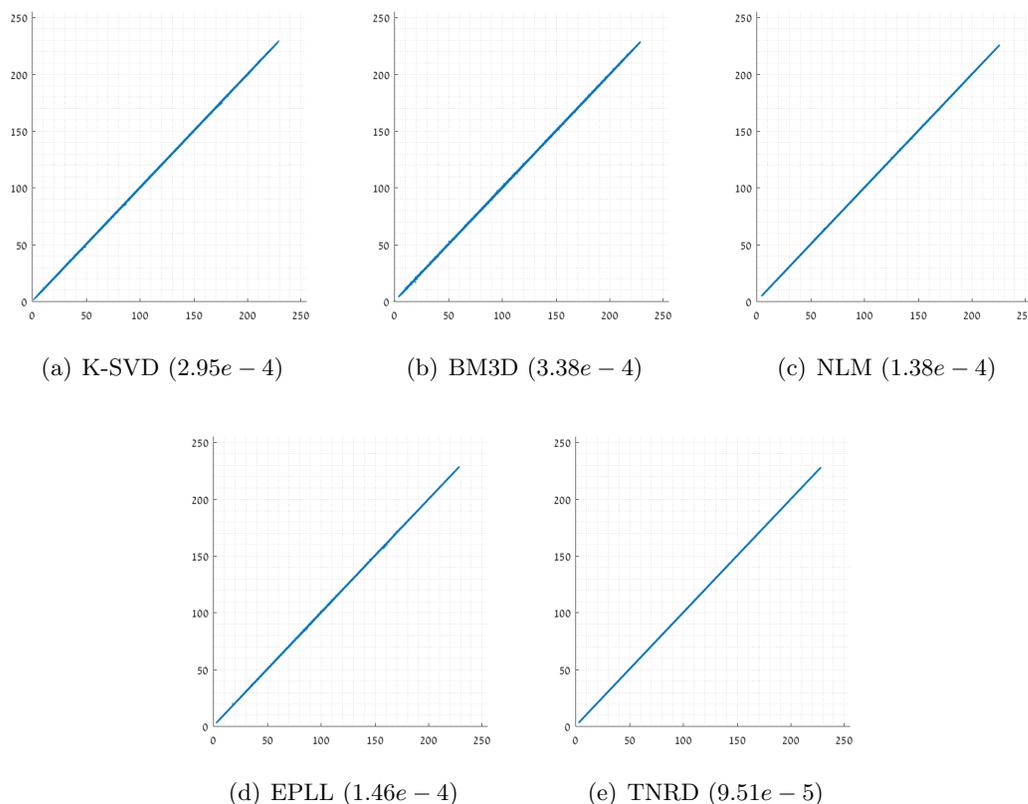


Figure 1. An empirical evaluation of the homogeneity property. These graphs show $f((1 + \epsilon)\mathbf{x})$ versus $(1 + \epsilon)f(\mathbf{x})$ as a scatter-plot for K-SVD, BM3D, NLM, EPLL, and TNRD. Equality implies satisfaction of the homogeneity, and the numbers in the brackets provide the STD of the difference. Note that these results were observed on various test images but are shown here for the image Peppers.

In order to bypass the need to explicitly obtain $\nabla_{\mathbf{x}}f(\mathbf{x})\mathbf{h}_k$, we rely on the first-order Taylor expansion again,

$$(21) \quad f(\mathbf{x} + \mathbf{h}) = f(\mathbf{x}) + \nabla_{\mathbf{x}}f(\mathbf{x}) \cdot \mathbf{h},$$

implying that $\nabla_{\mathbf{x}}f(\mathbf{x}) \cdot \mathbf{h} \approx f(\mathbf{x} + \mathbf{h}) - f(\mathbf{x})$, which holds true if $\|\mathbf{h}\|_2$ is small enough. Thus, our alternative Power-Method activates one denoising step per iteration,

$$(22) \quad \mathbf{h}_{k+1} \approx \frac{f(\mathbf{x} + \mathbf{h}_k) - f(\mathbf{x})}{\|f(\mathbf{x} + \mathbf{h}_k) - f(\mathbf{x})\|_2}.$$

The vector \mathbf{h}_k is normalized in each iteration, and thus $\|\mathbf{h}_k\|_2 = 1$. This vector is an image, and thus the gray values in it must be very small ($\mathbf{h}_k(j) \ll 1$), so as to lead to a sum of squares to be equal to 1. This agrees with the need for the perturbation $\mathbf{x} + \mathbf{h}_k$ to be small.

This algorithm has been applied to K-SVD, BM3D, NLM, EPLL, and TNRD (\mathbf{x} set to be the image Cameraman, $\sigma = 5$, number of iterations set to give an accuracy of $1e - 5$), all resulting with values smaller than or equal to 1, verifying the passivity of these filters.

4. Regularization by Denoising (RED).

4.1. The image-adaptive Laplacian. The new and alternative framework we propose relies on a form of an image-adaptive Laplacian which builds a powerful (empirical) prior that can be used to regularize a variety of inverse problems. As a place to start and motivate this definition, let's go back to the description of the denoiser given in (17), namely⁴ $\mathbf{W}(\mathbf{x})\mathbf{x}$. We may think of this pseudo-linear filter as one where a set of coefficients (depending on \mathbf{x}) is first computed in the matrix \mathbf{W} and then applied to the image \mathbf{x} . From this we can construct the Laplacian form,

$$(23) \quad \rho_L(\mathbf{x}) = \frac{1}{2} \mathbf{x}^T \mathbf{L}(\mathbf{x})\mathbf{x} = \frac{1}{2} \mathbf{x}^T (\mathbf{I} - \mathbf{W}(\mathbf{x}))\mathbf{x} = \frac{1}{2} \mathbf{x}^T [\mathbf{x} - \mathbf{W}(\mathbf{x})\mathbf{x}].$$

This definition by itself is not novel, as it is similar to ideas brought up in a series of recent contributions [47, 48, 49, 50, 51, 52, 53, 54, 55, 56]. This expression relies on using an image-adaptive Laplacian—one that draws its definition from the image itself.

Observing the obtained expression, we note that it can be interpreted as the *unnormalized cross-correlation* between the image \mathbf{x} and its corresponding residual $\mathbf{x} - \mathbf{W}(\mathbf{x})\mathbf{x}$. As a *prior* expression should give low values for likely images, in our case this would be achieved in one of two ways (or their combination):

- A small value is obtained for $\rho_L(\mathbf{x})$ if the residual is very small, implying that the image \mathbf{x} serves as a near fixed point of the denoising engine, $\mathbf{x} \approx \mathbf{W}(\mathbf{x})\mathbf{x}$.
- A small value is obtained for $\rho_L(\mathbf{x})$ if the cross-correlation of the residual to the image itself is small, a feature that implies that the residual behaves like white noise or, alternatively, if it does not contain elements from the image itself. Interestingly, this concept has been harnessed successfully by some denoising algorithms such as the Dantzig Selector [62] and by image denoising boosting techniques [51, 63, 64]. Indeed, enforcing orthogonality between the signal and its treated residual is the underlying force behind the normal equations in statistical estimation (e.g., least squares and Kalman filtering).

Given the above prior, we return to the general inverse problem posed in (6) and define our new objective,

$$(24) \quad \hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{Argmin}} \ell(\mathbf{y}, \mathbf{x}) + \frac{\lambda}{2} \mathbf{x}^T [\mathbf{x} - \mathbf{W}(\mathbf{x})\mathbf{x}].$$

The prior expression, while exhibiting a possibly complicated dependency on the unknown \mathbf{x} , is well-defined and clear. Nevertheless, an attempt to apply any gradient-based algorithm for solving the above minimization task encounters an immediate problem, due to the need to differentiate $\mathbf{W}(\mathbf{x})$ w.r.t. to \mathbf{x} . We overcome this problem by observing that $\mathbf{W}(\mathbf{x})\mathbf{x}$ is in fact the activation of the image denoising engine on \mathbf{x} , i.e., $f(\mathbf{x}) = \mathbf{W}(\mathbf{x})\mathbf{x}$. This observation inspires the following more general definition of the Laplacian regularizer, which is the prime message of this paper:

$$(25) \quad \rho_L(\mathbf{x}) = \frac{1}{2} \mathbf{x}^T [\mathbf{x} - f(\mathbf{x})].$$

⁴Note that we conveniently assume that the prior is applied to the clean image \mathbf{x} , a matter that will be clarified as we dive into our explanations.

This is the Regularization by Denoising (RED) paradigm that this work advocates. In this expression, the residual is defined more generally for any filter $f(\mathbf{x})$ even if it cannot be written in the familiar (pseudo-) linear form. Note that all the preceding intuition about the meaning of this prior remains intact; namely, the value is low if the cross-correlation between the image and its denoising residual is small or if the residual itself is small due to \mathbf{x} being a fixed point of f .

Surprisingly, while this expression is more general, it leads to a better-managed optimization problem, due to the careful properties we have outlined in section 3 on our denoising engines f . The overall energy functional to minimize is

$$(26) \quad E(\mathbf{x}) = \ell(\mathbf{y}, \mathbf{x}) + \frac{\lambda}{2} \mathbf{x}^T (\mathbf{x} - f(\mathbf{x})),$$

and the gradient of this expression is readily available by

$$(27) \quad \begin{aligned} \nabla_{\mathbf{x}} E(\mathbf{x}) &= \nabla_{\mathbf{x}} \ell(\mathbf{y}, \mathbf{x}) + \frac{\lambda}{2} \nabla_{\mathbf{x}} \{ \mathbf{x}^T (\mathbf{x} - f(\mathbf{x})) \} \\ &= \nabla_{\mathbf{x}} \ell(\mathbf{y}, \mathbf{x}) + \frac{\lambda}{2} \nabla_{\mathbf{x}} \mathbf{x}^T \mathbf{x} - \frac{\lambda}{2} \nabla_{\mathbf{x}} [\mathbf{x}^T f(\mathbf{x})] \\ &= \nabla_{\mathbf{x}} \ell(\mathbf{y}, \mathbf{x}) + \lambda \mathbf{x} - \frac{\lambda}{2} [f(\mathbf{x}) + \nabla_{\mathbf{x}} f(\mathbf{x}) \mathbf{x}]. \end{aligned}$$

Based on our prior assumption regarding the availability of a directional derivative for the denoising engine, the term $\nabla_{\mathbf{x}} f(\mathbf{x}) \mathbf{x}$ can be replaced⁵ by $\nabla_{\mathbf{x}} f(\mathbf{x}) \mathbf{x} = f(\mathbf{x})$, based on (14), implying that the gradient expression is further simplified to be

$$(28) \quad \nabla_{\mathbf{x}} E(\mathbf{x}) = \nabla_{\mathbf{x}} \ell(\mathbf{y}, \mathbf{x}) + \lambda \mathbf{x} - \lambda f(\mathbf{x}) = \nabla_{\mathbf{x}} \ell(\mathbf{y}, \mathbf{x}) + \lambda (\mathbf{x} - f(\mathbf{x})),$$

requiring only one activation of the denoising engine for the gradient evaluation. Interestingly, if we now bring back the pseudo-linear interpretation of the denoising engine, the gradient would be the residual, just as posed above, implying that

$$(29) \quad \nabla_{\mathbf{x}} E(\mathbf{x}) = \nabla_{\mathbf{x}} \ell(\mathbf{y}, \mathbf{x}) + \lambda (\mathbf{x} - \mathbf{W}(\mathbf{x}) \mathbf{x}).$$

Observe that this is a nontrivial derivation of the gradient of the original penalty function posed in (24).

4.2. Deploying the denoising engine for solving inverse problems. In the discussion above we have seen that the gradient of the energy functional to minimize (given in (26)) is easily computable, given in (28). We now turn to showing several options for using this in order to solve a general inverse problem. Common to all these methods is the fact that the eventual algorithm is iterative, in which each step is composed of applying the denoising engine (once or more), accompanied by other simpler calculations. In Figures 2, 3, and 4 we present pseudo-code for several such algorithms, all in the context of handling the case in which the likelihood function is given by $\ell(\mathbf{y}, \mathbf{x}) = \|\mathbf{H}\mathbf{x} - \mathbf{y}\|_2^2 / 2\sigma^2$.

⁵A better approximation can be applied in which we replace $\nabla_{\mathbf{x}} f(\mathbf{x}) \mathbf{x}$ by the difference $(f((1+\epsilon)\mathbf{x}) - f(\mathbf{x})) / \epsilon$, but this calls for two activations of the denoising engine per gradient evaluation.

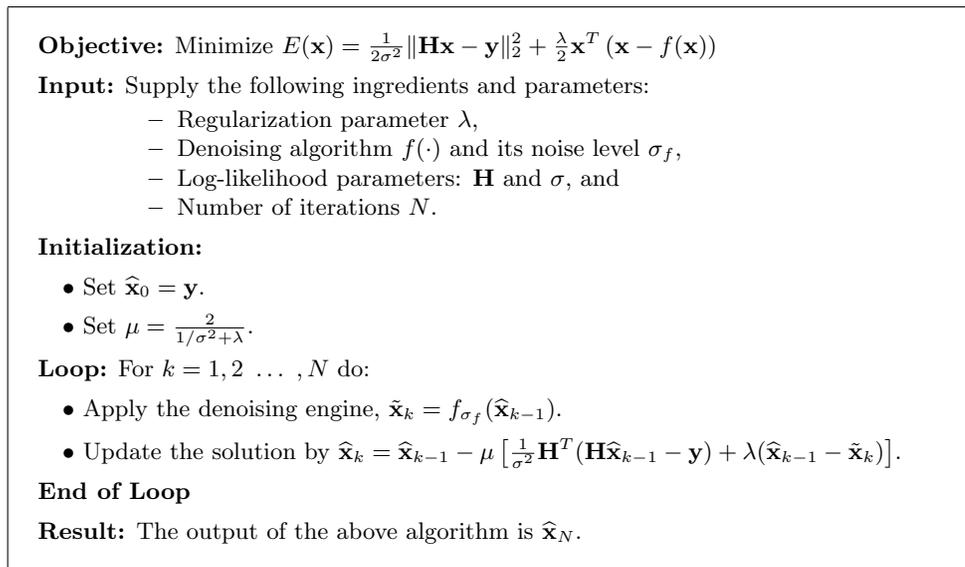


Figure 2. The proposed scheme (RED) via the SD method.

- **Gradient descent methods:** Given the gradient of the energy function $E(\mathbf{x})$, steepest descent (SD) is the simplest option that can be considered, and it amounts to the update formula

$$(30) \quad \begin{aligned} \hat{\mathbf{x}}_{k+1} &= \hat{\mathbf{x}}_k - \mu \nabla_{\mathbf{x}} E(\mathbf{x})|_{\hat{\mathbf{x}}_k} \\ &= \hat{\mathbf{x}}_k - \mu \left[\nabla_{\mathbf{x}} \ell(\mathbf{y}, \mathbf{x})|_{\hat{\mathbf{x}}_k} + \lambda(\hat{\mathbf{x}}_k - f(\hat{\mathbf{x}}_k)) \right]. \end{aligned}$$

Figure 2 describes this algorithm in more detail.

A line-search can be proposed in order to set μ dynamically per iteration, but this is necessarily more involved. For example, in the case of the Armijo rule, it requires a computation of the above gradient \mathbf{g}_k and then assessing the energy $E(\hat{\mathbf{x}}_k - \mu \mathbf{g}_k)$ for different values of μ in a retracting fashion, each of which calling for a computation of the denoising engine once.

One could envision using conjugate gradient (CG) to speed this method or, better yet, applying the sequential subspace optimization (SESOP) algorithm [69]. SESOP holds the current gradient and the last several update directions as the columns of a matrix \mathbf{V}_k (referring to the k th iteration), and seeks the best linear combination of these columns as an update direction to the current solution, namely, $\mathbf{x}_{k+1} = \mathbf{x}_k + \mathbf{V}_k \boldsymbol{\alpha}_k$. When restricted to have only one column, this reduces to a simple SD with line-search. When using two columns, it has the flavor (and strength) of CG, and when using more columns, this method can lead to much faster convergence in nonquadratic problems. The key points of SESOP are (i) the matrix \mathbf{V} is updated easily from one iteration to another by discarding the last direction, bringing in the last one, and adding the new gradient; and (ii) the unknown weight vector $\boldsymbol{\alpha}_k$ is low-dimensional, and thus updating it can be done using a Newton method. Naturally, one should evaluate the

first and second derivatives of the penalty function w.r.t. α_k , and these will leverage the relations established above. We shall not dive deeper into this option because it will not be included in our experiments.

One possible shortcoming of the gradient approach (in all its manifestations) is the fact that per activation of the denoising engine, the likelihood is updated rather mildly as a simple step toward the current log-likelihood gradient. This may imply that the overall algorithm will require many iterations to converge. The next two methods propose a way to overcome this limitation, by treating the log-likelihood term more “aggressively.”

- **ADMM:** Addressing the optimization task given in (26), we can imitate the path taken by the P^3 scheme, and apply variable splitting and ADMM. The steps would follow the description given in section 2 almost exactly, with one major difference—the prior in our case is explicit, and, therefore, the stage termed “update of \mathbf{v} ” would become

$$(31) \quad \hat{\mathbf{v}} = \underset{\mathbf{v}}{\text{Argmin}} \frac{\lambda}{2} \mathbf{v}^T (\mathbf{v} - f(\mathbf{v})) + \frac{\beta}{2} \|\mathbf{v} - \mathbf{x} - \mathbf{u}\|_2^2.$$

Rather than applying an arbitrary denoising engine to compute $\hat{\mathbf{v}}$ as a replacement to the actual minimization, we should target this minimization directly by some iterative scheme. For example, setting the gradient of the above expression to zero leads to the equation

$$(32) \quad \lambda(\mathbf{v} - f(\mathbf{v})) + \beta(\mathbf{v} - \mathbf{x} - \mathbf{u}) = \mathbf{0},$$

which can be solved iteratively using the fixed-point strategy, by

$$(33) \quad \begin{aligned} \lambda(\mathbf{v}_j - f(\mathbf{v}_{j-1})) + \beta(\mathbf{v}_j - \mathbf{x} - \mathbf{u}) &= \mathbf{0} \\ \rightarrow \mathbf{v}_j &= \frac{1}{\beta + \lambda} (f(\mathbf{v}_{j-1}) + \beta(\mathbf{x} + \mathbf{u})). \end{aligned}$$

This means that our approach in this case is computationally more expensive, as it will require several activations of the denoising engine. However, a common approach to speed up the convergence (in terms of runtime) of ADMM is called “early termination” [57], suggesting approximating the solution of the \mathbf{v} -update stage. We found this approach useful for our setting, especially because the application of a denoiser is computationally expensive. To this end, we may choose to apply only one iteration of the iterative process described in (33), which amounts to one operation of a denoising algorithm. Figure 3 describes this specific algorithm in more detail. If one changes all of **Part 2** (in Figure 3) with the computation $\hat{\mathbf{v}}_k = f_{1/\sqrt{\beta}}(\mathbf{z}^*)$, we obtain the P^3 scheme for the same choice of the denoising engine. While this difference is quite delicate, we should remind the reader that (i) this bridge between the two approaches is valid only when we deploy ADMM on our scheme, and (ii) as opposed to the P^3 method, our method is guaranteed to converge to the global optimum of the overall penalty function, as will be described hereafter.

We should point out that when using ADMM, the update of \mathbf{x} applies an aggressive inversion of the log-likelihood term, which is followed by the above optimization task.

Objective: Minimize $E(\mathbf{x}) = \frac{1}{2\sigma^2} \|\mathbf{H}\mathbf{x} - \mathbf{y}\|_2^2 + \frac{\lambda}{2} \mathbf{x}^T (\mathbf{x} - f(\mathbf{x}))$

Input: Supply the following ingredients and parameters:

- Regularization parameter λ ,
- Denoising algorithm $f(\cdot)$ and its noise level σ_f ,
- Log-likelihood parameters: \mathbf{H} and σ ,
- Number of outer and inner iterations, N , m_1 , and m_2 , and
- ADMM coefficient β .

Initialization: Set $\hat{\mathbf{x}}_0 = \mathbf{y}$, $\hat{\mathbf{v}}_0 = \mathbf{y}$, and $\hat{\mathbf{u}}_0 = \mathbf{0}$.

Outer Loop: For $k = 1, 2 \dots, N$ do:

Part 1: Solve $\hat{\mathbf{x}}_k = \underset{\mathbf{z}}{\text{Argmin}} \frac{1}{2\sigma^2} \|\mathbf{H}\mathbf{z} - \mathbf{y}\|_2^2 + \frac{\beta}{2} \|\mathbf{z} - \hat{\mathbf{v}}_{k-1} + \hat{\mathbf{u}}_{k-1}\|_2^2$ by

- Initialization: $\mathbf{z}_0 = \hat{\mathbf{x}}_{k-1}$, and define $\mathbf{z}^* = \hat{\mathbf{v}}_{k-1} - \hat{\mathbf{u}}_{k-1}$.
- Inner Iteration: For $j = 1, 2 \dots, m_1$ do:
 - Compute the gradient $\mathbf{e}_j = \frac{1}{\sigma^2} \mathbf{H}^T (\mathbf{H}\mathbf{z}_{j-1} - \mathbf{y}) + \beta(\mathbf{z}_{j-1} - \mathbf{z}^*)$.
 - Compute $\mathbf{r}_j = \frac{1}{\sigma^2} \mathbf{H}^T \mathbf{H} \mathbf{e}_j + \beta \mathbf{e}_j$.
 - Compute the step size $\mu = \mathbf{e}_j^T \mathbf{e}_j / \mathbf{e}_j^T \mathbf{r}_j$.
 - Update the solution by $\mathbf{z}_j = \mathbf{z}_{j-1} + \mu \mathbf{e}_j$.
 - Project the result to the interval $[0, 255]$.
- End of Inner Loop
- Set $\hat{\mathbf{x}}_k = \mathbf{z}_{m_1}$.

Part 2: Solve $\hat{\mathbf{v}}_k = \underset{\mathbf{z}}{\text{Argmin}} \lambda \mathbf{z}^T (\mathbf{z} - f_{\sigma_f}(\mathbf{z})) + \frac{\beta}{2} \|\mathbf{z} - \hat{\mathbf{x}}_k - \hat{\mathbf{u}}_{k-1}\|_2^2$ by

- Initialization: $\mathbf{z}_0 = \hat{\mathbf{v}}_{k-1}$, and define $\mathbf{z}^* = \hat{\mathbf{x}}_k + \hat{\mathbf{u}}_{k-1}$.
- Inner Iteration: For $j = 1, 2 \dots, m_2$ do:
 - Apply the denoising engine, $\tilde{\mathbf{z}}_j = f_{\sigma_f}(\hat{\mathbf{z}}_{j-1})$.
 - Compute the gradient $\mathbf{z}_j = \frac{1}{\beta + \lambda} (\lambda \tilde{\mathbf{z}}_j + \beta \mathbf{z}^*)$.
- End of Inner Loop
- Set $\hat{\mathbf{v}}_k = \mathbf{z}_{m_2}$.

Part 3: Update $\hat{\mathbf{u}}_k = \hat{\mathbf{u}}_{k-1} + \hat{\mathbf{x}}_k - \hat{\mathbf{u}}_{k-1}$.

End of Outer Loop

Result: The output of the above algorithm is $\hat{\mathbf{x}}_N$.

Figure 3. The proposed scheme (RED) via the ADMM method.

Thus, the shortcoming mentioned above regarding the lack of balance between the treatments given to the likelihood and the prior is mitigated.

- **Fixed-point strategy:** An appealing alternative to the above exists, obtained via the fixed-point strategy. As our aim is to find \mathbf{x} that nulls the gradient, this could be posed as an implicit equation to be solved directly,

$$(34) \quad \nabla_{\mathbf{x}} \ell(\mathbf{y}, \mathbf{x}) + \lambda(\mathbf{x} - f(\mathbf{x})) = 0.$$

Using the fixed-point strategy, this could be handled by the iterative formula

$$(35) \quad \nabla_{\mathbf{x}} \ell(\mathbf{y}, \mathbf{x}_{k+1}) + \lambda(\mathbf{x}_{k+1} - f(\mathbf{x}_k)) = 0.$$

As an example, in the case of linear degradation model and Gaussian white additive noise, this equation would be

$$(36) \quad \frac{1}{\sigma^2} \mathbf{H}^T (\mathbf{y} - \mathbf{H}\mathbf{x}_{k+1}) + \lambda(\mathbf{x}_{k+1} - f(\mathbf{x}_k)) = 0,$$

leading to the recursive update relation

$$(37) \quad \mathbf{x}_{k+1} = \left[\frac{1}{\sigma^2} \mathbf{H}^T \mathbf{H} + \lambda \mathbf{I} \right]^{-1} \left[\frac{1}{\sigma^2} \mathbf{H}^T \mathbf{y} + \lambda f(\mathbf{x}_k) \right].$$

This formula suggests one activation of the denoising per iteration, followed by what seems to be a plain Wiener filtering computation.⁶ The matrix inversion itself could be done in the Fourier domain for block-circulant \mathbf{H} , or iteratively using, for example, the Richardson algorithm: Defining

$$(38) \quad \mathbf{A} = \frac{1}{\sigma^2} \mathbf{H}^T \mathbf{H} + \lambda \mathbf{I} \quad \text{and} \quad \mathbf{b} = \frac{1}{\sigma^2} \mathbf{H}^T \mathbf{y} + \lambda f(\mathbf{x}_k),$$

our goal is to solve the linear system $\mathbf{A}\mathbf{x} = \mathbf{b}$. This is achieved by a variant of the SD method,⁷ $\mathbf{x}_{j+1} = \mathbf{x}_j - \mu(\mathbf{A}\mathbf{x}_j - \mathbf{b}) = \mathbf{x}_j - \mu \mathbf{e}_j$, where we have defined $\mathbf{e}_j = \mathbf{A}\mathbf{x}_j - \mathbf{b}$. By setting the step size to be $\mu_j = \mathbf{e}_j^T \mathbf{A} \mathbf{e}_j / \mathbf{e}_j^T \mathbf{A}^T \mathbf{A} \mathbf{e}_j$, we greedily optimize the potential of each iteration.

Convergence of the above algorithm is guaranteed, since

$$(39) \quad \left\| \left[\frac{1}{\sigma^2} \mathbf{H}^T \mathbf{H} + \lambda \mathbf{I} \right]^{-1} \lambda \nabla_{\mathbf{x}} f(\mathbf{x}_k) \right\| < 1.$$

This approach, similarly to ADMM, has the desired balance mentioned above between the likelihood and the regularization terms, matching the efforts dedicated to both. A pseudo-code describing this algorithm appears in Figure 4.

A basic question that has not been discussed so far is how to set the parameters of $f(\mathbf{x})$ in defining the regularization term. More specifically, assuming that the denoising engine depends on one parameter—the noise standard-deviation σ_f —the question is which value to use. While one could envision using varying values as the iterations progress and the outcome improves, the approach we take in this work is to set this parameter to be a small and fixed value. Our intuition for this choice is the desire to have a clear and fixed regularization term, which in turn implies a clear cost function to work with. Furthermore, the prior we propose should encapsulate in it our desire to get to a final image that is a stable point of such a weak denoising engine, $\mathbf{x} \approx f(\mathbf{x})$. Clearly, more work is required to better understand the influence of this parameter and its automatic setting.

⁶Note that (33) refers to the same problem posed here under the choice $\mathbf{H} = \mathbf{I}$ and $\beta = 1/\sigma^2$.

⁷All of this refers to a specific iteration k within which we apply inner iterations to solve the linear system and thus the use of the different index j .

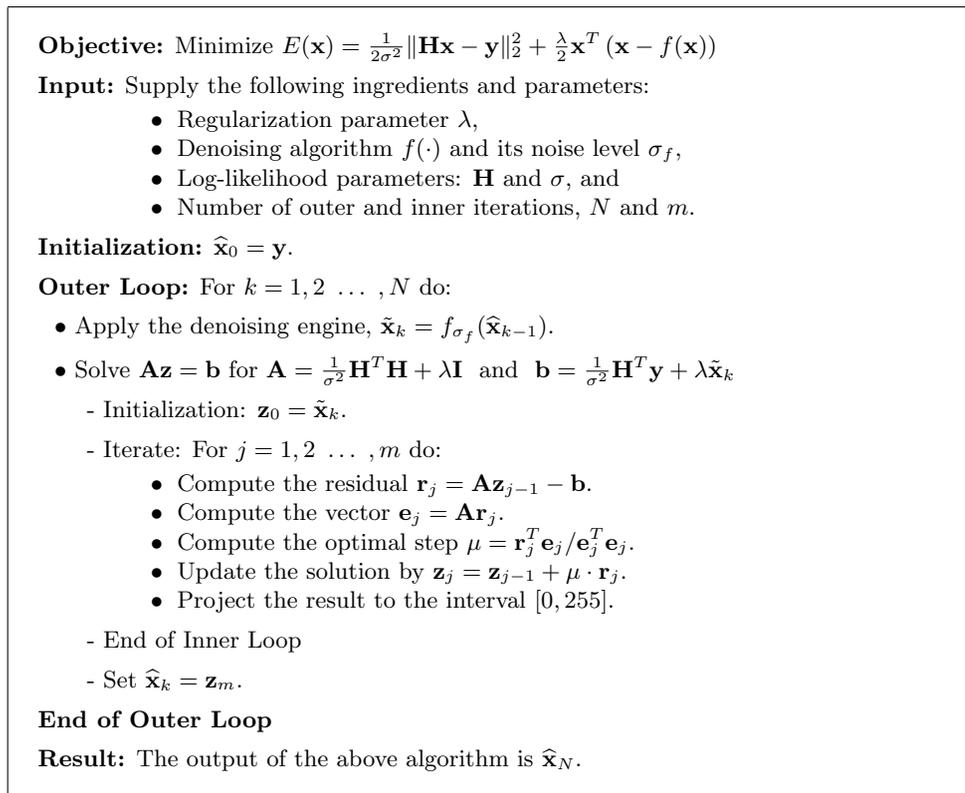


Figure 4. The proposed scheme (RED) via the fixed-point method.

5. Analysis.

5.1. Convexity. Is our proposed regularization function $\rho_L(\mathbf{x})$ convex? At first glance, this may seem like too much to expect. Nevertheless, it appears that for reasonably performing denoising engines obeying the conditions posed in section 3, this is exactly the case. For the function $\rho_L(\mathbf{x}) = \mathbf{x}^T (\mathbf{x} - f(\mathbf{x}))$ to be convex, we should demand that the second derivative be a positive semidefinite matrix [66]. We have already seen that the first derivative is simply $\mathbf{x} - f(\mathbf{x})$, which leads to the conclusion that the second derivative is given by $\mathbf{I} - \nabla_{\mathbf{x}} f(\mathbf{x})$.

As already mentioned earlier, in the context of some algorithms such as the NLM and the K-SVD, this is associated with the Laplacian $\mathbf{I} - \mathbf{W}(\mathbf{x})$, and it is positive semidefinite if \mathbf{W} has all its eigenvalues in the range⁸ $[0, 1]$. This is indeed the case for the NLM filter [1], the K-SVD-denoising algorithm [56], and many other denoising engines.

In the wider context of general image denoising engines, convexity is assured if the Jacobian $\nabla_{\mathbf{x}} f(\mathbf{x})$ of the denoising algorithm is stable, as indeed required in Condition 2 in section 3, $\eta(\nabla_{\mathbf{x}} f(\mathbf{x})) \leq 1$. In this case we have that $\rho_L(\cdot)$ is convex, and this implies that if the log-likelihood expression is convex as well, the proposed scheme is guaranteed to converge to the global optimum of our cost function in (6). In this respect the proposed algorithm is superior to the P^3 scheme in its most general form, which at best is known to get to a stable-

⁸We could in fact allow negative eigenvalues for \mathbf{W} , but this is unnatural in the context of denoising.

point [30, 32]. Furthermore, this result may seem similar to the one posed in [33, 38], as our two conditions for global convergence are homogeneity and passivity of f , while in these papers the requirements are passivity of f , along with a convex energy functional whom f minimizes. The second condition—having a convex origin to derive $f(\mathbf{x})$ —is in fact more restrictive than demanding homogeneity, as it is unclear which of the known denoisers meet this requirement.

5.2. An alternative prior. In section 4 we motivated the choice of the proposed prior by the desire to characterize the unknown image \mathbf{x} as one that is not affected by the denoising algorithm, namely, $\mathbf{x} \approx f(\mathbf{x})$. Rather than taking the route proposed, we could have suggested a prior of the form

$$(40) \quad \rho_Q(\mathbf{x}) = \|\mathbf{x} - f(\mathbf{x})\|_2^2.$$

This prior term makes sense intuitively, being based on the same desire to see the denoising residual being small. Indeed, this choice is somewhat related to the option we chose, since

$$(41) \quad \rho_Q(\mathbf{x}) = \|\mathbf{x} - f(\mathbf{x})\|_2^2 = \rho_L(\mathbf{x}) + f(\mathbf{x})^T(f(\mathbf{x}) - \mathbf{x}),$$

suggesting a symmetrization of our own expression.

In order to understand the deeper meaning of this alternative, we resort again to the pseudo-linear denoisers, for which this prior is nothing but

$$(42) \quad \rho_Q(\mathbf{x}) = \|\mathbf{x} - \mathbf{W}(\mathbf{x})\mathbf{x}\|_2^2 = \mathbf{x}^T(\mathbf{I} - \mathbf{W}(\mathbf{x}))^T(\mathbf{I} - \mathbf{W}(\mathbf{x}))\mathbf{x}.$$

This means that rather than regularizing with the Laplacian, we do so with its square. While this is a worthy possibility which has been considered in the literature under the term “fourth order regularization” [70], it is known to be more delicate. We leave this and other possibilities of formulating the regularization with the use of $f(\mathbf{x})$ for future work.

5.3. When is Plug-and-Play-Prior = RED? In section 4 we described the use of ADMM as one of the possible avenues for handling our proposed regularization. When handling the inverse problem posed in (26) with ADMM, we have shown that the only difference between this and the P^3 scheme resides in the update stage for \mathbf{v} . Here we aim to answer the following question: Assuming that the numerical algorithm used is indeed ADMM, under what conditions would the two methods (P^3 and ours) become equivalent? The answer to this question resides in the optimization task for updating \mathbf{v} , which is a denoising task. Thus, purifying this question, our goal is to find conditions on $f(\cdot)$ and λ such that the two treatments of this update stage coincide. Starting from our approach, we would seek the solution of

$$(43) \quad \hat{\mathbf{x}} = \underset{\mathbf{x}}{\text{Argmin}} \frac{\beta}{2} \|\mathbf{x} - \mathbf{y}\|_2^2 + \frac{\lambda}{2} \mathbf{x}^T(\mathbf{x} - f(\mathbf{x}))$$

or, putting it in terms of nulling the gradient of this energy, require

$$(44) \quad \beta(\mathbf{x} - \mathbf{y}) + \lambda(\mathbf{x} - f(\mathbf{x})) = \mathbf{0}.$$

The $\hat{\mathbf{x}}$ that is the solution of this equation is our updated image. On the other hand, the P^3 scheme would propose to simply compute⁹ $\hat{\mathbf{x}} = f(\mathbf{y})$ as a replacement to this minimization task. Therefore, for the two methods to coincide, we should demand that the gradient expression posed above be solved for the choice of the P^3 scheme, namely,

$$(45) \quad \beta(f(\mathbf{y}) - \mathbf{y}) + \lambda(f(\mathbf{y}) - f(f(\mathbf{y}))) = \mathbf{0},$$

or, posed slightly different,

$$(46) \quad f(\mathbf{y}) - f(f(\mathbf{y})) = \frac{\beta}{\lambda}(\mathbf{y} - f(\mathbf{y})).$$

This means that the denoising residual should remain the same (up to a constant) for the first activation of the denoising engine $\mathbf{y} - f(\mathbf{y})$ and the second one applied on the filtered image $f(\mathbf{y})$.

In order to get a better intuition towards this result, let's return to the pseudo-linear case, $f(\mathbf{y}) = \mathbf{W}\mathbf{y}$, with the assumption that \mathbf{W} is a fixed and diagonalizable matrix. Plugged into the above condition, this gives

$$(47) \quad \mathbf{W}\mathbf{y} - \mathbf{W}^2\mathbf{y} = \frac{\beta}{\lambda}(\mathbf{y} - \mathbf{W}\mathbf{y})$$

or, posed differently,

$$(48) \quad \left(\frac{\beta}{\lambda}\mathbf{I} - \mathbf{W}\right)(\mathbf{I} - \mathbf{W})\mathbf{y} = \mathbf{0}.$$

As the above equation should hold true for any image \mathbf{y} , we require that

$$(49) \quad \left(\frac{\beta}{\lambda}\mathbf{I} - \mathbf{W}\right)(\mathbf{I} - \mathbf{W}) = \mathbf{0}.$$

Without loss of generality, we can assume that \mathbf{W} is diagonal, after multiplying the above equation from the left and right by the diagonalizing matrix. With this simplification in mind, we now consider the eigenvalues of \mathbf{W} , and the above equation implies that exact equivalence between our scheme and the P^3 one is obtained only if our denoising engine has eigenvalues that are purely 1's or β/λ . Clearly, this is a very limiting case, which suggests that for all other cases, the two methods are likely to differ.

Interestingly, the above analysis is somewhat related to the one given in [31]. Both [31] and our treatment assume that the actual applied denoising engine is $f(\mathbf{y}) = \mathbf{W}\mathbf{y}$ within the ADMM scheme. While we ask for the conditions on \mathbf{W} to fit our regularization term $\mathbf{x}^T(\mathbf{x} - \mathbf{W}\mathbf{x})$, the author of [31] seeks the actual form of the prior to match this step, reaching the conclusion that the prior should be $\mathbf{x}^T(\mathbf{I} - \mathbf{W})\mathbf{W}^\dagger\mathbf{x}$. Bearing in mind that the conditions we get for the equivalence between the two methods are too restricting and rarely met, the result in [31] shows the actual gap between the two methods: While we regularize with the

⁹A delicate matter not considered here is that P^3 may apply $\frac{1}{c}f(c\mathbf{y})$ in order to tune to a specific noise level. We assume $c = 1$ for simplicity.

expression $\mathbf{x}^T(\mathbf{I} - \mathbf{W})\mathbf{x}$, an equivalence takes place only if P^3 modifies this to involve \mathbf{W}^\dagger , getting a far more complicated and less natural term.

Just before we conclude this section, we turn briefly to discuss the computational complexity of the proposed algorithm and its relation to the complexity of the P^3 scheme. Put very simply, RED and P^3 are roughly of the same computational cost. This is the case when RED is deployed via ADMM and assuming only one iteration in the update of \mathbf{v} , as shown above. Similarly, when using the fixed-point option, RED has the same cost as P^3 per iteration.

To conclude, we must state that this paper is about a more general framework rather than a comparison to P^3 . Indeed, one could consider this work as an attempt to provide more solid mathematical foundations for methods like P^3 . In addition, when comparing P^3 and RED, one can identify several major differences that are far more central than the complexity issue, such as (1) a lack of a clear objective function that P^3 serves, while our scheme has a very well-defined penalty; (2) the inability to claim much in terms of convergence of P^3 , while our penalty is shown to be convex; and (3) the complications of tuning the P^3 algorithm, which is very different from the experience we show with RED.

6. Results. In this section we compare the performance of the proposed framework to the P^3 approach, along with various other leading algorithms that are designed to tackle the image deblurring and super-resolution problems. To this end, we plug two substantially different denoising algorithms into the proposed scheme. The first is the (simple) *median filter*, which surprisingly turns out to act as a reasonable regularizer to our ill-posed inverse problems. This option is brought as a core demonstration of the idea that an arbitrary denoiser can be deployed in RED without difficulties. The second denoising engine we use is the state-of-the-art trainable nonlinear reaction diffusion (TNRD) [20] method. This algorithm trains a nonlinear reaction-diffusion model in a *supervised* manner. As such, in order to treat different restoration problems, one should retrain the underlying model for every specific task—something we aim to avoid. In the experiments below we build upon the published pretrained model by the authors of TNRD, tailored to denoise images that are contaminated by WGN with a fixed¹⁰ noise level, which is equal to 5. Leveraging this, we show how state-of-the-art deblurring and super-resolution results can be achieved simply by integrating the TNRD denoiser in RED. In all the experiments that follow, the parameters were manually set in order to enable each method to get its best possible results over the subset of images tested.

6.1. Image deblurring. In order to have a fair comparison to previous work, we follow the synthetic nonblind deblurring experiments conducted in the state-of-the-art work that introduced the nonlocally centralized sparse representation (NCSR) algorithm [9], which combines the self-similarity assumption [1] with the sparsity-inspired model [67]. More specifically, we degrade the test images, supplied by the authors of NCSR, by convolving them with two commonly used point spread functions (PSFs); the first is a 9×9 uniform blur, and the second is a 2D Gaussian function with a standard deviation of 1.6. In both cases, an additive Gaussian noise with $\sigma = \sqrt{2}$ is then added to the blurred images. Similarly to NCSR, restoring an RGB image is done by converting it to the YCbCr color-space, applying the deblurring algorithm

¹⁰In order to handle an arbitrary noise level, σ_f , we rely on the relation $f_{\sigma_f}(\mathbf{y}) = \frac{1}{c}f_5(c \cdot \mathbf{y})$, where $c = 5/\sigma_f$.

Table 1

Deblurring results measured in PSNR [dB] and evaluated on the set of images provided by the authors of NCSR [9]. P^3 and RED build upon TNRD [20] as the denoising engine. We also provide the results obtained by integrating the median filter with the SD RED scheme. PSNR scores being less than 0.01[dB] away from the highest result are highlighted. Note that P^3 does not converge when setting TNRD to be the denoising algorithm. Therefore, we run P^3 for a fixed number of iterations, chosen to achieve the best PSNR on average (otherwise the restoration quality would be significantly inferior). Please refer to Figure 10 and section 6.3 for more details regarding the sensitivity of P^3 to the choice of parameters.

Image	Butterfly	Boats	C. Man	House	Parrot	Lena	Barbara	Starfish	Peppers	Leaves	Average
Deblurring: Uniform kernel, $\sigma = \sqrt{2}$											
Total variation [71]	28.37	29.04	26.82	31.99	29.11	28.33	25.75	27.75	28.43	26.49	28.21
IDD-BM3D [27]	29.21	31.20	28.56	34.44	31.06	29.70	27.98	29.48	29.62	29.38	30.06
ASDS-Reg [72]	28.70	30.80	28.08	34.03	31.22	29.92	27.86	29.72	29.48	28.59	29.84
NCSR	29.68	31.08	28.62	34.31	31.95	29.96	28.10	30.28	29.66	29.98	30.36
P^3 -TNRD	30.32	31.19	28.73	33.90	31.86	30.13	27.21	30.27	30.11	30.08	30.38
RED: SD-median filter	26.10	28.03	25.57	29.81	28.67	27.29	25.62	27.84	27.40	25.45	27.18
RED: SD-TNRD	30.20	31.20	28.67	33.83	31.62	29.98	27.35	30.47	30.10	29.72	30.31
RED: ADMM-TNRD	30.40	31.12	28.71	33.77	31.86	30.03	27.27	30.58	30.11	30.12	30.40
RED: FP-TNRD	30.41	31.12	28.76	33.76	31.83	30.02	27.27	30.57	30.12	30.13	30.40
Deblurring: Gaussian kernel, $\sigma = \sqrt{2}$											
Total variation [71]	30.36	29.36	26.81	31.50	31.23	29.47	25.03	29.65	29.42	29.36	29.22
IDD-BM3D [27]	30.73	31.68	28.17	34.08	32.89	31.45	27.19	31.66	29.99	31.4	30.92
ASDS-Reg [72]	29.83	30.27	27.29	31.87	32.93	30.36	27.05	31.91	28.95	30.62	30.11
NCSR [9]	30.84	31.49	28.34	33.63	33.39	31.26	27.91	32.27	30.16	31.57	31.09
P^3 -TNRD	31.73	31.67	28.08	33.95	33.43	31.52	27.11	32.71	30.94	32.18	31.33
RED: SD-median filter	29.02	30.01	26.45	31.59	31.32	30.00	25.02	30.29	28.53	28.69	29.09
RED: SD-TNRD	31.57	31.53	28.31	33.71	33.19	31.47	26.62	32.46	29.98	31.95	31.08
RED: ADMM-TNRD	31.66	31.55	28.31	33.73	33.33	31.40	26.76	32.49	30.48	31.93	31.16
RED: FP-TNRD	31.66	31.55	28.38	33.74	33.33	31.39	26.76	32.49	30.51	31.93	31.17

on the luminance channel only and then converting the result back to the RGB domain.

Table 1 provides the restoration performance of the three RED schemes—the SD, the ADMM, and the fixed-point (FP) methods—along with the results of the¹¹ P^3 , the state-of-the-art NCSR and IDD-BM3D [27], and two additional baseline deblurring methods [71, 72]. For brevity, only the SD scheme is presented when considering the basic median filter as a denoiser. The performance is evaluated using the peak signal to noise ratio (PSNR) measure, higher is better, computed on the luminance channel of the ground truth and the estimated image. The parameters of the proposed approach, as well as the ones of P^3 , are tuned to achieve the best performance on this dataset; in the case of the TNRD denoiser, these are depicted in Table 2 and 3, respectively. In the setting of the median filter, which extracts the median value of a 3×3 window, we choose to run the suggested SD scheme for $N = 400$ iterations with $\lambda = 0.12$ for the uniform PSF and $N = 200$ with $\lambda = 0.225$ for the Gaussian PSF.

Several remarks are to be made with regard to the obtained results. When the image is degraded by a Gaussian blur kernel, integrating the median filter in the proposed framework leads to a surprising restoration performance that is similar to the total variation deblurring [71]. Furthermore, by choosing the state-of-the-art TNRD to be our denoising engine we achieve results that are competitive with the impressive NCSR and IDD-BM3D methods, which are specifically designed to tackle the deblurring task. Notice that the three versions of the proposed framework obtain a similar PSNR score. However, while ADMM and the fixed-point variants are of similar complexity, SD requires many more steps to converge and

¹¹We note that P^3 using TNRD has never appeared in an earlier publication, and it is brought here in order to let P^3 perform as best as it possibly can.

Table 2

The set of parameters being used in our framework, leading to the deblurring results reported in Table 1 when plugging the TNRD [20] denoiser.

PSF	Proposed approach: Deblurring			
	Parameter	Steepest descent	Fixed point	ADMM
Uniform	N	1500	200	200
	σ_f	3.25	3.25	3.25
	λ	0.02	0.02	0.02
	m_1	–	closed-form using FFT	closed-form using FFT
	m_2	–	–	1
	β	–	–	0.001
Gaussian	N	1500	200	200
	σ_f	4.1	4.1	4.1
	λ	0.01	0.01	0.01
	m_1	–	closed-form using FFT	closed-form using FFT
	m_2	–	–	1
	β	–	–	0.001

Table 3

The set of parameters being used by the P^3 method, leading to the deblurring results reported in Table 1 when plugging the TNRD [20] denoiser.

PSF	P^3 : Deblurring	
	Parameter	Value
Uniform	N	200
	α	1.02
	β_0	0.0007
	β_k	$\alpha^k \cdot \beta_0$
	λ	$512 \cdot \beta_0$
	σ_f	$\sqrt{\lambda/\beta_k}$
Gaussian	N	200
	α	1.02
	β_0	0.0007
	β_k	$\alpha^k \cdot \beta_0$
	λ	$320 \cdot \beta_0$
	σ_f	$\sqrt{\lambda/\beta_k}$

thereby more applications of the denoiser. As our last observation, based on the obtained results we conclude that the proposed approach is equivalent in quality to the alternative P^3 framework. However, tuning the parameters of the proposed algorithm is significantly simpler than the ones of P^3 ; while in P^3 the parameters should be modified throughout the iterations, in our approach these are always fixed (refer to section 6.3 for a broader discussion).

An illustration of the convergence of the proposed approach using the three different numerical schemes (SD, ADMM, and fixed-point) and the two denoisers (median filter and TNRD) is given in Figure 5. As can be observed, the three algorithms indeed converge, but at different rates; SD is the slowest, while ADMM is the fastest. The fixed-point strategy is slightly slower than the ADMM alternative, but requires only one application of the denoiser

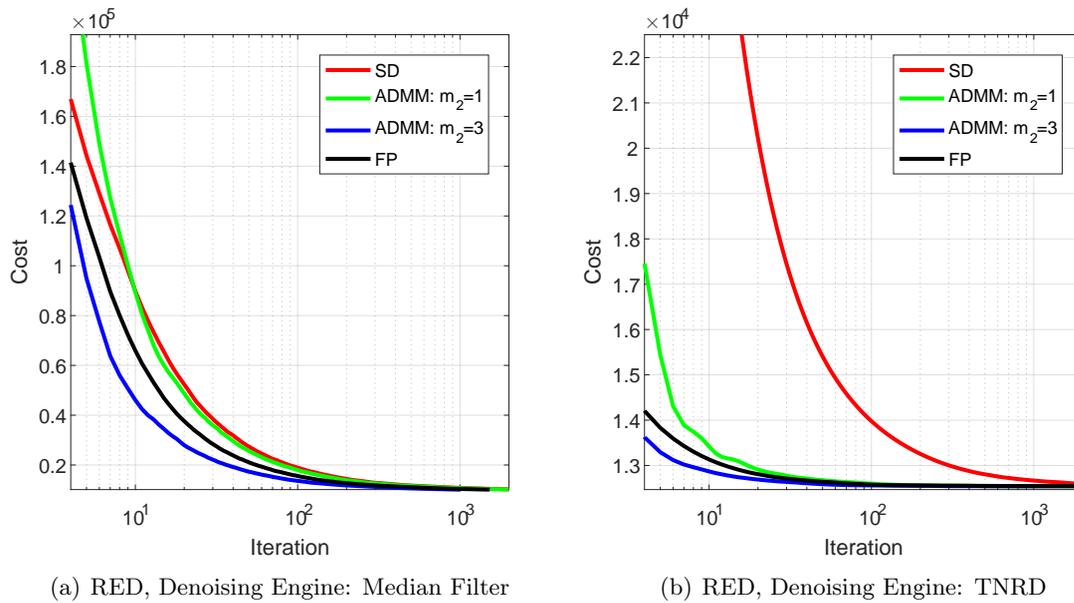


Figure 5. An illustration of the convergence of RED using three proposed numerical schemes—SD (red), ADMM with $m_2 = 1$ (green), ADMM with $m_2 = 3$ (blue), and fixed-point (black). These are applied on the image *Leaves*, degraded by a Gaussian PSF, when two denoising engines are tested: (a) the median filter and (b) TNRD [20].

per iteration while ADMM demands $m_2 = 3$ such operations. However, when comparing fixed-point to ADMM with the setting $m_2 = 1$ (now both have the same computational cost) we observe that fixed-point is faster.

The above discussion is also supported visually in Figures 6 and 7, comparing the proposed method to P^3 and NCSR both for uniform and Gaussian PSFs. As can be seen, by plugging the median filter into RED we improve the quality of the blurry image, yet the gap in performance between this simplistic approach and the state-of-the-art one is substantial. Once relying on TNRD, we obtain an efficient deblurring machine that has comparable results to P^3 , and both are competitive or even slightly better than NCSR.

6.2. Image super-resolution. Similarly to the previous subsection, we imitate the super-resolution experiments done in [9]. To this end, a low-resolution image is generated by blurring the ground-truth high-resolution one with a 7×7 Gaussian blur kernel with standard deviation 1.6, followed by downsampling by a factor of 3 in each axis. Next, WGN of standard deviation 5 is added to the low-resolution images. The upscaling of an RGB image is done by transforming it first to the YCbCr color-space, super-resolving the luminance channel using the proposed approach (or the baseline methods), while the chroma channels are upscaled by bicubic interpolation. Finally, the outcome is converted back to the RGB color-space.

In terms of PSNR, Table 4 presents the restoration performance of the three variants of the proposed approach in addition to the ones of the P^3 , the NCSR, and the ASDS-Reg [72] algorithms. Similarly to the deblurring case, the PSNR is computed on the luminance channel

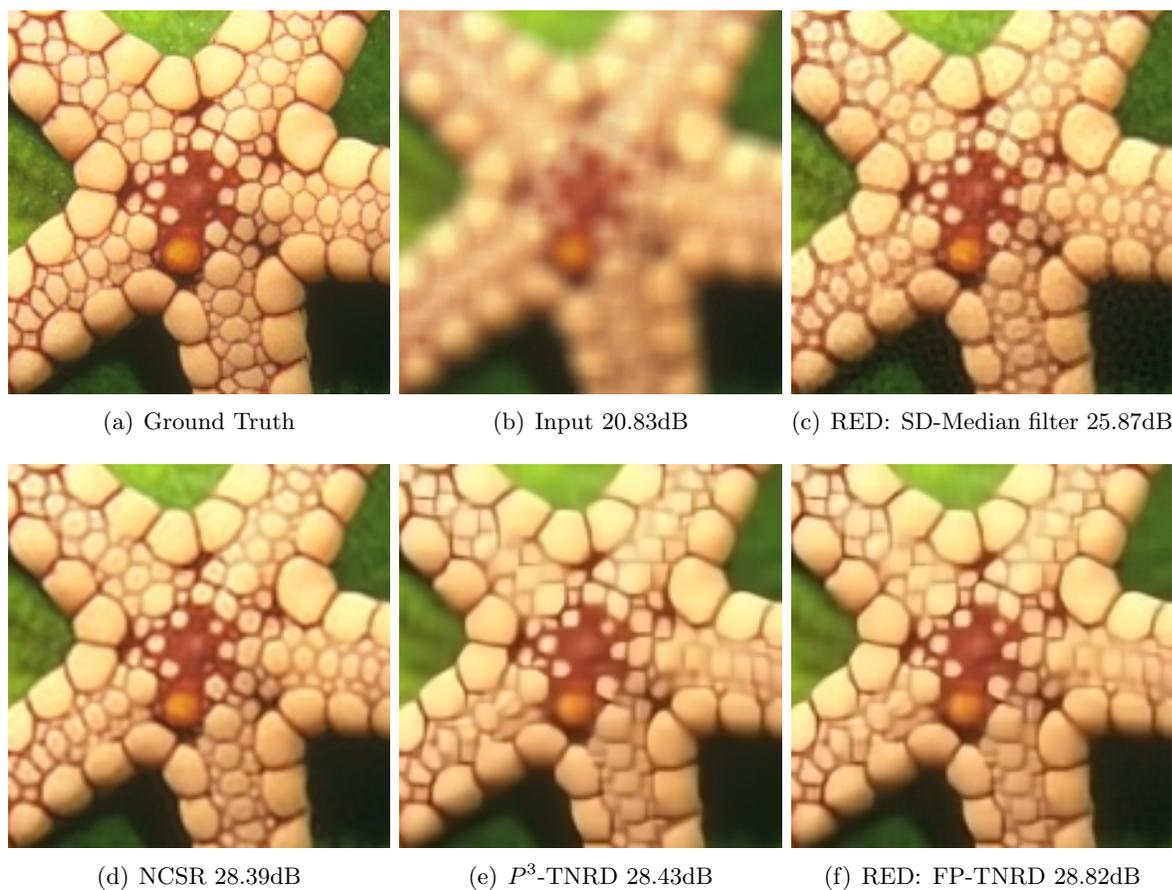


Figure 6. Visual comparison of deblurring a cropped area from the image *Starfish*, degraded by a uniform PSF, along with the corresponding PSNR [dB] score.

only. In the case of the TNRD denoiser, the parameters that are used in our approach and in P^3 are listed in Tables 5 and 6, respectively. In the simpler case of the median filter (defined by a 3×3 window), the number of iterations of the proposed SD algorithm is set to $N = 50$ with a parameter $\lambda = 0.0325$.

Interestingly, when setting the median filter to be our denoising engine we get a 2.19dB improvement (on average) over the bicubic interpolation. Alternatively, when choosing a stronger denoiser such as TNRD, we achieve state-of-the-art results. Notably, P^3 and the three variants of the proposed approach lead to similar restoration performance, consistent with the observation that was made in the context of the deblurring problem. These support once again the claim that our framework is a tangible alternative to P^3 . Figures 8 and 9 visually compare the proposed method to P^3 and also to the state-of-the-art NCSR. As shown, the three algorithms offer an impressive restoration with sharp and clear edges, complying with the quantitative results which are given in Table 4.

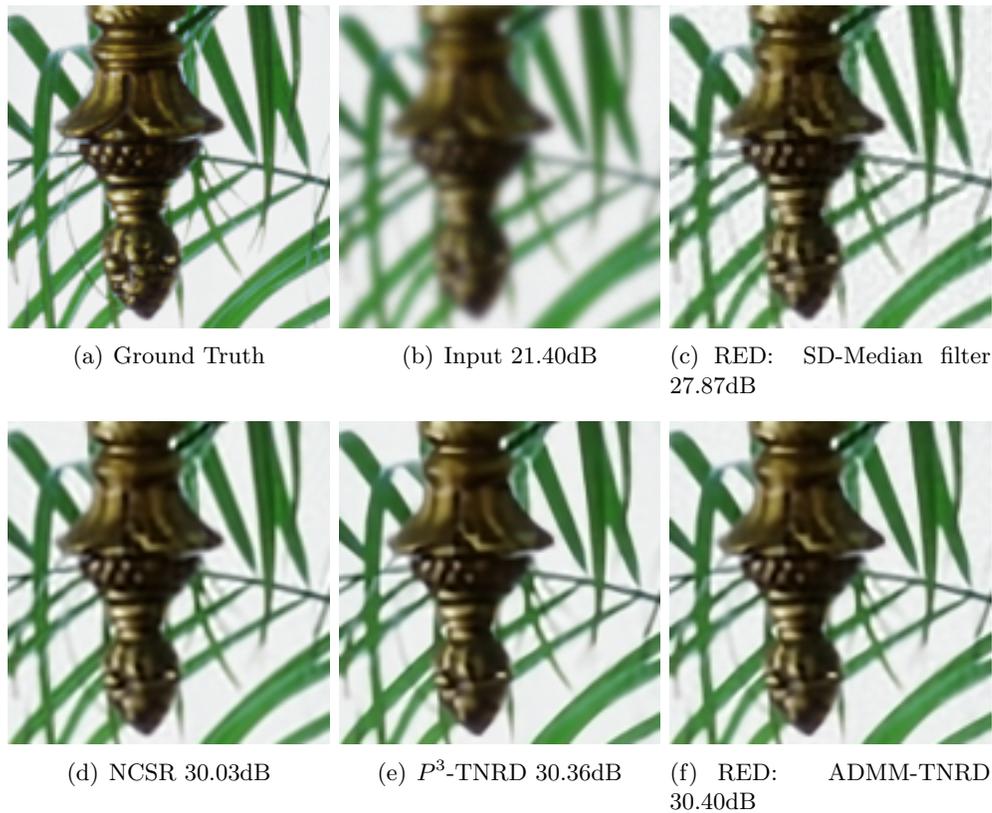


Figure 7. Visual comparison of deblurring a cropped area from the image *Leaves*, degraded by a Gaussian PSF, along with the corresponding PSNR [dB] score.

Table 4

Super-resolution results, measured in terms of PSNR [dB] and evaluated on the set of images provided by the authors of NCSR [9]. The P^3 and our SD, ADMM, and fixed-point (FP) methods build upon TNRD [20] as the denoising engine. We also provide the results obtained by integrating the median filter in the SD scheme. PSNR scores being less than 0.01[dB] away from the highest result are highlighted. Similarly to the deblurring case, the P^3 does not converge here as well. Therefore, we run it for a fixed number of iterations, manually tuned to achieve the highest PSNR on average. This behavior is demonstrated in detail in Figure 10 and discussed in section 6.3.

Super-resolution, scaling = 3, $\sigma = \sqrt{2}$										
Image	Butterfly	flower	Girl	Parth.	Parrot	Raccoon	Bike	Hat	Plants	Average
Bicubic	20.74	24.74	29.61	24.04	25.43	26.20	20.75	27.00	27.62	25.13
ASDS-Reg [72]	26.06	27.83	31.87	26.22	29.01	28.01	23.62	29.61	31.18	28.16
NCSR [9]	26.86	28.08	32.03	26.38	29.51	28.03	23.80	29.94	31.73	28.48
P^3 -TNRD	27.13	28.23	32.08	26.50	29.65	27.95	24.04	30.30	31.78	28.63
RED: SD-median filter	24.44	27.24	31.13	25.80	27.76	27.65	22.89	28.69	30.24	27.32
RED: SD-TNRD	27.37	28.23	32.08	26.54	29.43	27.98	24.04	30.36	31.79	28.65
RED: ADMM-TNRD	27.22	28.24	32.08	26.51	29.41	27.97	23.96	30.35	31.77	28.61
RED: FP-TNRD	27.26	28.24	32.08	26.52	29.42	27.97	23.97	30.35	31.77	28.62

6.3. Robustness to the choice of parameters. In this subsection we test the robustness of RED to the choice of its parameters and contrast it to P^3 . To this end, we choose the single image super-resolution problem as a case study, described in the previous subsection. Figure 10(a) plots the average PSNR obtained by the different approaches as a function of the

Table 5

The set of parameters being used in our framework, leading to the super-resolution results reported in Table 4 when plugging the TNRD [20] denoiser.

Proposed approach: Super-resolution			
Parameter	Steepest descent	Fixed point	ADMM
N	1500	200	200
σ_f	3	3	3
λ	0.008	0.008	0.008
m_1	–	200 (until convergence)	200 (until convergence)
m_2	–	–	1
β	–	–	0.001

Table 6

The set of parameters being used in the P^3 method, leading to the super-resolution results reported in Table 4 when plugging the TNRD [20] denoiser.

P^3 : Super-resolution	
Parameter	Value
N	200
α	1.02
β_0	0.001
β_k	$\alpha^k \beta_0$
λ	$360 \beta_0$
σ_f	$\sqrt{\lambda/\beta_k}$

outer iterations. One can observe that RED (in all its forms) converges to a similar PSNR value. Also, an increase of m_2 (the number denoising steps within each iteration) leads to an improved rate of convergence of ADMM. On the other hand, the curve describing P^3 shows an unstable behavior and tendency to decrease in the PSNR after the first 200 iterations. Note that no tool has been suggested in the literature so far to automatically stop P^3 for extracting the best performing outcome.

This unstable nature of P^3 appears again as we modify the values of α and β_0 . Figure 10(b) shows the behavior of P^3 for several settings of these two parameters, clearly exhibiting an erratic behavior. One could observe that for specific choices of these two parameters, convergence is obtained, as manifested by the flattened curves. However, this is a fake convergence, caused by a large enough value of β_k . We stress that, in principle, a change in β_0 is expected to modify the convergence rate of ADMM, but the steady-state outcome should remain the same. However, when observing the curves in Figure 10(b), it is clear that this is not the case in P^3 .

Back to RED, we repeat a similar experiment and test the sensitivity (or better yet the robustness) of ADMM to the choice of β . Figure 10(c) shows that different values of β indeed affect the convergence rate (as expected), but the PSNR of the final outcome is always the same. This figure also indicates that the more accurate the solution of **Part 2** in Figure 3 (obtained by increasing the value of m_2), the better the convergence rate.

The sensitivity of RED to the choice of σ_f —the input noise level to the denoiser—is depicted in Figure 10(d). Notice that the choice of σ_f directly affects the proposed regularizer,

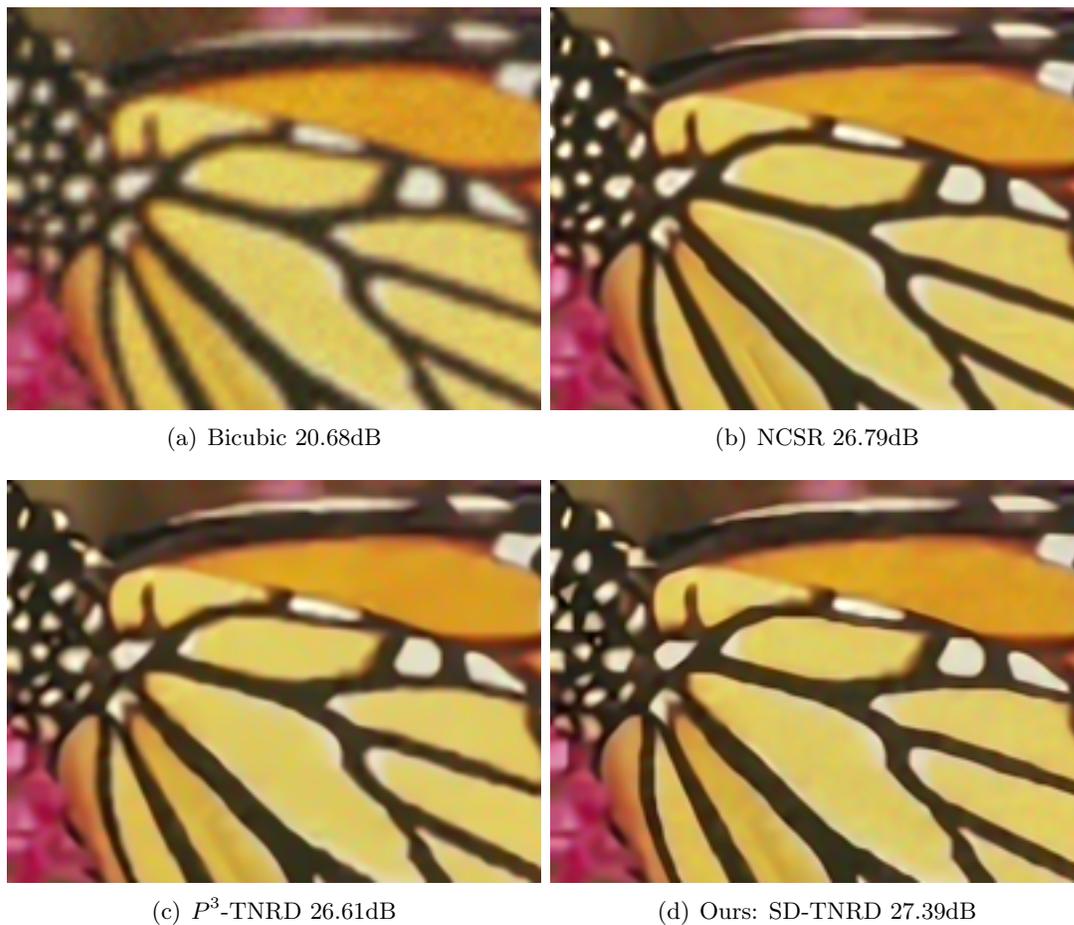


Figure 8. Visual comparison for upscaling by a factor of 3 for the image *Butterfly*, along with the corresponding PSNR [dB] score.

given by $\lambda\rho(\mathbf{x}, \sigma_f) = \frac{\lambda}{2} \mathbf{x}^T [\mathbf{x} - f_{\sigma_f}(\mathbf{x})]$. As such, a change in σ_f is expected to modify the objective and thereby the resulting PSNR, as shown empirically in Figure 10(d) for the fixed-point method.¹² Clearly, a similar behavior is expected to occur when modifying the weight of the regularizer, λ , in which we choose to omit from this experimental part for brevity.

7. Conclusions. The idea put forward in this paper is strange—using a denoising engine within the regularization term in order to handle general inverse problems. A surprising outcome of this proposal is the fact that differentiation of the regularization terms remains tractable, still using the very same denoiser engine and *not* its derivative. This led us to the proposed scheme, termed Regularization by Denoising (RED). We have shown and discussed various appealing properties of this approach, such as convexity and its relation to advanced Laplacian smoothing. We have contrasted this scheme with the Plug-and-Play Prior (P^3)

¹²One expects that we could tune this parameter using existing techniques such as SURE, but we leave this topic for future research.

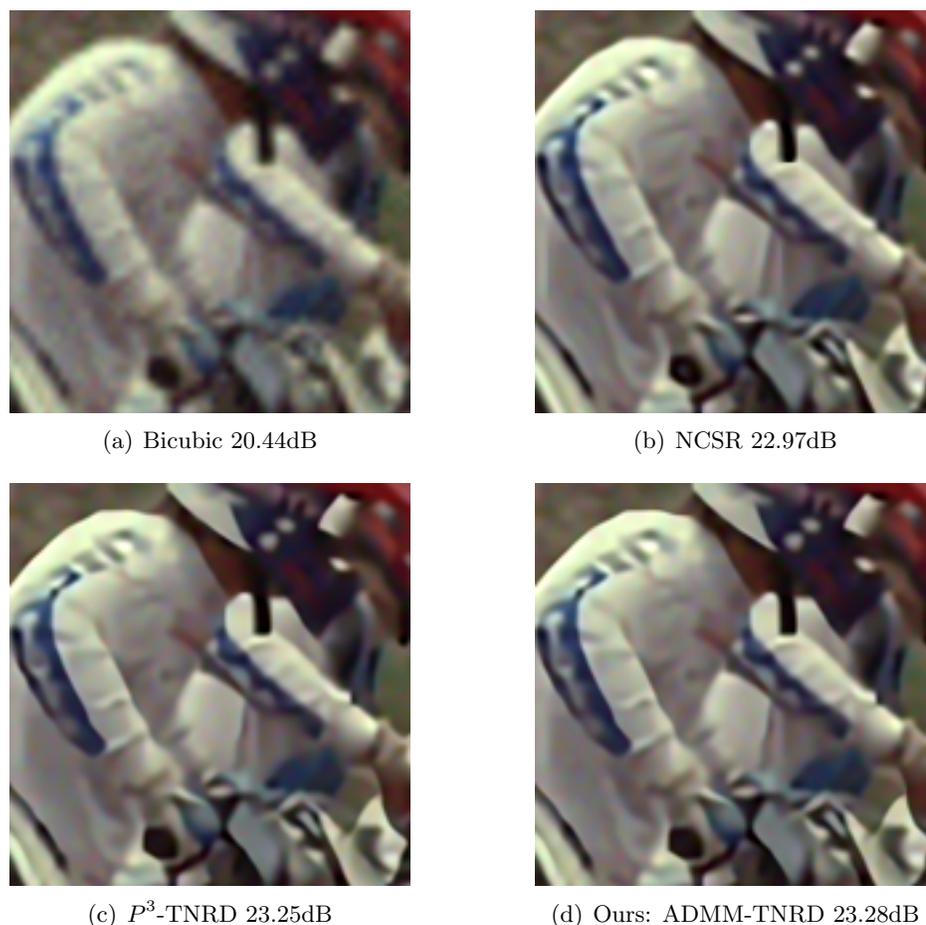


Figure 9. Visual comparison for upscaling by a factor of 3 for the image *Bike*, along with the corresponding PSNR [dB] score.

method [30], and we have provided experiments that demonstrate the validity of the regularization strategy and the resulting competitive performance.

Is there a wider message in this work? Could it be that one could use more general f functions and still form the proposed regularization? We have seen that all that it takes is the availability of the directional derivative of this function in order to follow all through. Strangely enough, we have shown how even a median filter could fit into this scheme, despite the fact that it does not have any relation to Gaussian denoising. More work is required to investigate the ability to further generalize RED to other and perhaps more daring regularization functionals.

A key question we have left open at this stage is the setting of the parameter σ_f . We chose this to be a fixed value, but we did not address the question of its influence on the overall performance or whether a varying value strategy could lead to a benefit. More work is required here as well.

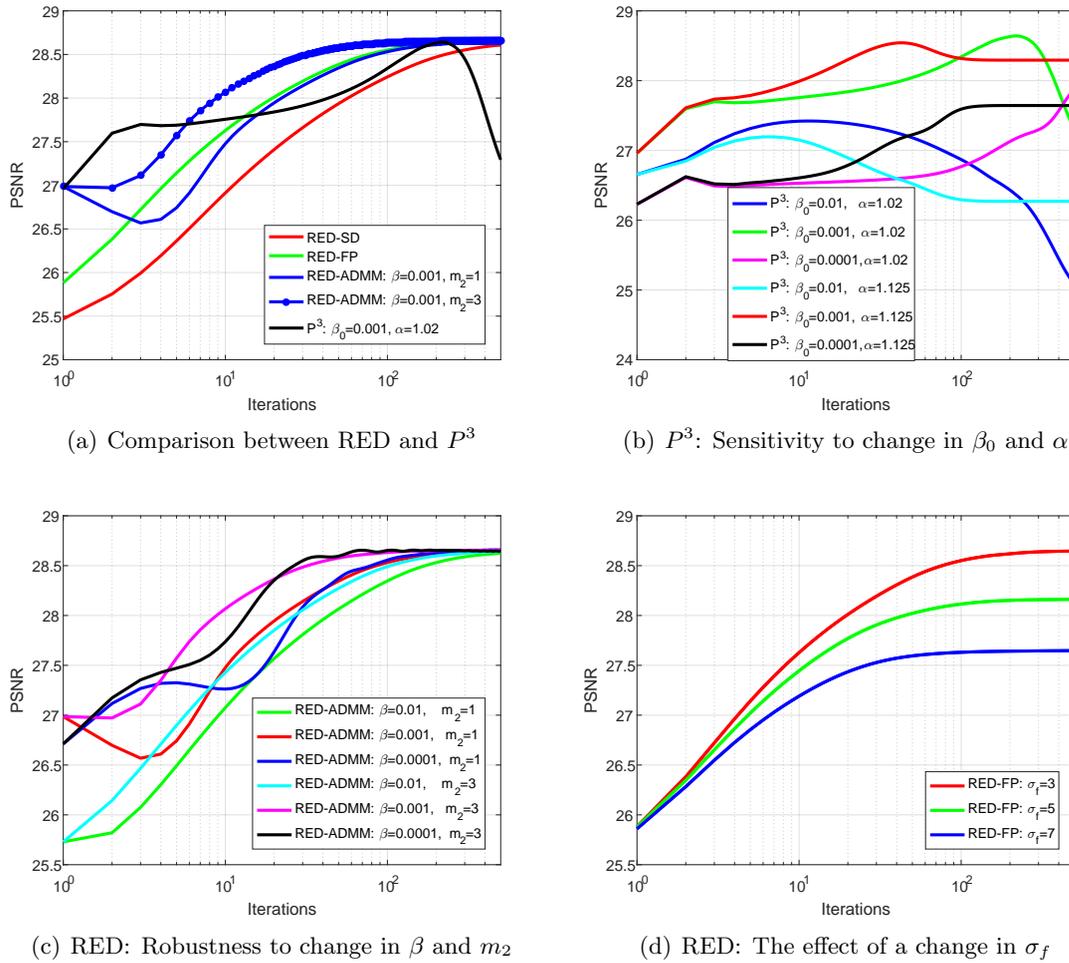


Figure 10. Image super-resolution. Average PSNR as a function of the iterations of the different methods. The test images are those of Table 4.

Appendix A. Can we mimic any prior? So far we have shown that denoisers $f(\mathbf{x})$ can be used to define the powerful family of priors $\mathbf{x}^T(\mathbf{x} - f(\mathbf{x}))$ that can regularize a wide variety of inverse problems. Now, let's consider the reverse question: For a given $\rho(\mathbf{x})$, what is the effective $f(\mathbf{x})$ behind it (if any)? More specifically, given $\rho(\mathbf{x})$, we wish to find $f(\mathbf{x})$ such that

$$(50) \quad \frac{1}{2} \mathbf{x}^T (\mathbf{x} - f(\mathbf{x})) = \rho(\mathbf{x}).$$

Recall that one of the key conditions we assumed for the denoiser $f(\mathbf{x})$ is that it is homogeneous of degree 1,

$$(51) \quad f(c\mathbf{x}) = cf(\mathbf{x}).$$

This immediately notifies us that the $\rho(\mathbf{x})$ we wish to mimic in (50) must be 2-homogeneous

because

$$(52) \quad \rho(c\mathbf{x}) = \frac{1}{2}(c\mathbf{x})^T(c\mathbf{x} - f(c\mathbf{x})) = \frac{1}{2}c^2\mathbf{x}^T(\mathbf{x} - f(\mathbf{x})) = c^2\rho(\mathbf{x}).$$

Of course, not all priors satisfy this condition, but the class of priors that do is wide. For instance, while the total-variation function $\text{TV}(\mathbf{x}) = \|\nabla x\|_1$ is only 1-homogeneous, its square $\|\nabla x\|_1^2$ keeps us in business. More generally, since any norm is *by definition* 1-homogeneous, all regularizers of the type $\rho(\mathbf{x}) = \|\mathbf{A}\mathbf{x}\|_q^2$ for $q = 1, 2, \dots, \infty$ are 2-homogeneous:¹³

$$(53) \quad \rho(c\mathbf{x}) = \|\mathbf{A}(c\mathbf{x})\|_q^2 = (c\|\mathbf{A}\mathbf{x}\|_q)^2 = c^2\|\mathbf{A}\mathbf{x}\|_q^2 = c^2\rho(\mathbf{x}).$$

Squared norms are not the only functions at our disposal. Beyond norms, any k -homogeneous function (such as homogeneous polynomials of degree k) can be made 2-homogeneous by raising to power $2/k$. As well, order statistics such as maximum, minimum, and median are 1-homogeneous and can be squared to the same end.

With the above discussion, we move forward with the assumption that we have a 2-homogeneous prior $\rho(\mathbf{x})$ in our hands. Let's recall that the task is to find $f(\mathbf{x})$ so that

$$(54) \quad \frac{1}{2}\mathbf{x}^T(\mathbf{x} - f(\mathbf{x})) = \rho(\mathbf{x}).$$

We proceed by differentiating both sides:

$$\begin{aligned} \mathbf{x} - \frac{1}{2}\nabla(\mathbf{x}^T f(\mathbf{x})) &= \nabla\rho(\mathbf{x}), \\ \mathbf{x} - \frac{1}{2}(f(\mathbf{x}) + \nabla f(\mathbf{x})\mathbf{x}) &= \nabla\rho(\mathbf{x}), \\ \mathbf{x} - \frac{1}{2}(f(\mathbf{x}) + f(\mathbf{x})) &= \nabla\rho(\mathbf{x}), \\ \mathbf{x} - f(\mathbf{x}) &= \nabla\rho(\mathbf{x}), \end{aligned}$$

where in the last step we have invoked the directional derivative expression developed in (14). The solution, $f(\mathbf{x})$, is a denoiser explicitly defined in terms of the prior,

$$(55) \quad f(\mathbf{x}) = \mathbf{x} - \nabla\rho(\mathbf{x}).$$

This is quite a natural result in retrospect—it resembles an SD step.¹⁴ To see this, consider the denoising problem regularized by $\rho(\mathbf{x})$:

$$(56) \quad \underset{\mathbf{x}}{\text{Argmin}} \frac{1}{2}\|\mathbf{x} - \mathbf{y}\|^2 + \rho(\mathbf{x}).$$

One step of SD with a step-size of 1 and initial condition $\mathbf{x}_0 = \mathbf{y}$ would read as

$$(57) \quad \mathbf{x}_1 = \mathbf{x}_0 - (\mathbf{x}_0 - \mathbf{y} + \nabla\rho(\mathbf{x}_0)) = \mathbf{y} - \nabla\rho(\mathbf{y}),$$

¹³The sparsity-inspired L_0 is an exception that cannot be treated this way.

¹⁴It is interesting to note the resemblance of this expression to a similar expression that holds for proximal operators (see [60]). In this context, the class of denoisers we have described using Conditions 1 and 2 is a more general form of proximal mappings.

just as shown above as a denoising on \mathbf{y} .

Appendix B. Kernelizing priors. We showed in the previous section that a prior with the right properties implies a denoiser beneath and that this denoiser is directly given by the gradient of the prior. Next, let's address a related, but more specific question: Given a prior $\rho(\mathbf{x})$, does it imply a denoising filter of the *pseudo-linear form* $f(\mathbf{x}) = \mathbf{W}(\mathbf{x})\mathbf{x}$? These types of filters are of course of great interest because they reveal the kind of weighted averaging done by the denoiser, and they are easy to implement. Given \mathbf{x} , we can compute the weights $\mathbf{W}(\mathbf{x})$ in one step and apply them as $\mathbf{W}(\mathbf{x})\mathbf{x}$ to the image in another step. Some of the most popular filters to date, such as bilateral, NLM, and K-SVD, are of this convenient form.

Before we go about finding the hidden filter matrix $\mathbf{W}(\mathbf{x})$, let's illustrate a useful property of the pseudo-linear filters. Take $f(\mathbf{x}) = \mathbf{W}(\mathbf{x})\mathbf{x}$, and again invoke the expression $f(\mathbf{x}) = \nabla f(\mathbf{x})\mathbf{x}$ developed in (14). Substitution gives

$$(58) \quad \mathbf{W}(\mathbf{x})\mathbf{x} = \nabla f(\mathbf{x})\mathbf{x} \quad \implies \quad (\nabla f(\mathbf{x}) - \mathbf{W}(\mathbf{x}))\mathbf{x} = 0$$

for all \mathbf{x} . From this we conclude that the gradient of pseudo-linear filters is in fact the weight matrix

$$(59) \quad \nabla f(\mathbf{x}) = \mathbf{W}(\mathbf{x}).$$

Now we can go after the weight matrix by taking the analysis from the previous section one step further and computing the second derivative of (50). Starting with the expression (55) that arose from the first derivative, we differentiate again,

$$\begin{aligned} f(\mathbf{x}) &= \mathbf{x} - \nabla \rho(\mathbf{x}), \\ \nabla f(\mathbf{x}) &= \mathbf{I} - \nabla(\nabla \rho(\mathbf{x})), \\ \nabla^2 f(\mathbf{x}) &= \mathbf{I} - \mathbf{H}(\rho(\mathbf{x})). \end{aligned}$$

Replacing $\nabla f(\mathbf{x}) = \mathbf{W}(\mathbf{x})$, we obtain the pleasing result that the weight matrix implied by the prior is the identity matrix minus the Hessian of the prior,

$$(60) \quad \mathbf{W}(\mathbf{x}) = \mathbf{I} - \mathbf{H}(\rho(\mathbf{x})),$$

or, posed another way, $\mathbf{L}(\mathbf{x}) = \mathbf{I} - \mathbf{W}(\mathbf{x}) = \mathbf{H}(\rho(\mathbf{x}))$ (i.e., the Laplacian filter is directly given by the Hessian of the prior, which is not surprising, bearing in mind that we seek a relation of the form $\rho(\mathbf{x}) = \mathbf{x}^T \mathbf{L}(\mathbf{x})\mathbf{x}$). What we have done here is to “kernelize” the regularizer and find an explicit expression for the weights of the implied filter. Several observations about this result are in order:

- **Convexity:** If $\rho(\mathbf{x})$ is convex, then its Hessian is symmetric positive semidefinite (PSD), and therefore $\mathbf{L}(\mathbf{x})$ is PSD. Furthermore, if $\eta(\mathbf{L}(\mathbf{x})) \leq 1$, we get that (i) $\mathbf{W}(\mathbf{x})$ has spectral radius equal or smaller than 1, implying that strong passivity condition (Condition 2) is guaranteed; and (ii) $\mathbf{W}(\mathbf{x}) = \mathbf{I} - \mathbf{L}(\mathbf{x})$ is also PSD—a desirable property [19].

- **Homogeneity:** Since $\rho(\mathbf{x})$ is 2-homogeneous, its gradient is 1-homogeneous. We can show this by differentiating:

$$\begin{aligned}\rho(c\mathbf{x}) &= c^2\rho(\mathbf{x}), \\ \nabla_{\mathbf{x}}\rho(c\mathbf{x}) &= c^2\nabla_{\mathbf{x}}\rho(\mathbf{x}), \\ \frac{\partial c\mathbf{x}}{\partial \mathbf{x}} \nabla_{c\mathbf{x}}\rho(c\mathbf{x}) &= c^2\nabla_{\mathbf{x}}\rho(\mathbf{x}), \\ c\nabla_{c\mathbf{x}}\rho(c\mathbf{x}) &= c^2\nabla_{\mathbf{x}}\rho(\mathbf{x}), \\ \nabla_{c\mathbf{x}}\rho(c\mathbf{x}) &= c\nabla_{\mathbf{x}}\rho(\mathbf{x}).\end{aligned}$$

Similarly the Hessian $\mathbf{H}(\rho(\mathbf{x}))$ is *invariant*¹⁵ to scaling of the image \mathbf{x} , and so is the implied filter matrix $\mathbf{W}(\mathbf{x})$. Consequently, the applied filter $\mathbf{W}(\mathbf{x})\mathbf{x}$ is 1-homogeneous, which again is consistent with our earlier conditions.

- **Row stochasticity:** Let's recall the expression we derived above for the filter in terms of the Hessian of the prior,

$$(61) \quad \mathbf{W}(\mathbf{x}) = \mathbf{I} - \mathbf{H}(\rho(\mathbf{x})).$$

If a filter defined this way is to be row-stochastic, we would have (for every \mathbf{x}), $\mathbf{W}(\mathbf{x})\mathbf{1} = (\mathbf{I} - \mathbf{H}(\rho(\mathbf{x})))\mathbf{1} = \mathbf{1}$, or equivalently

$$(62) \quad \mathbf{H}(\rho(\mathbf{x}))\mathbf{1} = \mathbf{0}.$$

This relation does not hold in general. However, consider defining the prior in terms of the gradient of the image instead of the image itself. Namely, this involves a change of variables in the prior $\rho(\mathbf{x})$ from \mathbf{x} to $\mathbf{D}\mathbf{x}$, where \mathbf{D} is the gradient (e.g., difference) operator. For instance, instead of $\rho(\mathbf{x}) = \|\mathbf{x}\|_1^2$, consider $\rho_{\mathbf{D}}(\mathbf{x}) = \|\mathbf{D}\mathbf{x}\|_1^2$. The Hessian of the prior under this linear transformation is given by the chain rule as

$$(63) \quad \mathbf{H}(\rho_{\mathbf{D}}(\mathbf{x})) = \mathbf{D}^T\mathbf{H}(\rho(\mathbf{x}))\mathbf{D}.$$

This Hessian, when applied to the constant vector $\mathbf{1}$, will vanish for all \mathbf{x} , since $\mathbf{D}\mathbf{1} = \mathbf{0}$:

$$(64) \quad \mathbf{D}^T\mathbf{H}(\rho(\mathbf{x}))\mathbf{D}\mathbf{1} = \mathbf{0},$$

so the filter resulting from this Hessian is row-stochastic. In fact, the same can be said for column stochasticness, since $\mathbf{1}^T\mathbf{D}^T = \mathbf{0}$.

Appendix C. More on homogeneity. The concept of homogeneity of a filter played an important role in the development of the key results of the paper. So it is worth saying a bit more about it and to question whether this condition is satisfied for some popular and familiar filters.

¹⁵Or 0-homogeneous.

C.1. Nonlocal means and bilateral filter. A general construction of a denoising filter could be based on a symmetric positive semidefinite kernel $\mathbf{K}_{i,j}(\mathbf{x}) = \mathbf{K}(x_i, x_j) \geq 0$ from which the filter matrix $\mathbf{W}(\mathbf{x})$ is constructed by normalization. More specifically,

$$(65) \quad \mathbf{W}_{i,j}(\mathbf{x}) = \frac{\mathbf{K}_{i,j}(\mathbf{x})}{\sum_{i=1}^n \mathbf{K}_{i,j}(\mathbf{x})}.$$

Whether such a denoiser is homogeneous very much depends on the choice of the kernel \mathbf{K} . For instance, if the kernel is homogeneous of *any* degree p , then the resulting filter matrix is invariant to scaling through cancellation,

$$(66) \quad \mathbf{W}_{i,j}(c\mathbf{x}) = \frac{\mathbf{K}_{i,j}(c\mathbf{x})}{\sum_{i=1}^n \mathbf{K}_{i,j}(c\mathbf{x})} = \frac{c^p \mathbf{K}_{i,j}(\mathbf{x})}{\sum_{i=1}^n c^p \mathbf{K}_{i,j}(\mathbf{x})} = \mathbf{W}_{i,j}(\mathbf{x}).$$

Examples of homogeneous kernels include (homogeneous) polynomials and several others [45] which are not in common use in image processing. Most commonly used kernels are the exponentials (Gaussian to be exact), which are used in the bilateral or NLM cases. The Gaussian function is not homogeneous, but as we will show below, the resulting pseudo-linear filter is nearly so. We will show that for $c = 1 + \epsilon$ with very small ϵ we have 1-homogeneity for the NLM-type filters, namely,

$$(67) \quad f(c\mathbf{x}) = \mathbf{W}(c\mathbf{x})(c\mathbf{x}) = c\mathbf{W}(\mathbf{x})\mathbf{x} = cf(\mathbf{x}).$$

The i, j th element of the filter weight matrix for the NLM filter¹⁶ is

$$(68) \quad \mathbf{W}_{i,j}(\mathbf{x}; \sigma) = \frac{e_{ij}(\sigma)}{d_j(\sigma)},$$

where

$$e_{ij}(\sigma) = \exp(-\|\mathbf{R}_i\mathbf{x} - \mathbf{R}_j\mathbf{x}\|^2/2\sigma^2),$$

$$d_j(\sigma) = \sum_{i=1}^n \exp(-\|\mathbf{R}_i\mathbf{x} - \mathbf{R}_j\mathbf{x}\|^2/2\sigma^2),$$

where $\mathbf{R}_i\mathbf{x}$ is a patch centered at pixel position i , extracted from the image; the normalization constant $d_j(\sigma)$ is given by summing across the rows of the kernel matrix.

Do these weights change much when the image \mathbf{x} is replaced by a scaled version $c\mathbf{x}$? First, note that if σ is nearly zero, then all weights are essentially equal to $1/n$ and therefore they are automatically invariant to scaling of the image. Next, let's consider the other extreme where the value of σ is away from zero. Now, note that the scaling in \mathbf{x} can be absorbed in the parameter σ as follows:

$$(69) \quad \exp(-\|c\mathbf{R}_i\mathbf{x} - c\mathbf{R}_j\mathbf{x}\|^2/2\sigma^2) = \exp(-\|\mathbf{R}_i\mathbf{x} - \mathbf{R}_j\mathbf{x}\|^2/2(\sigma/c)^2).$$

¹⁶The bilateral filter, which includes a spatial distance weight, can be treated similarly, since the spatial weights are invariant to scaling of the values of the image in any case.

Second, the effect of this (multiplicative) scaling can be approximated by an additive perturbation,

$$(70) \quad \frac{\sigma}{c} = \frac{\sigma}{1 + \epsilon} \approx \sigma - \epsilon\sigma = \sigma + \delta.$$

We now compute an approximation to $\mathbf{W}_{i,j}(\mathbf{x}, \sigma + \delta)$ using a Taylor series:

$$\begin{aligned} \mathbf{W}_{i,j}((1 + \epsilon)\mathbf{x}; \sigma) &\approx \mathbf{W}_{i,j}(\mathbf{x}; \sigma + \delta) \\ &\approx \mathbf{W}_{i,j}(\mathbf{x}; \sigma) + \delta \frac{\partial \mathbf{W}_{i,j}(\mathbf{x}; \sigma)}{\partial \sigma}. \end{aligned}$$

The derivative of the weight values will be calculated in terms of the functions $e_{ij}(\sigma)$ and $d_j(\sigma)$ as follows:

$$\begin{aligned} \frac{\partial \mathbf{W}_{i,j}(\mathbf{x}; \sigma)}{\partial \sigma} &= \frac{\partial}{\partial \sigma} \left(\frac{e_{ij}(\sigma)}{d_j(\sigma)} \right) \\ &= \frac{e'_{ij}(\sigma)d_j(\sigma) - e_{ij}(\sigma)d'_j(\sigma)}{d_j^2(\sigma)} \\ &= \frac{e'_{ij}(\sigma)}{d_j(\sigma)} - \frac{e_{ij}(\sigma)}{d_j(\sigma)} \frac{d'_j(\sigma)}{d_j(\sigma)} \\ &= \frac{\|\mathbf{R}_i\mathbf{x} - \mathbf{R}_j\mathbf{x}\|^2}{\sigma^3} \frac{e_{ij}(\sigma)}{d_j(\sigma)} - \frac{e_{ij}(\sigma)}{d_j(\sigma)} \frac{d'_j(\sigma)}{d_j(\sigma)} \\ &= \frac{\|\mathbf{R}_i\mathbf{x} - \mathbf{R}_j\mathbf{x}\|^2}{\sigma^3} \mathbf{W}_{i,j}(\mathbf{x}; \sigma) - \frac{d'_j(\sigma)}{d_j(\sigma)} \mathbf{W}_{i,j}(\mathbf{x}; \sigma) \\ &= \left(\frac{\|\mathbf{R}_i\mathbf{x} - \mathbf{R}_j\mathbf{x}\|^2}{\sigma^3} - \frac{d'_j(\sigma)}{d_j(\sigma)} \right) \mathbf{W}_{i,j}(\mathbf{x}; \sigma). \end{aligned}$$

Therefore,

$$(71) \quad \mathbf{W}_{i,j}(\mathbf{x}; \sigma + \delta) \approx \left[1 + \delta \left(\frac{\|\mathbf{R}_i\mathbf{x} - \mathbf{R}_j\mathbf{x}\|^2}{\sigma^3} - \frac{d'_j(\sigma)}{d_j(\sigma)} \right) \right] \mathbf{W}_{i,j}(\mathbf{x}; \sigma).$$

Replacing $\delta = -\epsilon\sigma$, we obtain

$$\begin{aligned} \mathbf{W}_{i,j}((1 + \epsilon)\mathbf{x}; \sigma) &\approx \left[1 - \epsilon \left(\frac{\|\mathbf{R}_i\mathbf{x} - \mathbf{R}_j\mathbf{x}\|^2}{\sigma^2} - \sigma \frac{d'_j(\sigma)}{d_j(\sigma)} \right) \right] \mathbf{W}_{i,j}(\mathbf{x}; \sigma) \\ &= (1 - \epsilon \phi(\sigma)) \mathbf{W}_{i,j}(\mathbf{x}; \sigma). \end{aligned}$$

We can simplify further, but this is not necessary, since $\phi(\sigma)$ does not depend on ϵ . For its part, $\phi(\sigma)$ behaves like n/σ^2 . To see this note that the first term in the definition of ϕ is at worst n/σ^2 , since $\|\mathbf{R}_i\mathbf{x} - \mathbf{R}_j\mathbf{x}\|^2$ is bounded by n , given that the pixel values are in the range $[0, 255]$.

Similarly, d is on the order of $n \cdot \exp(\|\mathbf{R}_i\mathbf{x} - \mathbf{R}_j\mathbf{x}\|^2/2\sigma^2)$, and its derivative d' is on the order of $n \cdot \|\mathbf{R}_i\mathbf{x} - \mathbf{R}_j\mathbf{x}\|^2 \cdot d/\sigma^3$. Consequently, the second term $\sigma \cdot d'/d$ behaves like

$\|\mathbf{R}_i \mathbf{x} - \mathbf{R}_j \mathbf{x}\|^2 / \sigma^2$ is also on the order of n / σ^2 . Therefore, choosing $\epsilon = 1/n$, for sufficiently large σ , the term $\epsilon \phi(\sigma)$ becomes negligible.

What we have shown is that the filter weights change very little as a result of the scaling $(1 + \epsilon)\mathbf{x}$ as long as ϵ is very small. Therefore, the NLM (and bilateral) filters are (almost exactly) 1-homogeneous, as we had hoped.

C.2. Tikhonov regularizer and Wiener filtering. We now turn to show that Tikhonov regularization obeys the homogeneity condition. In this case, the denoised image is the solution of

$$(72) \quad \hat{\mathbf{x}} = \underset{\mathbf{x}}{\text{Argmin}} \frac{1}{2\sigma^2} \|\mathbf{x} - \mathbf{y}\|_2^2 + \frac{\lambda}{2} \|\mathbf{B}\mathbf{x}\|_2^2,$$

where \mathbf{B} , for example, can be a discrete approximation of a derivative operator. The closed-form expression of the above minimization is given by

$$(73) \quad \begin{aligned} \hat{\mathbf{x}} = f(\mathbf{y}) &= (\mathbf{I} + \lambda\sigma^2 \mathbf{B}^T \mathbf{B})^{-1} \mathbf{y} \\ &= \mathbf{W}\mathbf{y}. \end{aligned}$$

When we feed this linear denoiser with the scaled image $(1 + \epsilon)\mathbf{y}$, it gives

$$(74) \quad f((1 + \epsilon)\mathbf{y}) = (\mathbf{I} + \lambda\sigma^2 \mathbf{B}^T \mathbf{B})^{-1} (1 + \epsilon)\mathbf{y},$$

which is trivially the same as $(1 + \epsilon)f(\mathbf{y})$.

A more challenging case is obtained when the denoiser adapts to the input such that when we apply $f((1 + \epsilon)\mathbf{y})$, the denoiser modifies λ to be $\lambda(1 + \epsilon)^2$. In what follows we return to (74), but this time with the modified λ , and study the behavior of the filter for infinitesimal change in ϵ , i.e.,

$$(75) \quad \lim_{\epsilon \rightarrow 0} \mathbf{W}(\epsilon) = \lim_{\epsilon \rightarrow 0} (\mathbf{I} + \lambda(1 + \epsilon)^2 \sigma^2 \mathbf{B}^T \mathbf{B})^{-1}.$$

By relying on the first-order Taylor expansion of the above expression and taking the limit $\epsilon \rightarrow 0$ we get the following:

$$(76) \quad \begin{aligned} \mathbf{W}(\epsilon) &\approx \mathbf{W}(0) - \epsilon \mathbf{W}(\epsilon) (\lambda(1 + \epsilon)^2 \sigma^2 \mathbf{B}^T \mathbf{B})^T \mathbf{W}(\epsilon)|_{\epsilon=0} \\ &= \mathbf{W} - \epsilon 2\lambda\sigma^2 \mathbf{W} \mathbf{B}^T \mathbf{B} \mathbf{W} \\ &= (\mathbf{I} - \epsilon 2\lambda\sigma^2 \mathbf{W} \mathbf{B}^T \mathbf{B}) \mathbf{W}, \end{aligned}$$

where $\mathbf{W} = \mathbf{W}(0)$. Therefore, we obtain that the adaptive filter $\mathbf{W}(\epsilon)$ changes linearly with ϵ . Moreover, $\|\sigma^2 \mathbf{W} \mathbf{B}^T \mathbf{B}\| < 1$, and thus if $\lambda \ll 1$, the second term becomes negligible (similar to what we have shown for the NLM filter). Therefore, we conclude that the image-adaptive Wiener filtering satisfies the homogeneity condition under mild conditions.

C.3. Patch-based denoising. We now turn to discuss state-of-the-art patch-based denoising algorithms. These methods clean the input noisy image by (i) breaking it into small overlapping patches, (ii) applying a local denoising step, and (iii) reconstructing the output

by averaging the denoised overlapping patches to form the final image. The wisdom of these algorithms relies on the choice of the local nonlinear prior. Moreover, in most cases, the denoising process can be divided into two parts; the first contains all the nonlinear decisions made by the prior, while the second is nothing but a linear filter that cleans the noisy patches followed by a patch-averaging step. Clearly, the latter fulfills the homogeneity condition due to its linearity. In what follows we argue that the nonlinear part of various popular denoisers is stable to an infinitesimal change in scale, thus leading to an overall satisfaction of the homogeneity property.

Gaussian mixture model. We start by EPLL, which can be considered as an iterative GMM denoising algorithm. The nonlinear part of the GMM prior is the choice of a (pretrained) Gaussian model that best fits the input noisy patch, and its linear part is the subsequent Wiener filtering step. Consider a Gaussian model for an n -dimensional patch $\mathbf{x}_i \sim N(\mathbf{0}, \mathbf{\Sigma}_k)$, where $\mathbf{\Sigma}_k \in \mathbf{R}^{n \times n}$ is the k th Gaussian, taken from the mixture. Following the derivations in [43], the MAP estimate is formulated by

$$(77) \quad \begin{aligned} \hat{\mathbf{x}}_i^k &= \underset{\mathbf{x}_i}{\text{Argmin}} \frac{1}{2} \|\mathbf{x}_i - \mathbf{y}_i\|_2^2 + \sigma^2 \mathbf{x}_i^T \mathbf{\Sigma}_k^{-1} \mathbf{x}_i \\ &= (\mathbf{I} + \sigma^2 \mathbf{\Sigma}_k^{-1})^{-1} \mathbf{y}_i \\ &= \mathbf{W}_i^k \mathbf{y}_i, \end{aligned}$$

which is nothing but the Wiener filter that cleans the i th noisy patch \mathbf{y}_i . The best model for the i th patch, k_i^* , is the one that maximizes the MAP over all the possible models, given by

$$(78) \quad k_i^* = \underset{k}{\text{Argmin}} \|\hat{\mathbf{x}}_i^k - \mathbf{y}_i\|_2^2 + \sigma^2 (\hat{\mathbf{x}}_i^k)^T \mathbf{\Sigma}_k^{-1} (\hat{\mathbf{x}}_i^k) + \sigma^2 \log |\mathbf{\Sigma}_k|.$$

By plugging (77) into the above we get

$$(79) \quad \begin{aligned} k_i^* &= \underset{k}{\text{Argmin}} \|(\mathbf{W}_i^k - \mathbf{I})\mathbf{y}_i\|_2^2 + \sigma^2 (\mathbf{W}_i^k \mathbf{y}_i)^T \mathbf{\Sigma}_k^{-1} (\mathbf{W}_i^k \mathbf{y}_i) + \sigma^2 \log |\mathbf{\Sigma}_k| \\ &= \underset{k}{\text{Argmin}} \Psi(\mathbf{y}, \sigma, \mathbf{\Sigma}_k). \end{aligned}$$

When we feed the denoiser with $(1 + \epsilon)\mathbf{y}_i$, the above can be written as

$$(80) \quad \begin{aligned} k_i^{*,\epsilon} &= \underset{k}{\text{Argmin}} (1 + \epsilon)^2 \|(\mathbf{W}_i^k - \mathbf{I})\mathbf{y}_i\|_2^2 + (1 + \epsilon)^2 \sigma^2 (\mathbf{W}_i^k \mathbf{y}_i)^T \mathbf{\Sigma}_k^{-1} (\mathbf{W}_i^k \mathbf{y}_i) + \sigma^2 \log |\mathbf{\Sigma}_k| \\ &= \underset{k}{\text{Argmin}} \|(\mathbf{W}_i^k - \mathbf{I})\mathbf{y}_i\|_2^2 + \sigma^2 (\mathbf{W}_i^k \mathbf{y}_i)^T \mathbf{\Sigma}_k^{-1} (\mathbf{W}_i^k \mathbf{y}_i) + \frac{\sigma^2}{(1 + \epsilon)^2} \log |\mathbf{\Sigma}_k|. \end{aligned}$$

By relying on the relation $1/(1 + \epsilon)^2 \approx 1 - 2\epsilon$ we further simplify the above and obtain

$$(81) \quad \begin{aligned} k_i^{*,\epsilon} &= \underset{k}{\text{Argmin}} \|(\mathbf{W}_i^k - \mathbf{I})\mathbf{y}_i\|_2^2 + \sigma^2 (\mathbf{W}_i^k \mathbf{y}_i)^T \mathbf{\Sigma}_k^{-1} (\mathbf{W}_i^k \mathbf{y}_i) + \sigma^2 \log |\mathbf{\Sigma}_k| - 2\epsilon \sigma^2 \log |\mathbf{\Sigma}_k| \\ &= \underset{k}{\text{Argmin}} \Psi(\mathbf{y}, \sigma, \mathbf{\Sigma}_k) - 2\epsilon \sigma^2 \log |\mathbf{\Sigma}_k|. \end{aligned}$$

Now we turn to compare (79) to the one derived above and get the following condition on ϵ that guarantees that $k_i^{*,\epsilon} = k_i^*$:

$$(82) \quad \Psi(\mathbf{y}, \sigma, \mathbf{\Sigma}_k) - 2\epsilon \sigma^2 \log |\mathbf{\Sigma}_k| > \Psi(\mathbf{y}, \sigma, \mathbf{\Sigma}_{k^*}) - 2\epsilon \sigma^2 \log |\mathbf{\Sigma}_{k^*}|.$$

Since $\Psi(\mathbf{y}, \sigma, \Sigma_k) > \Psi(\mathbf{y}, \sigma, \Sigma_{k^*})$, one can always choose $\epsilon \rightarrow 0$, which will keep this inequality intact. In an extremely rare case, when $\Psi(\mathbf{y}, \sigma, \Sigma_k) = \Psi(\mathbf{y}, \sigma, \Sigma_{k^*})$, we can modify the GMM denoiser and propose a simple rule for choosing the model that has a smaller $\log |\Sigma_k|$ term, ensuring that $k_i^{*,\epsilon} = k_i^*$. To conclude, the nonlinear part of the GMM denoising algorithm is stable to small scaling of the input.

Sparsity-inspired denoisers: K-SVD. Given a dictionary, the nonlinear mechanism of K-SVD is the orthogonal matching pursuit (OMP) [61] algorithm, which estimates the sparse representation of an input noisy patch. This is a greedy method, aiming to approximate the solution of

$$(83) \quad \hat{\alpha}_i = \underset{\alpha_i}{\operatorname{Argmin}} \|\alpha_i\|_0 \quad \text{s.t.} \quad \|\mathbf{y}_i - \mathbf{D}\alpha_i\|_2^2 \leq n\sigma^2,$$

where n is the size of the patch, and $\mathbf{D} \in \mathbf{R}^{n \times m}$ is a (possibly redundant, $m > n$, and nonorthogonal) dictionary. At each step, denoted by k , OMP picks a new atom (a column from \mathbf{D}) that minimizes the residual. Formally, the rule for choosing the first atom d_{j_1} can be written as

$$(84) \quad j_1 = \underset{j}{\operatorname{Argmax}} |d_j^T \mathbf{y}_i|,$$

while in the t th step of the OMP, it is the one that maximizes

$$(85) \quad j_t = \underset{j}{\operatorname{Argmax}} |d_j^T \mathbf{r}_i^t|,$$

where $\mathbf{r}_i^t = \mathbf{y}_i - \mathbf{D}_{S_i^t} \alpha_i^t$ is the residual. We denote by S_i^t the set of chosen atoms, obtained in the previous steps, and by $\mathbf{D}_{S_i^t} \in \mathbf{R}^{n \times |S_i^t|}$ the corresponding dictionary—a matrix having the chosen atoms as its columns. This expression can be further simplified by relying on the fact that the representation is the outcome of a least-squares solution, given by

$$(86) \quad \alpha_i^t = \left(\mathbf{D}_{S_i^t}^T \mathbf{D}_{S_i^t} \right)^{-1} \mathbf{D}_{S_i^t}^T \mathbf{y}_i.$$

Substituting the above in (85) results in

$$(87) \quad j_t = \underset{j}{\operatorname{Argmax}} |d_j^T \mathbf{r}_i^t| = \underset{j}{\operatorname{Argmax}} \left| d_j^T \left(\mathbf{I} - \left(\mathbf{D}_{S_i^t}^T \mathbf{D}_{S_i^t} \right)^{-1} \mathbf{D}_{S_i^t}^T \right) \mathbf{y}_i \right|.$$

This process is repeated until reaching the error constraint.

Trivially, scaling the input patch would not modify the result of (84) and (87). However, scaling the input may modify the stopping rule and thereby the number of atoms that the OMP picks. This will happen only when $\|\mathbf{y}_i - \mathbf{D}\alpha_i\|_2^2 = n\sigma^2$, which is an extremely rare case. Notice that given the final set of chosen atoms, S_i^* , the cleaned patch is given by $\hat{\mathbf{x}}_i = \mathbf{D}_{S_i^*} \left(\mathbf{D}_{S_i^*}^T \mathbf{D}_{S_i^*} \right)^{-1} \mathbf{D}_{S_i^*}^T \mathbf{y}_i$, which clearly satisfies the homogeneity condition. To conclude, we showed that the nonlinear part of OMP is stable to an infinitesimal change in scale, and since the cleaned patch is simply obtained by a linear projection onto the chosen atoms we have that with high probability OMP (and thus K-SVD) fulfills our hope for homogeneity.

Appendix D. The differentiability requirement. In the opening of section 3.1, we required the denoising engine $f(\mathbf{x})$ to be differentiable. Why? There are several benefits for having this behavior:

1. The directional derivative property, $\nabla_{\mathbf{x}}f(\mathbf{x})\mathbf{x} = f(\mathbf{x})$, which emerges from the homogeneity condition, becomes possible.
2. The passivity condition, which refers to the spectral radius of $\nabla_{\mathbf{x}}f(\mathbf{x})$, stands on solid ground.
3. The convergence of the fixed-point algorithm discussed in section 4.2 is guaranteed.
4. The convexity of the regularization term $\rho_L(\mathbf{x})$, discussed in section 5.1 and emerging from Condition 2, becomes possible.

All these are good reasons for demanding differentiability of the denoiser $f(\mathbf{x})$, and yet the question is whether it is too limiting or whether this requirement could be circumvented.

Observe that point 1 above could rely on the availability of a far weaker requirement of the availability of all directional derivatives of the form $\nabla_{\mathbf{x}}f(\mathbf{x})\mathbf{x}$. Indeed, the passivity mentioned in point 2 could also be posed in terms of a directional derivative, replacing (15) by the somewhat weaker requirement

$$\forall \mathbf{x} \neq 0, \quad \frac{\mathbf{x}^T \nabla_{\mathbf{x}}f(\mathbf{x})\mathbf{x}}{\mathbf{x}^T \mathbf{x}} \leq 1,$$

under the assumption that $\nabla_{\mathbf{x}}f(\mathbf{x})$ is symmetric.

The question we leave as open at this stage is whether points 3 and 4 above (convergence of the fixed-point algorithm and the convexity of our regularization) could rely on the existence of all directional derivatives. At worst, if indeed the denoising engine is not differentiable but has all its directional derivatives, points 3 and 4 are lost, and the behavior of the proposed algorithm is not clear from a theoretical standpoint.

A different and more practical question is whether a differentiability assumption on $f(\mathbf{x})$ is feasible in existing algorithms. NLM and the Bilateral filter are clearly differentiable. TNRD is differentiable as well, since its nonlinearities are smooth (linear combination of Gaussian RBFs). EPLL, BM3D, and K-SVD are more challenging, since their nonlinear parts include sharp decisions. Each of these methods could be ϵ -modified to have a fuzzy decision, thus rendering all of them differentiable with hardly any change in their actual behavior.

REFERENCES

- [1] A. BUADES, B. COLL, AND J.M. MOREL, *A non-local algorithm for image denoising*, in CVPR, 2005.
- [2] C. KERVRANN AND J. BOULANGER, *Optimal spatial adaptation for patch-based image denoising*, IEEE Trans. Image Process., 15 (2006), pp. 2866–2878.
- [3] T. TADZIZEN, *Principal neighborhood dictionaries for nonlocal means image denoising*, IEEE Trans. Image Process., 18 (2009), pp. 2649–2660.
- [4] M. ELAD AND M. AHARON, *Image denoising via sparse and redundant representations over learned dictionaries*, IEEE Trans. Image Process., 15 (2006), pp. 3736–3745.
- [5] K. DABOV, A. FOI, V. KATKOVNIK, AND K. EGIAZARIAN, *Image denoising by sparse 3-D transform-domain collaborative filtering*, IEEE Trans. Image Process., 16 (2007), pp. 2080–2095.
- [6] H. TAKEDA, S. FARSIU, AND P. MILANFAR, *Kernel regression for image processing and reconstruction*, IEEE Trans. Image Process., 16 (2007), pp. 349–366.

- [7] J. MAIRAL, F. BACH, J. PONCE, G. SAPIRO, AND A. ZISSERMAN, *Non-local sparse models for image restoration*, in ICCV, 2009.
- [8] W. DONG, X. LI, L. ZHANG, AND G. SHI, *Sparsity-based image denoising via dictionary learning and structural clustering*, in CVPR, 2011.
- [9] W. DONG, L. ZHANG, G. SHI, AND X. LI, *Nonlocally centralized sparse representation for image restoration*, IEEE Trans. Image Process., 22 (2013), pp. 1620–1630.
- [10] P. CHATTERJEE AND P. MILANFAR, *Patch-based near-optimal image denoising*, IEEE Trans. Image Process., 21 (2012), pp. 1635–1649.
- [11] H.C. BURGER, C.J. SCHULER, AND S. HARMELING, *Image denoising: Can plain neural networks compete with BM3D?*, in CVPR, 2012.
- [12] M. LEBRUN, M. COLOM, A. BUADES, AND J.-M. MOREL, *Secrets of image denoising cuisine*, Acta Numer., 21 (2012), pp. 475–576.
- [13] M. LEBRUN, A. BUADES, AND J.-M. MOREL, *Implementation of the “non-local Bayes” (NL-Bayes) image denoising algorithm*, IPOL, 3 (2013), pp. 1–42.
- [14] C. KNAUS AND M. ZWICKER, *Dual-domain image denoising*, in ICIP, 2013.
- [15] H. TALEBI AND P. MILANFAR, *Global image denoising*, IEEE Trans. Image Process., 23 (2014), pp. 755–768.
- [16] S. GU, L. ZHANG, W. ZUO, AND X. FENG, *Weighted nuclear norm minimization with application to image denoising*, in CVPR, 2014.
- [17] W. DONG, G. SHI, Y. MA, AND X. LI, *Image restoration via simultaneous sparse coding: Where structured sparsity meets Gaussian scale mixture*, Int. J. Comput. Vis., 114 (2015), pp. 217–232.
- [18] N. PIERAZZO, M. RAIS, J.-M. MOREL, AND G. FACCIOLO, *DA3D: Fast and data adaptive dual domain denoising*, in ICIP, 2015.
- [19] P. MILANFAR, *Symmetrizing smoothing filters*, SIAM J. Imaging Sci., 6 (2013), pp. 263–284, <https://doi.org/10.1137/120875843>.
- [20] Y. CHEN AND T. POCK, *Trainable Nonlinear Reaction Diffusion: A Flexible Framework for Fast and Effective Image Restoration*, preprint, <https://arxiv.org/abs/1508.02848>, 2015.
- [21] P. CHATTERJEE AND P. MILANFAR, *Is denoising dead?*, IEEE Trans. Image Process., 19 (2010), pp. 895–911.
- [22] A. LEVIN AND B. NADLER, *Natural image denoising: Optimality and inherent bounds*, in CVPR, 2011.
- [23] A. LEVIN, B. NADLER, F. DURAND, AND W.T. FREEMAN, *Patch complexity, finite pixel correlations and optimal denoising*, in ECCV, 2012.
- [24] M. PROTTER, M. ELAD, H. TAKEDA, AND P. MILANFAR, *Generalizing the nonlocal-means to superresolution reconstruction*, IEEE Trans. Image Process., 18 (2009), pp. 36–51.
- [25] S. RAVISHANKAR AND Y. BRESLER, *MR image reconstruction from highly undersampled k-space data by dictionary learning*, IEEE Trans. Med. Imag., 30 (2011), pp. 1028–1041.
- [26] M. AKCAKAYA, T.A. BASHA, B. GODDU, L.A. GOEPFERT, K.V. KISSINGER, V. TAROKH, W.J. MANNING, AND R. NEZAFAT, *Low-dimensional structure self-learning and thresholding: Regularization beyond compressed sensing for MRI reconstruction*, Magn. Reson. Med., 66 (2011), pp. 756–767.
- [27] A. DANIELYAN, V. KATKOVNIK, AND K. EGIAZARIAN, *BM3D frames and variational image deblurring*, IEEE Trans. Image Process., 21 (2012), pp. 1715–1728.
- [28] W. DONG, L. ZHANG, R. LUKAC, AND G. SHI, *Sparse representation based image interpolation with nonlocal autoregressive modeling*, IEEE Trans. Image Process., 22 (2013), pp. 1382–1394.
- [29] C.A. METZLER, A. MALEKI, AND R.G. BARANIUK, *Optimal recovery from compressive measurements via denoising-based approximate message passing*, in SAMPTA, 2015.
- [30] S.V. VENKATAKRISHNAN, C.A. BOUMAN, AND B. WOHLBERG, *Plug-and-play priors for model based reconstruction*, in GlobalSIP, 2013.
- [31] S.H. CHAN, *Algorithm-Induced Prior for Image Restoration*, preprint, <https://arxiv.org/abs/1602.00715>, 2016.
- [32] S.H. CHAN, X. WANG, AND O.A. ELGENDY, *Plug-and-Play ADMM for Image Restoration: Fixed Point Convergence and Applications*, preprint, preprint, <https://arxiv.org/abs/1605.01710>, 2016.
- [33] S. SREEHARI, S.V. VENKATAKRISHNAN, B. WOHLBERG, L.F. DRUMMY, J.P. SIMMONS, AND C.A. BOUMAN, *Plug-and-Play Priors for Bright Field Electron Tomography and Sparse Interpolation*, preprint, <https://arxiv.org/abs/1512.07331>, 2015.

- [34] Y. DAR, A.M. BRUCKSTEIN, M. ELAD, AND R. GIRYES, *Postprocessing of compressed images via sequential denoising*, IEEE Trans. Image Process., 25 (2016), pp. 3044–3058.
- [35] M.U. SADIQ, J.P. SIMMONS, AND C.A. BOUMAN, *Model based image reconstruction with physics based priors*, in ICIP, 2016.
- [36] A.M. TEODORO, J.M. BIOCAS-DIAS, AND M.A.T. FIGUEIREDO, *Image restoration with locally selected class-adapted models*, in ICIP, 2016.
- [37] A. BRIFMAN, Y. ROMANO, AND M. ELAD, *Turning a denoiser into a super-resolver using plug and play priors*, in ICIP, 2016.
- [38] A.M. TEODORO, J.M. BIOCAS-DIAS, AND M.A.T. FIGUEIREDO, *Hyperspectral Sharpening Using Scene-Adapted Gaussian Mixture Priors*, preprint, <https://arxiv.org/abs/1702.02445>, 2017.
- [39] R.L. LEGENDIJK AND J. BIEMOND, *Iterative Identification and Restoration of Images*, Springer, 1990.
- [40] L.I. RUDIN, S. OSHER, AND E. FATEMI, *Nonlinear total variation based noise removal algorithms*, Phys. D, 60 (1992), pp. 259–268.
- [41] S. MALLAT, *A Wavelet Tour of Signal Processing*, Academic Press, 1999.
- [42] D. ZORAN AND Y. WEISS, *From learning models of natural image patches to whole image restoration*, in CVPR, 2011.
- [43] G. YU, G. SAPIRO, AND S. MALLAT, *Solving inverse problems with piecewise linear estimators: From Gaussian mixture models to structured sparsity*, IEEE Trans. Image Process., 21 (2012), pp. 2481–2499.
- [44] A.M. BRUCKSTEIN, D.L. DONOHO, AND M. ELAD, *From sparse solutions of systems of equations to sparse modeling of signals and images*, SIAM Rev., 51 (2009), pp. 34–81, <https://doi.org/10.1137/060657704>.
- [45] A. VEDALDI AND A. ZISSERMAN, *Efficient additive kernels via explicit feature maps*, IEEE Trans. Pattern Anal. Mach. Intell., 34 (2012), pp. 480–492.
- [46] V. POPYAN, Y. ROMANO, AND M. ELAD, *Convolutional Neural Networks Analyzed via Convolutional Sparse Coding*, preprint, <https://arxiv.org/abs/1607.08194>, 2016.
- [47] A. ELMOATAZ, O. LEZORAY, AND S. BOUGLEUX, *Nonlocal discrete regularization on weighted graphs: A framework for image and manifold processing*, IEEE Trans. Image Process., 17 (2008), pp. 1047–1060.
- [48] A.D. SZLAM, M. MAGGIONI, AND R.R. COIFMAN, *Regularization on graphs with function-adapted diffusion processes*, J. Mach. Learn. Res., 9 (2008), pp. 1711–1739.
- [49] S. BOUGLEUX, A. ELMOATAZ, AND M. MELKEMI, *Local and nonlocal discrete regularization on weighted graphs for image and mesh processing*, Int. J. Comput. Vis., 84 (2009), pp. 220–236.
- [50] G. PEYRÉ, S. BOUGLEUX, AND L.D. COHEN, *Non-local regularization of inverse problems*, Inverse Probl. Imaging, 5 (2011), pp. 511–530.
- [51] P. MILANFAR, *A tour of modern image filtering: New insights and methods, both practical and theoretical*, IEEE Signal Process. Mag., 30 (2013), pp. 106–128.
- [52] F.G. MEYER AND X. SHEN, *Perturbation of the eigenvectors of the graph Laplacian: Application to image denoising*, Appl. Comput. Harmon. Anal., 36 (2014), pp. 326–334.
- [53] X. LIU, D. ZHAI, D. ZHAO, G. ZHAI, AND W. GAO, *Progressive image denoising through hybrid graph Laplacian regularization: A unified framework*, IEEE Trans. Image Process., 23 (2014), pp. 1491–1503.
- [54] S.M. HAQUE, G. PAI, AND V.M. GOVINDU, *Symmetric smoothing filters from global consistency constraints*, IEEE Trans. Image Process., 24 (2014), pp. 1536–1548.
- [55] A. KHERADMAND AND P. MILANFAR, *A general framework for regularized, similarity-based image restoration*, IEEE Trans. Image Process., 23 (2014), pp. 5136–5151.
- [56] Y. ROMANO AND M. ELAD, *Boosting of image denoising algorithms*, SIAM J. Imaging Sci., 8 (2015), pp. 1187–1219, <https://doi.org/10.1137/140990978>.
- [57] S. BOYD, N. PARIKH, E. CHU, B. PELEATO, AND J. ECKSTEIN, *Distributed optimization and statistical learning via the alternating direction method of multipliers*, Found. Trends Mach. Learn., 3 (2011), pp. 1–122.
- [58] M. AFONSO, J. BIOCAS-DIAS, AND M.A.T. FIGUEIREDO, *Fast image recovery using variable splitting and constrained optimization*, IEEE Trans. Image Process., 19 (2010), pp. 2345–2356.
- [59] M. AFONSO, J. BIOCAS-DIAS, AND M.A.T. FIGUEIREDO, *An augmented Lagrangian approach to the constrained optimization formulation of imaging inverse problems*, IEEE Trans. Image Process., 20 (2011), pp. 681–695.
- [60] N. PARIKH AND S. BOYD, *Proximal algorithms*, Found. Trends Optim., 1 (2014), pp. 123–231.

- [61] Y.C. PATI, R. REZAIIFAR, AND P.S. KRISHNAPRASAD, *Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition*, in IEEE Asilomar Conf. Signals, Syst. Comput., 1993, pp. 40–44.
- [62] E. CANDÉS AND T. TAO, *The Dantzig selector: Statistical estimation when p is much larger than n* , Ann. Statist., 35 (2007), pp. 2313–2351.
- [63] Y. ROMANO AND M. ELAD, *Improving K -SVD denoising by post-processing its method-noise*, in ICIP, 2013.
- [64] H. TALEBI, X. ZHU, AND P. MILANFAR, *How to SAIF-ly boost denoising performance*, IEEE Trans. Image Process., 22 (2013), pp. 1470–1485.
- [65] T.M. APOSTOL, *Calculus*, 2nd ed., Vol. 1, Blaisdell, Waltham, MA, 1967.
- [66] S. BOYD AND L. VANDENBERGHE, *Convex Optimization*, Cambridge University Press, 2004.
- [67] M. ELAD, *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing*, 1st ed., Springer, 2010.
- [68] G.H. GOLUB AND C.F. VAN LOAN, *Matrix Computations*, 3rd ed., John Hopkins University Press, Baltimore, 1996.
- [69] M. ELAD, B. MATALON, AND M. ZIBULEVSKY, *Coordinate and subspace optimization methods for linear least squares with non-quadratic regularization*, Appl. Comput. Harmon. Anal., 23 (2007), pp. 346–367.
- [70] M. DROSKE AND A. BERTOZZI, *Higher-order feature-preserving geometric regularization*, SIAM J. Imaging Sci., 3 (2010), pp. 21–51, <https://doi.org/10.1137/090751694>.
- [71] A. BECK AND M. TEOULLE, *Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems*, IEEE Trans. Image Process., 18 (2009), pp. 2419–2434.
- [72] W. DONG, L. ZHANG, G. SHI, AND X. WU, *Image deblurring and super-resolution by adaptive sparse domain selection and adaptive regularization*, IEEE Trans. Image Process., 20 (2011), pp. 1838–1857.