

CLEAR: Covariant LEAsT-Square Refitting with Applications to Image Restoration*

Charles-Alban Deledalle[†], Nicolas Papadakis[†], Joseph Salmon[‡], and Samuel Vaiter[§]

Abstract. In this paper, we propose a new framework to remove parts of the systematic errors affecting popular restoration algorithms, with a special focus on image processing tasks. Generalizing ideas that emerged for ℓ_1 regularization, we develop an approach refitting the results of standard methods toward the input data. Total variation regularizations and nonlocal means are special cases of interest. We identify important covariant information that should be preserved by the refitting method and emphasize the importance of preserving the Jacobian (w.r.t. the observed signal) of the original estimator. Then, we provide an approach that has a “twicing” flavor and allows refitting the restored signal by adding back a local affine transformation of the residual term. We illustrate the benefits of our method on numerical simulations for image restoration tasks.

Key words. inverse problems, variational methods, refitting, twicing, boosting, debiasing

AMS subject classifications. 49N45, 65K10, 68U10

DOI. 10.1137/16M1080318

1. Introduction. Restoring an image from its single noisy and incomplete observation necessarily requires enforcing regularity or a model *prior* on the targeted properties of the sought solution. Regularity properties such as sparsity and gradient sparsity of an image are difficult to enforce in general and notably lead to combinatorial and nonconvex problems. When one is willing to guarantee such kinds of features, convex relaxation is a popular path. This is typically done using the ℓ_1 norm instead of the ℓ_0 pseudonorm, as for the Lasso [45] or the total variation [39]. Nevertheless, such relaxations are well known to create solutions with a larger bias.

Typically, for the Lasso, using the ℓ_1 convex relaxation of the ℓ_0 pseudonorm leads large coefficients to be shrunk toward zero.

For the total variation, the same relaxation on the jump amplitudes induces a loss of contrast in the recovered signal; see Figure 1(a) for an illustration in this case.

In the Lasso case, a well-known refitting scheme consists in performing a posteriori a least-square re-estimation of the nonzero coefficients of the solution. This post-refitting technique

*Received by the editors June 16, 2016; accepted for publication (in revised form) December 2, 2016; published electronically March 2, 2017.

<http://www.siam.org/journals/siims/10-1/M108031.html>

Funding: The work of the authors was supported by the GDR 720 ISIS (Information, Signal, Image et viSion), and by the French State, managed by the French National Research Agency (ANR) in the frame of the “Investments for the future” Programme IdEx Bordeaux - CPU (ANR-10-IDEX-03-02).

[†]IMB, CNRS, Université de Bordeaux, Bordeaux INP, F-33405 Talence, France (charles-alban.deledalle@math.u-bordeaux.fr, nicolas.papadakis@math.u-bordeaux.fr).

[‡]LTCl, CNRS, Télécom ParisTech, Université Paris-Saclay, 75013 Paris, France (joseph.salmon@telecom-paristech.fr).

[§]IMB, CNRS, Université de Bourgogne, 21078 Dijon, France (samuel.vaiter@u-bourgogne.fr).




Figure 1. (a) Solutions of 1D-TV and our refitting on a noisy signal. (b) Illustration of the invariant refitting in a denoising problem of dimension $p = 3$. The gray surface is a piecewise affine mapping that models the evolution of \hat{x} in an extended neighborhood of y . The light red affine plane is the model subspace, i.e., the set of images sharing the same jumps as those of the solution $\hat{x}(y)$. The red triangle is the restriction of the model subspace to images that can be produced by TV. Finally, the pink dot represents the refitting $\mathcal{R}_x^{inv}(y)$ as the orthogonal projection of y on $\mathcal{M}_{\hat{x}}(y)$.

became popular under various names in the literature, such as hybrid Lasso [20], Lasso–Gauss [37], OLS post-Lasso [1], and debiased Lasso (see [29, 1] for extensive details on the subject). For the anisotropic total variation (aniso-TV), the same post-refitting approach can be performed to re-estimate the amplitudes of the jumps, provided their locations have been correctly identified.

In this paper, we introduce a generalization of this refitting technique that aims at re-enhancing the estimation toward the data without altering the desired properties imposed by the model prior. The underlying reason is that if the user chooses an estimator for which theoretical guarantees have been proven (as for the Lasso and total variation), refitting should preserve these guarantees. To that end, we introduce the *Covariant LEAs-t-square Refitting* (CLEAR). Though this method was originally elaborated with ℓ_1 analysis problems in mind, it has the ability to generalize to a wider family, while in simple cases such as the Lasso or the aniso-TV, it recovers the classical post-refitting solution described earlier. For instance, our methodology successfully applies to Tikhonov regularization [47], isotropic total variation (iso-TV), nonlocal means [3], block matching 3D (BM3D) [11], and dual domain image denoising (DDID) [27]. In common variational contexts, e.g., $\ell_1 - \ell_2$ analysis [21] (encompassing Lasso, group Lasso [30, 55, 33], aniso-TV, and iso-TV), we show that our refitting technique can be performed with a complexity overload of about twice that of the original algorithm. In other cases, we introduce a scheme requiring about three times the complexity of the original algorithm.

While our covariant refitting technique recovers the classical post-refitting solution in specific cases, the proposed algorithm offers more stable solutions. Unlike the Lasso post-refitting technique, ours does not require identifying a posteriori the support of the solution,

i.e., the set of nonzero coefficients. In the same vein, it does not require identifying the jump locations of the aniso-TV solution. Since the Lasso and the aniso-TV solutions are usually obtained through iterative algorithms, stopped at a prescribed convergence accuracy, the support or jump numerical identification might be imprecise (all the more so for ill-posed problems). Such erroneous support identifications lead to results that strongly deviate from the sought refitting. Our covariant refitting jointly estimates the re-enhanced solution during the iterations of the original algorithm and, as a by-product, produces more robust solutions.

This work follows a preliminary study [15] that attempted to suppress the bias emerging from the choice of the method (e.g., ℓ_1 relaxation), while leaving unchanged the bias due to the choice of the model (e.g., sparsity). While the approach from [15]—hereafter referred to as *invariant refitting*—provides interesting results, it is however limited to a class of restoration algorithms that satisfy restrictive local properties.

In particular, invariant refitting cannot handle iso-TV. In this case, invariant refitting is unsatisfactory as it removes some desired aspects enforced by the prior, such as smoothness, and suffers from a significant increase of variance. A simple illustration of this phenomenon for iso-TV is provided in Figure 2(d) where artificial oscillations are wrongly amplified near the boundaries.

While covariant and invariant refitting both correspond to the least-square post-refitting step in the case of aniso-TV, the two techniques do not match for iso-TV. Indeed, CLEAR outputs a more relevant solution than the one from invariant refitting. Figure 2(e) shows the benefit of our proposed solution with respect to (w.r.t.) the (naive) invariant refitting displayed in Figure 2(d).

It is worth mentioning that covariant refitting is also strongly related to boosting methods reinjecting useful information remaining in the residual (i.e., the map of the pointwise difference between the original signal and its prediction). Such approaches can be traced back to *twicing* [48] and have recently been thoroughly investigated: boosting [4], Bregman iterations and nonlinear inverse scale spaces [34, 5, 53, 35], unsharp residual iteration [8], SAIF-boosting [31, 44], ideal spectral filtering in the analysis sense [24], and SOS-boosting [38] are some of the most popular ones. Most of these methods can be performed iteratively, leading to a difficult choice for the number of steps to consider in practice. Our method has the noticeable advantage that it is by construction a two-step one. Iterating more would not be beneficial (see subsection 4.2). Unlike refitting, these later approaches aim at improving the overall image quality by authorizing the re-enhanced result to deviate strongly from the original biased solution. In particular, they do not recover the aforementioned post-refitting technique in the Lasso case. Our objective is not to guarantee the image quality to be improved but rather to generalize the refitting approach with the ultimate goal of reducing the bias while preserving the structure and the regularity of the original biased solution.

Interestingly, we have also realized that our scheme presents some similarities to the classical shrinking estimators introduced in [42], especially as presented in [23]. Indeed, the step performed by CLEAR is similar to a shrinkage step with a data-driven residual correction weight (later referred to as ρ in our approach; see Definition 14) when performing shrinkage as in [23, section 3.1].

Last but not least, it is well known that bias reduction is not always favorable in terms of mean square error (MSE) because of the so-called bias-variance trade-off. It is important to

highlight that a refitting procedure is expected to reinject part of the variance; therefore, it could lead to an increase of residual noise. Hence, the MSE is not expected to be improved by refitting techniques (unlike the aforementioned boosting-like methods that attempt to improve the MSE). We will show in our numerical experiments that refitting is in practice beneficial when the signal of interest fits well the model imposed by the prior. In other scenarios, when the model mismatches the sought signal, the original biased estimator remains favorable in terms of MSE. Refitting is nevertheless essential in the latter case for applications where the image intensities have a physical sense and critical decisions are taken from their values.

2. Background models and notation. We tackle the problem of estimating an unknown vector $x_0 \in \mathbb{R}^p$ from noisy and incomplete observations

$$(1) \quad y = \Phi x_0 + w,$$

where Φ is a linear operator from \mathbb{R}^p to \mathbb{R}^n and $w \in \mathbb{R}^n$ is the realization of a noisy random vector. This linear model is widely used in statistics and in imagery (e.g., for encoding degradations such as entrywise masking and convolution). Typically, the inverse problem associated to (1) is ill-posed, and one should add additional information to recover at least an approximation of x_0 from y . We denote by $\hat{x} : y \mapsto \hat{x}(y)$ the procedure that provides such an estimation of x_0 . In this paper, we consider a popular class of estimators relying on a variational formulation that involves a data fidelity term $F(x, y)$ and a regularizing term $G(x)$:¹

$$(2) \quad \hat{x}(y) \in \operatorname{argmin}_{x \in \mathbb{R}^p} F(x, y) + G(x).$$

Another kind of estimator can be defined as the output of an iterative algorithm $(k, y) \mapsto x^k$, e.g., a numerical solver. For a chosen iteration k , we define the final estimator $\hat{x}(y) = x^k$. Such a framework includes, for instance, proximal splitting methods, as well as discretization of partial differential equations, though we do not investigate this latter method in detail.

2.1. Notation. For a matrix M , M^+ is its Moore–Penrose pseudoinverse. For a (closed convex) set C , Π_C is the Euclidean projection over C , and ι_C is its indicator function defined by $\iota_C(u) = 0$ if $u \in C$ and $+\infty$ otherwise.

For any integer $d \in \mathbb{N}^*$, we write $[d]$ for the set $\{1, \dots, d\}$. For any subset \mathcal{I} of $[d]$, its complement in $[d]$ is written \mathcal{I}^c . For any vector v , $v_{\mathcal{I}} \in \mathbb{R}^{|\mathcal{I}|}$ is the subvector whose elements are indexed by $\mathcal{I} \subset \mathbb{N}$, and $|\mathcal{I}|$ is its cardinality. For any matrix M , $M_{\mathcal{I}}$ is the submatrix whose columns are indexed by \mathcal{I} . We denote, respectively, by $\operatorname{Im}[A]$ and $\operatorname{Ker}[A]$ the image space and the kernel space of an operator A . For any vector $x \in \mathbb{R}^d$ and $q \in [1, +\infty]$ we denote by $\|x\|_q$ its standard ℓ_q norm and by $\|x\|_0$ the number of nonzero entries of x , i.e., $\|x\|_0 = |\{i \in [d] : x_i \neq 0\}|$.

2.2. Main investigated estimators. Here, we provide several canonical examples of estimators of the form (2) that help us illustrate our methodology.

¹Often the solution of (2) is nonunique, but for simplicity we only consider a selection of such solutions, and we assume that the selected path $\hat{x} : y \mapsto \hat{x}(y)$ is differentiable almost everywhere.

Affine constrained least-squares, used when x_0 belongs to the affine subspace $C = b + \text{Im}[A]$ with $b \in \mathbb{R}^p$ and $A \in \mathbb{R}^{p \times n}$, are a particular case of (2) where $F(x, y) = \frac{1}{2}\|y - \Phi x\|_2^2$ and $G(x) = \iota_C(x)$. The solution of the minimum Euclidean norm is unique and is given by $\hat{x}(y) = b + A(\Phi A)^+(y - \Phi b)$ (see [Appendix A.1](#)).

Tikhonov regularization [47] (or Ridge regression [26]), used when $\|\Gamma x_0\|_2$ is small with $\Gamma \in \mathbb{R}^{m \times p}$, is an instance of (2) where $F(x, y) = \frac{1}{2}\|y - \Phi x\|_2^2$ and $G(x) = \frac{\lambda}{2}\|\Gamma x\|_2^2$ for some parameters $\lambda > 0$. Provided $\text{Ker } \Phi \cap \text{Ker } \Gamma = \{0\}$, $\hat{x}(y)$ exists and is uniquely defined as $\hat{x}(y) = (\Phi^\top \Phi + \lambda \Gamma^\top \Gamma)^{-1} \Phi^\top y$ (see subsection [A.2](#)).

Hard-thresholding [17], used when $\Phi = \text{Id}$ and x_0 is sparse, is a solution of (2) where $F(x, y) = \frac{1}{2}\|y - x\|_2^2$ and $G(x) = \frac{\lambda^2}{2}\|x\|_0$ for some parameter $\lambda > 0$. The hard-thresholding operation writes as $\hat{x}(y)_{\mathcal{I}} = y_{\mathcal{I}}$ and $\hat{x}(y)_{\mathcal{I}^c} = 0$, where $\mathcal{I} = \{i \in [p] : |y_i| > \lambda\}$ (see subsection [A.3](#)).

Soft-thresholding [17], used when $\Phi = \text{Id}$ and x_0 is sparse, is a solution of (2) where $F(x, y) = \frac{1}{2}\|y - x\|_2^2$ and $G(x) = \lambda\|x\|_1$ for some parameter $\lambda > 0$. The soft-thresholding operation writes as $\hat{x}(y)_{\mathcal{I}} = y_{\mathcal{I}} - \lambda \text{sign}(y_{\mathcal{I}})$ and $\hat{x}(y)_{\mathcal{I}^c} = 0$, where \mathcal{I} is defined as for hard-thresholding (see subsection [A.4](#)).

ℓ_1 *synthesis* (or Lasso [45, 17]), used when x_0 is sparse, is a solution of (2) where $F(x, y) = \frac{1}{2}\|\Phi y - x\|_2^2$ and $G(x) = \lambda\|x\|_1$ for some parameter $\lambda > 0$. Provided the solution is unique (see, for instance, [46]), it reads

$$(3) \quad \hat{x}(y)_{\mathcal{I}} = (\Phi_{\mathcal{I}})^+ y_{\mathcal{I}} - \lambda((\Phi_{\mathcal{I}})^\top \Phi_{\mathcal{I}})^{-1} s_{\mathcal{I}} \quad \text{and} \quad \hat{x}(y)_{\mathcal{I}^c} = 0$$

for almost all y and where $\mathcal{I} = \text{supp}(\hat{x}(y)) = \{i \in [p] : \hat{x}(y)_i \neq 0\}$ is the support of $\hat{x}(y)$, and $s_{\mathcal{I}} = \text{sign}(\hat{x}(y)_{\mathcal{I}})$.

ℓ_1 *analysis*, used when Γx_0 is sparse with $\Gamma \in \mathbb{R}^{m \times p}$, is a solution of (2) where

$$(4) \quad F(x, y) = \frac{1}{2}\|\Phi x - y\|_2^2 \quad \text{and} \quad G(x) = \lambda\|\Gamma x\|_1 \quad \text{for some } \lambda > 0.$$

Provided $\text{Ker } \Phi \cap \text{Ker } \Gamma = \{0\}$, there exists a solution given implicitly (see [49]) as

$$(5) \quad \hat{x}(y) = U(\Phi U)^+ y - \lambda U(U^\top \Phi^\top \Phi U)^{-1} U^\top (\Gamma^\top)_{\mathcal{I}} s_{\mathcal{I}}$$

for almost all y and where $\mathcal{I} = \text{supp}(\Gamma \hat{x}(y)) = \{i \in [m] : (\Gamma \hat{x}(y))_i \neq 0\}$ is called the Γ -support of the solution, $s_{\mathcal{I}} = \text{sign}((\Gamma \hat{x}(y))_{\mathcal{I}})$, U is a matrix whose columns form a basis of $\text{Ker}[\Gamma_{\mathcal{I}^c}]$, and ΦU has full column rank.

Aniso-TV [39] is an instance of (4) where $x_0 \in \mathbb{R}^p$ is identified to a b -dimensional signal, for which $\Gamma = \nabla : \mathbb{R}^p \rightarrow \mathbb{R}^{p \times b}$ is the discrete gradient operator and $\|\nabla x\|_1 = \sum_{i=1}^p \|(\nabla x)_i\|_1$. Aniso-TV promotes piecewise constant solutions with large constant regions and few sharp transitions. Here, \mathcal{I} is the set of indexes where the solution has discontinuities. The ℓ_1 norm of the gradient field induces an anisotropic effect by favoring the jumps to be aligned with the canonical directions (so it favors squared-like structures rather than rounded ones).

$\ell_1 - \ell_2$ *analysis* [30, 55, 33], used when Γx_0 is block sparse with $\Gamma \in \mathbb{R}^p \mapsto \mathbb{R}^{m \times b}$, is a solution of (2) with

$$(6) \quad F(x, y) = \frac{1}{2}\|\Phi x - y\|_2^2 \quad \text{and} \quad G(x) = \lambda\|\Gamma x\|_{1,2}$$




Figure 2. (a) A piecewise constant signal. (b) Its noisy version. (c) Solution of iso-TV with $\lambda = 10$ on the noisy signal. (d) Solution of the invariant refitting of iso-TV. (e) Solution of the covariant refitting of iso-TV. Red points indicate locations where the discrete gradient is nonzero.

for some parameter $\lambda > 0$ and where $\|\Gamma x\|_{1,2} = \sum_{i=1}^m \|(\Gamma x)_i\|_2$. Iso-TV [39] is a particular instance of (6) where $x_0 \in \mathbb{R}^p$ can be identified to a b -dimensional signal, for which $\Gamma = \nabla : \mathbb{R}^p \rightarrow \mathbb{R}^{p \times b}$ and $\|\nabla x\|_{1,2} = \sum_{i=1}^p \|(\nabla x)_i\|_2$. Like aniso-TV, it promotes solution with large constant regions, but some transition regions can be smooth (see, e.g., [6])—typically those with high curvature in the input image y ; see Figure 2(a)–(c). A major difference is that the $\ell_1 - \ell_2$ norm induces an isotropic effect by favoring rounded-like structures rather than squared ones.

(Blockwise) nonlocal means [3], used when the image $x_0 \in \mathbb{R}^p$ (with $p = p_1 \times p_2$) is composed of many redundant patterns, is the solution of (2) for $F(x, y) = \frac{1}{2} \sum_{i,j} w_{i,j} \|\mathcal{P}_i x - \mathcal{P}_j y\|_2^2$ and $G(x) = 0$, where \mathcal{P}_i is the linear operator extracting a patch (i.e., a small window) at pixel i of size $(2b+1)^2$. The index $i \in [p_1] \times [p_2]$ spans the image domain, and $j - i \in [-s, s]^2$ spans a limited search window. The weights $w_{i,j}$ are defined as $w_{i,j} = \varphi(\|\mathcal{P}_i y - \mathcal{P}_j y\|_2^2)$ where the kernel $\varphi : \mathbb{R}^+ \rightarrow [0, 1]$ is a decreasing function. Assuming periodical boundary conditions, its solution is given as

$$(7) \quad \hat{x}(y)_i = \frac{\sum_j \bar{w}_{i,j} y_j}{\sum_j \bar{w}_{i,j}} \quad \text{with} \quad \bar{w}_{i,j} = \sum_k w_{i-k, j-k},$$

where $k \in [-d, d]^2$ spans the patch domain (see subsection A.5).

2.3. Limitations. It is important to note that some of the previously introduced estimators are known to suffer from systematic drawbacks.

For instance, Lasso contracts large coefficients toward 0 by the shift $\lambda((\Phi_{\mathcal{I}})^\top \Phi_{\mathcal{I}})^{-1} s_{\mathcal{I}}$ whose expression simplifies to $\lambda \text{sign}(y_{\mathcal{I}})$ for soft-thresholding. In the same vein, ℓ_1 analysis contracts the signal toward $\text{Ker}[\Gamma]$ by the shift $\lambda U(U^\top \Phi^\top \Phi U)^{-1} U^\top (\Gamma^\top)_{\mathcal{I}} s_{\mathcal{I}}$. In the particular case of aniso-TV, this last quantity is the well-known systematic loss of contrast of total variation: a shift of intensity on each piece depending on its surroundings and the ratio between its perimeter and its area (see [43] for a thorough study for the one-dimensional (1D) case). A simple illustration is provided in Figure 1 where total variation is used for denoising a 1D piecewise constant signal in $[0, 192]$ and damaged by additive white Gaussian noise (AWGN) with a standard deviation $\sigma = 10$. Even though total variation has perfectly retrieved the support of ∇x_0 with one more extra jump, the intensities of some regions are systematically under- or overestimated. Iso-TV also suffers from a systematic loss of contrast, as illustrated in Figure 2(c).

In the following we investigate possible solutions to reduce such artifacts.

3. Invariant least-square refitting. As mentioned earlier, practitioners have realized that a systemic contraction affects estimators like Lasso and aniso-TV. In the Lasso case, a simple remedy (presented in the introduction) is to perform a posteriori a least-square refitting step of the nonzero coefficients, i.e., constrained to the support \mathcal{I} of the Lasso solution $\hat{x}(y)$, given by

$$(8) \quad \underset{x; \text{supp}(x) \subseteq \mathcal{I}}{\text{argmin}} \quad \frac{1}{2} \|\Phi x - y\|_2^2.$$

In this section, we present a refitting procedure, discussed in [15], that generalizes this approach to a broad family of estimators.

3.1. Refitting through model subspace least-squares. We investigate a refitting procedure well suited for estimators almost everywhere differentiable (a.e. differentiable). It relies on the notion of model subspace, which requires Jacobian matrix computations. From now on, we consider only estimators $y \mapsto \hat{x}(y)$ from \mathbb{R}^n to \mathbb{R}^p and a.e. differentiable.

Many estimation procedures rely on a structural prior of the data. Such structures include smoothness, sparsity, autosimilarity, or the fact that the signal is piecewise constant. Such priors can be captured by the following notion of model subspace.

Definition 1. The *model subspace* associated to an a.e. differentiable estimator \hat{x} is defined at almost all points $y \in \mathbb{R}^n$ by the affine subspace of \mathbb{R}^p

$$(9) \quad \mathcal{M}_{\hat{x}}(y) = \hat{x}(y) + \text{Im} [J_{\hat{x}}(y)],$$

where $J_{\hat{x}}(y) \in \mathbb{R}^{p \times n}$ is the Jacobian matrix of \hat{x} taken at y .

The model subspace captures what is linearly invariant through \hat{x} w.r.t. small perturbations of y ; typically, for Lasso, it will encode the set of signals sharing the same support. In order to generalize the refitting step, it is thus natural to cast it as a constrained optimization procedure preserving the model subspace.

Definition 2. The *invariant refitting* associated to an a.e. differentiable estimator $y \mapsto \hat{x}(y)$ is given for almost all $y \in \mathbb{R}^n$ by

$$(10) \quad \mathcal{R}_{\hat{x}}^{\text{inv}}(y) = \hat{x}(y) + J(\Phi J)^+(y - \Phi \hat{x}(y)) \in \underset{x \in \mathcal{M}_{\hat{x}}(y)}{\operatorname{argmin}} \frac{1}{2} \|\Phi x - y\|_2^2,$$

where $J = J_{\hat{x}}(y)$ is the Jacobian matrix of \hat{x} at the point y . In the following, we use the notation J when no ambiguity is possible.

Remark 3. We only consider $F(x, y) = \frac{1}{2} \|\Phi x - y\|_2^2$, but extension to a general F (e.g., logistic regression) is straightforward and reads $\mathcal{R}_{\hat{x}}^{\text{inv}}(y) \in \operatorname{argmin}_{x \in \mathcal{M}_{\hat{x}}(y)} F(x, y)$.

Remark 4. When $\hat{x}(y) \in \operatorname{Im}[J]$, then $\mathcal{M}_{\hat{x}}(y) = \operatorname{Im}[J]$ and $\mathcal{R}_{\hat{x}}^{\text{inv}}(y) = J(\Phi J)^+(y)$.

3.2. Refitting examples. We now exemplify the previous definitions for the various variational estimators introduced in subsection 2.2.

Affine constrained least-squares have everywhere the same Jacobian matrix $J = A(\Phi A)^+$, and its affine model subspace is $\mathcal{M}_{\hat{x}}(y) = b + \operatorname{Im}[A(\Phi A)^+]$ (as $\operatorname{Im}[M^+] = \operatorname{Im}[M^+]$). Taking $C = \mathbb{R}^p$, $n = p$, $A = \operatorname{Id}$, and $b = 0$ leads to $\hat{x}(y) = \Phi^+ y$, whose model subspace is $\mathcal{M}_{\hat{x}}(y) = \operatorname{Im}[\Phi^+]$, reducing to \mathbb{R}^p when Φ has full column rank. In this case, the invariant refitting is $\mathcal{R}_{\hat{x}}^{\text{inv}}(y) = \hat{x}(y)$.

Tikhonov regularization has everywhere the same Jacobian matrix $J = (\Phi^T \Phi + \lambda \Gamma^T \Gamma)^{-1} \Phi^T$ and everywhere the same affine model subspace $\mathcal{M}_{\hat{x}}(y) = \operatorname{Im}[J]$. It follows that $\mathcal{R}_{\hat{x}}^{\text{inv}}(y) = J(\Phi J)^+ y$. In particular, when Φ has full column rank, $\mathcal{M}_{\hat{x}}(y) = \mathbb{R}^p$ and $\mathcal{R}_{\hat{x}}^{\text{inv}}(y) = \Phi^+ y$.

Soft- and hard-thresholding share a.e. the same Jacobian matrix given by $J = \operatorname{Id}_{\mathcal{I}} \in \mathbb{R}^{p \times |\mathcal{I}|}$. Their model subspace reads as $\mathcal{M}_{\hat{x}}(y) = \operatorname{Im}[\operatorname{Id}_{\mathcal{I}}] = \operatorname{Im}[J]$, and the invariant refitting is the same for both hard-thresholding and soft-thresholding.

ℓ_1 *synthesis* (or Lasso) has for Jacobian matrix $J = \operatorname{Id}_{\mathcal{I}}(\Phi_{\mathcal{I}})^+$ a.e., and when $\Phi_{\mathcal{I}}$ has full column rank, it shares the same model subspace $\mathcal{M}_{\hat{x}}(y) = \operatorname{Im}[\operatorname{Id}_{\mathcal{I}}] = \operatorname{Im}[J]$ as the hard- and soft-thresholding. Its invariant refitting is in this case $\mathcal{R}_{\hat{x}}^{\text{inv}}(y) = \operatorname{Id}_{\mathcal{I}}(\Phi_{\mathcal{I}})^+ y$. While Lasso systematically underestimates the amplitude of the signal by a shift $\lambda \operatorname{Id}_{\mathcal{I}}((\Phi_{\mathcal{I}})^t \Phi_{\mathcal{I}})^{-1} s_{\mathcal{I}}$, the refitting $\mathcal{R}_{\hat{x}}^{\text{inv}}(y)$ is free of such a contraction.

ℓ_1 *analysis* has for Jacobian matrix $J = U(\Phi U)^+$ a.e. (since the Γ -support, \mathcal{I} , and the sign $s_{\mathcal{I}}$ are a.e. stable w.r.t. small perturbations [49]). It results that the model subspace reads as $\mathcal{M}_{\hat{x}}(y) = \operatorname{Ker}[\Gamma_{\mathcal{I}^c}] = \operatorname{Im}[J]$. The generalization of the Lasso refitting leads to the least-square estimator constrained on the Γ -support of $\hat{x}(y)$:

$$(11) \quad \mathcal{R}_{\hat{x}}^{\text{inv}}(y) = U(\Phi U)^+ y.$$

For aniso-TV denoising (i.e., $\Phi = \operatorname{Id}$ and $\Gamma = \nabla$), the model subspace is the space of images whose jumps are included in those of the solution $\hat{x}(y)$. The refitting procedure $\mathcal{R}_{\hat{x}}^{\text{inv}}$ is the projector $\Pi_{\operatorname{Im}[J]} = UU^+$ that performs a piecewise average of its input on each plateau of the solution (see Figure 1(a) for an illustration in one dimension).

$\ell_1 - \ell_2$ *analysis* admits a.e. the following Jacobian operator [50]:

$$(12) \quad J = U(U^T \Phi^T \Phi U + \lambda U^T \Gamma^T \Omega \Gamma U)^{-1} U^T \Phi^T y,$$

$$\text{where } \Omega : z \in \mathbb{R}^{m \times b} \mapsto \begin{cases} \frac{1}{\|(\Gamma \hat{x}(y))_i\|_2} \left(z_i - \left\langle z_i, \frac{(\Gamma \hat{x}(y))_i}{\|(\Gamma \hat{x}(y))_i\|_2} \right\rangle \frac{(\Gamma \hat{x}(y))_i}{\|(\Gamma \hat{x}(y))_i\|_2} \right) & \text{if } i \in \mathcal{I}, \\ 0 & \text{otherwise,} \end{cases}$$




Figure 3. (a) A noisy image $y = x_0 + w$ where x_0 contains six overlapping squares. (b) The solution $\hat{x}(y)$ of aniso-TV with $\lambda = 10$ and (c) its invariant refitting $\mathcal{R}_{\hat{x}}^{\text{inv}}(y)$ (which coincides in this case with the covariant one $\mathcal{R}_{\hat{x}}(y)$). (d) The solution $\hat{x}(y)$ of iso-TV with $\lambda = 10$, (e) its invariant refitting $\mathcal{R}_{\hat{x}}^{\text{inv}}(y)$, and (f) its covariant one $\mathcal{R}_{\hat{x}}(y)$.

with U and \mathcal{I} defined exactly as for the ℓ_1 analysis case. Note that (12) is well founded as soon as $\text{Ker}[\Gamma_{\mathcal{I}^c}] \cap \text{Ker} \Phi = \{0\}$. For weaker conditions, see [50, Example 26]. In the particular case where ΦU has full column rank, the model subspace matches with that of the ℓ_1 analysis, i.e., $\mathcal{M}_{\hat{x}}(y) = \text{Ker}[\Gamma_{\mathcal{I}^c}] = \text{Im}[J]$, and the refitting is also given by (11); hence $\mathcal{R}_{\hat{x}}^{\text{inv}}(y) = U(\Phi U)^+ y$. As a consequence, for the iso-TV denoising, i.e., when $\Phi = \text{Id}$ and $\Gamma = \nabla$, the invariant refitting step consists again in performing a piecewise average on each plateau of the solution.

Nonlocal means has a Jacobian matrix with a complex structure [51, 19]. In particular, computing the projection on the model subspace is challenging in this case, and so is the computation of the invariant refitting. Note that a greedy procedure was proposed in [15] to compute the invariant refitting.

3.3. Results and limitations. Figure 1(a) illustrates the invariant refitting in the case of a 1D total variation denoising example (ℓ_1 analysis estimator). It recovers the jumps of the underlying signal (adding an extra one) but systematically underestimates their amplitudes. As expected, refitting re-enhances the amplitudes of all plateaus toward the data. Figure 1(b) gives a geometrical interpretation in dimension $p = 3$ of the model subspace and the invariant refitting. The model subspace is represented as the tangent plane of \hat{x} at y , and its refitting is the projection of y on this plane. All elements of this plane share the same jumps with the solution $\hat{x}(y)$. Figure 3(a)–(c) gives a similar illustration for a two-dimensional (2D) aniso-TV denoising example.




Figure 4. (a) A noisy and incomplete image $y = \Phi x_0 + w$ where Φ is a masking operator encoding missing pixel values (in red) and x_0 is a smooth signal. (b) Tikhonov regularization $\hat{x}(y)$ with $\Gamma = \nabla$ and $\lambda = 20$, (c) its invariant refitting $\mathcal{R}_{\hat{x}}^{\text{inv}}(y)$, and (d) its covariant one $\mathcal{R}_{\hat{x}}(y)$.

While the invariant refitting acts properly for the ℓ_1 analysis estimator, it is less appealing in other scenarios. [Figure 2\(c\)–\(d\)](#) and [Figure 3\(d\)–\(e\)](#) give two illustrations of the invariant refitting of a 2D iso-TV denoising example. As for aniso-TV, the invariant refitting is the projection of y on the space of signals whose jumps are located at the same position as those of $\hat{x}(y)$. But unlike the anisotropic case, $\hat{x}(y)$ is not piecewise constant. Instead of being composed of distinct flat regions, it reveals smoothed transitions with dense supports (referred to as extended supports in [\[6\]](#)); see [Figure 2\(c\)](#). As a consequence, the invariant refitting reintroduces a large amount of the original noisy signal in these smooth but nonconstant areas, creating the artifacts observed in [Figure 2\(d\)](#) and [Figure 3\(e\)](#).

[Figure 4\(a\)–\(c\)](#) gives another illustration of the invariant refitting of a 2D Tikhonov masking example (with Φ a diagonal matrix with 0 or 1 elements on the diagonal and $\Gamma = \nabla$). While the dynamic of the Tikhonov solution $\hat{x}(y)$ has been strongly reduced, the refitting $\mathcal{R}_{\hat{x}}^{\text{inv}}(y)$ refits clearly the solution toward the original intensities. However, such a refitting is not satisfactory as it does not preserve the smoothness of the solution $\hat{x}(y)$.

In fact, the model subspace captures only what is linearly invariant through \hat{x} w.r.t. small perturbations of y . This includes the support of the solution for iso-TV and the absence of variations inside $\text{Im}[J]^\perp$ for Tikhonov regularization. In particular, it fails at capturing some of the desirable relationships between the entries of y and the entries of $\hat{x}(y)$, which we call the *covariants*. These relationships typically encode some of the local smoothness and nonlocal interactions between the entries of the solution $\hat{x}(y)$. Such crucial information is not encoded in the linear model subspace, but interestingly the Jacobian matrix captures by definition how much the entries of \hat{x} linearly vary w.r.t. all the entries of y . This is at the heart of *covariant refitting*, defined in the next section, and, for comparison, it produces the results given in [Figure 2\(e\)](#), [Figure 3\(f\)](#), and [Figure 4\(d\)](#).

4. Covariant LEast-square Refitting (CLEAR). The objective of this section is to present our main contribution, the introduction of the covariant refitting procedure CLEAR. We particularly aim at solving the issues raised in [subsection 3.3](#). Toward this goal, we put a stronger emphasis on the first-order behavior of the original estimator by imposing the conservation of its Jacobian, at least locally. To define the refitting procedure we first need to introduce a procedure which, loosely speaking, takes as input the original estimator and

a guess of Φx_0 and outputs a new estimator with some targeted properties. We next define our Covariant LEast-square Refitting (CLEAR) by choosing a specific guess for Φx_0 ; see [Definition 14](#).

4.1. Local approach and desired properties for a suitable refitting. In this subsection, our objective is to define, from the original estimator \hat{x} and a guess $z \in \mathbb{R}^n$ of Φx_0 , a new estimator $\mathcal{D}_{\hat{x},z} : \mathbb{R}^n \rightarrow \mathbb{R}^p$ that satisfies several desirable properties and shares with \hat{x} some first-order properties. Afterward, we will consider the choice $z = y$, and the resulting estimator is going to be our covariant refitting $\mathcal{R}_{\hat{x}}$. We are now equipped to introduce such a guess-based refitting.

Definition 5. Let $\hat{x} : \mathbb{R}^n \rightarrow \mathbb{R}^p$ be differentiable at $z \in \mathbb{R}^n$. An estimator $\mathcal{D}_{\hat{x},z} : \mathbb{R}^n \rightarrow \mathbb{R}^p$ is a guess-based covariant least-square refitting of \hat{x} for z if

$$(13) \quad \mathcal{D}_{\hat{x},z} \in \underset{h \in \mathcal{H}}{\operatorname{argmin}} \|\Phi h(z) - z\|_2^2,$$

where \mathcal{H} is the set of maps $h : \mathbb{R}^n \rightarrow \mathbb{R}^p$ satisfying, for all $y \in \mathbb{R}^n$,

1. *Affine map:* $h(y) = Ay + b$ for some $A \in \mathbb{R}^{p \times n}, b \in \mathbb{R}^p$.
2. *Covariant preserving:* $J_h(z) = \rho J_{\hat{x}}(z)$ for some $\rho \in \mathbb{R}$.
3. *Coherent map:* $h(\Phi \hat{x}(z)) = \hat{x}(z)$.

[Definition 5](#) is natural as it states that a guess-based refitting of \hat{x} for z should be, in prediction, as close as possible to z . Of course, it should satisfy some extra conditions. First, the estimator should be easy to compute, and so we choose a first-order approximation, leading to a locally affine estimator. Second, as highlighted in [subsection 3.3](#), the relative variation of the original estimator w.r.t. the input should be preserved to capture not only the invariant features of the estimator but also its first-order behavior, capturing both its singularities and smoothness. Third, applying a refitting step to the prediction obtained by the original estimator at z should not modify it. The purpose of refitting is to be close to y while also preserving the structure of $\hat{x}(z)$. Hence, if $y = \Phi \hat{x}(z)$, the result should be unaltered.

The next theorem provides a unique closed form expression for $\mathcal{D}_{\hat{x},z}(y)$.

Theorem 6. Let \hat{x} be an estimator from \mathbb{R}^n to \mathbb{R}^p differentiable at $z \in \mathbb{R}^n$. Then, for $\delta = z - \Phi \hat{x}(z)$, the guess-based covariant least-square refitting, defined in [Definition 5](#), exists, is unique if $\Phi J \delta \neq 0$, and is given by

$$(14) \quad \mathcal{D}_{\hat{x},z}(y) = \hat{x}(z) + \rho J(y - \Phi \hat{x}(z)), \quad \text{where } \rho = \begin{cases} \frac{\langle \Phi J \delta, \delta \rangle}{\|\Phi J \delta\|_2^2} & \text{if } \Phi J \delta \neq 0, \\ 1 & \text{otherwise,} \end{cases}$$

where $J = J_{\hat{x}}(z)$ is the Jacobian matrix of \hat{x} at the point z .

Proof. Let h be a mapping satisfying properties 1, 2, and 3 in the previous definition. Observe that properties 1 and 2 of the set \mathcal{H} together ensure that the estimator is of the form $h(y) = \rho J y + b$ for some $\rho \in \mathbb{R}$ and $b \in \mathbb{R}^p$. Plugging in condition 3 gives that $b = (\operatorname{Id} - \rho J \Phi) \hat{x}(z)$; hence $h(y) = \hat{x}(z) + \rho J(y - \Phi \hat{x}(z))$. Reciprocally, it is easy to see that any estimator of the form $h(y) = \hat{x}(z) + \rho J(y - \Phi \hat{x}(z))$ satisfies properties 1, 2, and 3. It thus remains to find ρ .

Note that problem (13) can be recast as a 1D problem,

$$(15) \quad \min_{\rho \in \mathbb{R}} \left\{ \|\Phi(\hat{x}(z) + \rho J(z - \Phi\hat{x}(z))) - z\|_2^2 = \|(\text{Id} - \rho\Phi J)(\Phi\hat{x}(z) - z)\|_2^2 \right\},$$

whose unique solution, if $\Phi J(\Phi\hat{x}(z) - z) \neq 0$, is given by (14), and $\rho = 1$ otherwise. \blacksquare

The case where ΦJ is an orthogonal projector leads to interesting properties, such as when \hat{x} is associated to constrained least-squares, Lasso, or aniso-TV.

Proposition 7. *If ΦJ is an orthogonal projector, then $\rho = 1$.*

Proof. If ΦJ is an orthogonal projector, then it follows that $\|\Phi J\delta\|_2^2 = \langle \Phi J\delta, \Phi J\delta \rangle = \langle (\Phi J)^\top \Phi J\delta, \delta \rangle = \langle \Phi J\delta, \delta \rangle$. Injecting this into (14) gives $\rho = 1$. \blacksquare

Statistical interpretation. For a random vector with expectation Φx_0 and finite second-order moment, the bias and covariance of $\mathcal{D}_{\hat{x},z}$ are given in closed form.

Proposition 8. *Let Y be a random vector in \mathbb{R}^n such that $\mathbb{E}[Y] = \Phi x_0$ and $\text{Cov}[Y] = \Sigma \in \mathbb{R}^{n \times n}$. Then $y \mapsto \mathcal{D}_{\hat{x},z}(y)$ satisfies*

$$(16) \quad \mathbb{E}[\mathcal{D}_{\hat{x},z}(Y)] - x_0 = (\text{Id} - \rho J\Phi)(\hat{x}(z) - x_0),$$

$$(17) \quad \text{Cov}[\mathcal{D}_{\hat{x},z}(Y), Y] = \rho J\Sigma,$$

$$(18) \quad \text{Cov}[\mathcal{D}_{\hat{x},z}(Y)] = \rho^2 J\Sigma J^\top,$$

where the cross covariance is $\text{Cov}[X, Y] = \mathbb{E}[XY^\top] - \mathbb{E}[X]\mathbb{E}[Y]^\top$ for any random column vectors X and Y (not necessarily of the same size) and $\text{Cov}[Y] = \text{Cov}[Y, Y]$.

Proof. The first equality is a direct consequence of the linearity of the expectation operator. The second equality arises from

$$(19) \quad \mathbb{E}[(\hat{x}(z) + \rho J(Y - \Phi\hat{x}(z)))Y^\top] - \mathbb{E}[(\hat{x}(z) + \rho J(Y - \Phi\hat{x}(z)))]\mathbb{E}[Y]^\top \\ = \rho J \left(\mathbb{E}[YY^\top] - \mathbb{E}[Y]\mathbb{E}[Y]^\top \right)$$

since J and ρ are constant w.r.t. y as they depend only on the guess z . The third equation follows the same sketch by expanding the expression of $\text{Cov}[\mathcal{D}_{\hat{x},z}(Y)]$. \blacksquare

Proposition 8 provides a closed form expression for the bias, the cross-covariance, and the covariance of $\mathcal{D}_{\hat{x},z}$. The derivation of these quantities for a nonlinear estimator \hat{x} is a much more intricate task. Nevertheless, the next corollary shows how these quantities relate to those of the first-order Taylor expansion of the original estimator \hat{x} .

Corollary 9. *Let $\mathcal{T}_{\hat{x},z}(y)$ be the tangent estimator of \hat{x} at $z \in \mathbb{R}^n$ defined as*

$$(20) \quad \mathcal{T}_{\hat{x},z}(y) = \hat{x}(z) + J(y - z).$$

Let Y be a random vector in \mathbb{R}^n such that $\mathbb{E}[Y] = \Phi x_0$ and $\text{Cov}[Y] = \Sigma$. Then $y \mapsto \mathcal{T}_{\hat{x},z}(y)$

and $y \mapsto \mathcal{D}_{\hat{x},z}(y)$ satisfy

$$(21) \quad \mathbb{E}[\mathcal{T}_{\hat{x},z}(Y)] - x_0 = (\hat{x}(z) - x_0) + J(\Phi x_0 - z),$$

$$(22) \quad \text{Cov}[\mathcal{D}_{\hat{x},z}(Y), Y] = \rho \text{Cov}[\mathcal{T}_{\hat{x},z}(Y), Y],$$

$$(23) \quad \text{Cov}[\mathcal{D}_{\hat{x},z}(Y)] = \rho^2 \text{Cov}[\mathcal{T}_{\hat{x},z}(Y)],$$

$$(24) \quad \text{Corr}[\mathcal{D}_{\hat{x},z}(Y), Y] = \text{Corr}[\mathcal{T}_{\hat{x},z}(Y), Y],$$

$$(25) \quad \text{Corr}[\mathcal{D}_{\hat{x},z}(Y)] = \text{Corr}[\mathcal{T}_{\hat{x},z}(Y)],$$

where the cross-correlation matrix is defined as

$$\text{Corr}[X, Y]_{i,j} = \text{Cov}[X, Y]_{i,j} / \sqrt{\text{Cov}[X]_{i,i} \text{Cov}[Y]_{j,j}}$$

for any random column vectors X and Y (not necessarily of the same size), and $\text{Corr}[Y] = \text{Corr}[Y, Y]$.

Proof. The first relation holds from the expression of $\mathcal{T}_{\hat{x},z}$ and that J does not depend on y . It follows that $\text{Cov}[\mathcal{T}_{\hat{x},z}(Y), Y] = J\Sigma$ and $\text{Cov}[\mathcal{T}_{\hat{x},z}(Y)] = J\Sigma J^\top$. These, jointly with [Proposition 8](#), conclude the proof. \blacksquare

[Corollary 9](#) is essential in this work as it states that, by preserving the Jacobian structure, $\mathcal{D}_{\hat{x},z}(Y)$ cannot depart from the tangent estimator of \hat{x} at z in terms of (cross-)correlations. As a consequence, one can expect that they only differ in terms of expectation, i.e., in terms of bias. The next propositions state that when ΦJ is a projector, the bias in prediction is guaranteed to be reduced by our refitting.

[Proposition 10.](#) *Let Y be a random vector of \mathbb{R}^n such that $\mathbb{E}[Y] = \Phi x_0$. Assume ΦJ is an orthogonal projector; then $y \mapsto \mathcal{D}_{\hat{x},z}(y)$ satisfies*

$$(26) \quad \|\Phi(\mathbb{E}[\mathcal{D}_{\hat{x},z}(Y)] - x_0)\|_2 \leq \|\Phi(\mathbb{E}[\mathcal{T}_{\hat{x},z}(Y)] - x_0)\|_2.$$

Proof. As ΦJ is an orthogonal projector, by virtue of [Proposition 7](#), $\rho = 1$; then

$$(27) \quad \begin{aligned} \|\Phi(\mathbb{E}[\mathcal{T}_{\hat{x},z}(Y)] - x_0)\|_2^2 &= \|\Phi(\hat{x}(z) + J(\Phi x_0 - z) - x_0)\|_2^2 \\ &= \|\Phi(\hat{x}(z) + J\Phi(x_0 - \hat{x}(z)) - x_0)\|_2^2 + \|\Phi J(\Phi \hat{x}(z) - z)\|_2^2 \\ &= \|\Phi(\mathbb{E}[\mathcal{D}_{\hat{x},z}(Y)] - x_0)\|_2^2 + \|\Phi J\delta\|_2^2, \end{aligned}$$

which concludes the proof. \blacksquare

[Proposition 10](#) is a bit restrictive as it requires ΦJ to be a projector. Nevertheless, this assumption can be relaxed when z satisfies a more technical assumption, as shown in the next proposition.

[Proposition 11.](#) *Let Y be a random vector of \mathbb{R}^n such that $\mathbb{E}[Y] = \Phi x_0$. Let $\rho_0 = \frac{\langle \delta_0, \Phi J \delta_0 \rangle}{\|\Phi J \delta_0\|_2^2}$ and $\delta_0 = \Phi(x_0 - \hat{x}(z))$. Assume there exists $\alpha \in [0, 1]$ such that*

$$(28) \quad \left| \frac{\rho - \rho_0}{\rho_0} \right| \leq \sqrt{1 - \alpha}$$

$$(29) \quad \text{and} \quad \|\Phi J(\delta - \delta_0)\|_2^2 + 2\langle \delta_0, \Phi J(\delta - \delta_0) \rangle \geq -\alpha \frac{\langle \delta_0, \Phi J \delta_0 \rangle^2}{\|\Phi J \delta_0\|_2^2}.$$

Then, $y \mapsto \mathcal{D}_{\hat{x},z}(y)$ satisfies

$$(30) \quad \|\Phi(\mathbb{E}[\mathcal{D}_{\hat{x},z}(Y)] - x_0)\|_2 \leq \|\Phi(\mathbb{E}[\mathcal{T}_{\hat{x},z}(Y)] - x_0)\|_2.$$

Proof. It follows from [Proposition 8](#) and [Corollary 9](#) that $\|\Phi(\mathbb{E}[\mathcal{D}_{\hat{x},z}(Y)] - x_0)\|_2 = \|(\text{Id} - \rho\Phi J)\delta_0\|_2$ and $\|\Phi(\mathbb{E}[\mathcal{T}_{\hat{x},z}(Y)] - x_0)\|_2 = \|\delta_0 + \Phi J(\delta - \delta_0)\|_2$. Subsequently, we get that (30) holds true if

$$(31) \quad \|(\text{Id} - \rho\Phi J)\delta_0\|_2^2 \leq \|\delta_0 + \Phi J(\delta - \delta_0)\|_2^2,$$

$$(32) \quad \text{i.e., } \rho^2\|\Phi J\delta_0\|_2^2 - 2\rho\langle\delta_0, \Phi J\delta_0\rangle \leq \|\Phi J(\delta - \delta_0)\|_2^2 + 2\langle\delta_0, \Phi J(\delta - \delta_0)\rangle.$$

Using assumption (29), a sufficient condition for (30) to hold is

$$(33) \quad \rho^2\|\Phi J\delta_0\|_2^2 - 2\rho\langle\delta_0, \Phi J\delta_0\rangle + \alpha \frac{\langle\delta_0, \Phi J\delta_0\rangle^2}{\|\Phi J\delta_0\|_2^2} \leq 0.$$

The roots of this polynomial are given by $(1 \pm \sqrt{1 - \alpha})\rho_0$, which concludes the proof. \blacksquare

Remark 12. Note that requiring (28) is quite natural, as it states that ρ should be close enough to the optimal ρ_0 , minimizing the discrepancy with regard to Φx_0 (i.e., minimizing $\|\Phi h(z) - \Phi x_0\|_2^2$ for $h \in \mathcal{H}$ defined as in [Definition 5](#)). While the condition (29) sounds more technical, it holds true in several interesting cases. For instance, when $z = \Phi x_0$, assumption (29) holds true, as it would read $0 \geq -\alpha\langle\delta_0, \Phi J\delta_0\rangle^2/\|\Phi J\delta_0\|_2^2$ (since $\delta = \delta_0$ and $\rho = \rho_0$). Another case of interest is when ΦJ is an orthogonal projector for which (29) holds true, as it would read $\|\Phi J\delta\|_2^2 - \|\Phi J\delta_0\|_2^2 \geq -\|\Phi J\delta_0\|_2^2$ (using that $\rho = \rho_0 = 1$, $\langle\cdot, \Phi J\cdot\rangle = \|\Phi J\cdot\|_2^2$, and choosing $\alpha = 1$). Hence, [Proposition 11](#) recovers [Proposition 10](#).

Remark 13. Using the same sketch of proof as we did for [Proposition 11](#), we find that the condition $|(\rho - \rho_0)/\rho_0| \leq 1$ is sufficient to get $\|\Phi(\mathbb{E}[\mathcal{D}_{\hat{x},z}(Y)] - x_0)\|_2 \leq \|\Phi(\hat{x}(z) - x_0)\|_2$. In other words, even though ρ has a relative error of 100% w.r.t. ρ_0 , the estimator $y \mapsto \mathcal{D}_{\hat{x},z}(y)$ still reduces the bias of the constant estimator $y \mapsto \hat{x}(z)$. This result remains valid when comparing $y \mapsto \mathcal{D}_{\hat{x},z}(y)$ to the pseudo-oracle estimator $y \mapsto \hat{x}(z) + J(y - \Phi x_0)$, with the notable difference that they moreover share the same correlation structure.

While it is difficult to state a general result, we can reasonably claim from [Proposition 10](#), [Proposition 11](#), and [Remark 12](#) that the bias tends to be reduced by our refitting provided ΦJ is almost a projector (i.e., has eigenvalues concentrated around 0 and 1) and/or z is not too far from Φx_0 . In such cases, the estimator $\mathcal{D}_{\hat{x},z}$ can be considered as a debiasing procedure of \hat{x} , in the sense that it reduces the bias of $\mathcal{T}_{\hat{x},z}$ while preserving its correlation structure (according to [Corollary 9](#)).

4.2. Definitions and properties. Using $\mathcal{D}_{\hat{x},z}$ defined in [Theorem 6](#), we can now give an explicit definition of CLEAR as $\mathcal{R}_{\hat{x}}(y) = \mathcal{D}_{\hat{x},y}(y)$.

Definition 14 (CLEAR). The *Covariant LEast-square Refitting* associated to an a.e. differentiable estimator $y \mapsto \hat{x}(y)$ is, for almost all $y \in \mathbb{R}^n$, given by

$$(34) \quad \mathcal{R}_{\hat{x}}(y) = \hat{x}(y) + \rho J(y - \Phi\hat{x}(y)) \quad \text{with} \quad \rho = \begin{cases} \frac{\langle\Phi J\delta, \delta\rangle}{\|\Phi J\delta\|_2^2} & \text{if } \Phi J\delta \neq 0, \\ 1 & \text{otherwise,} \end{cases}$$

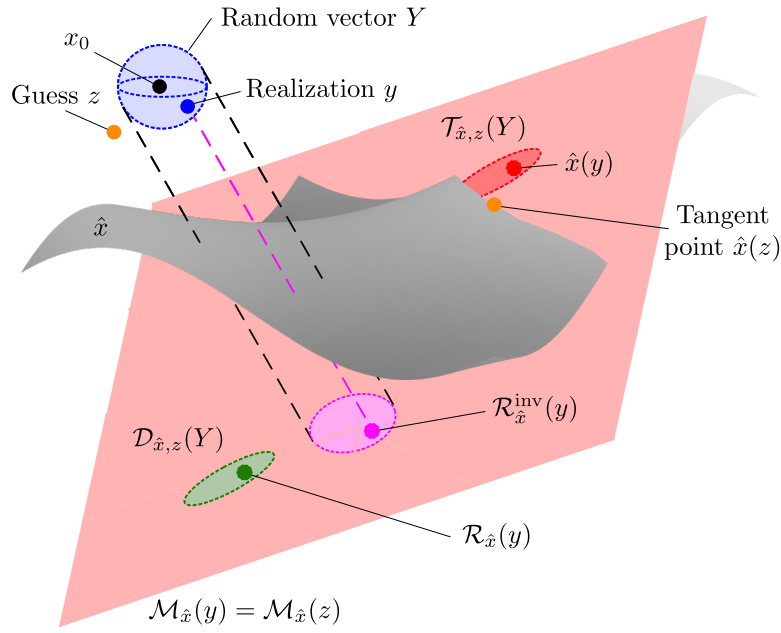


Figure 5. Geometrical illustration of the covariant refitting in a denoising problem of dimension $p = 3$. We assume that $\mathcal{M}_{\hat{x}}(z) = \mathcal{M}_{\hat{x}}(y)$ for the sake of clarity. The gray surface is the manifold modeling the evolution of \hat{x} in an extended neighborhood of y . The light red affine plane is the model subspace tangent at z . The ellipses represent the positive-definite symmetric covariance matrices of some random vectors, as defined in Proposition 8 and Corollary 9.

where $\delta = y - \Phi\hat{x}(y)$ and $J = J_{\hat{x}}(y)$ is the Jacobian matrix of \hat{x} at the point y .

Figure 5 gives a geometrical interpretation of CLEAR for a denoising task in dimension $p = 3$. One can observe that if Y varies isotropically, so will its projection on the model subspace. Contrarily, the tangent estimator at a guess z can present an anisotropic behavior along the model subspace, and the guess-based refitting, which is closer to z , will respect this anisotropy in order to capture the local regularity of \hat{x} . Finally, the covariant refitting is obtained from the guess-based refitting at $z = y$. For clarity, we assumed that $\mathcal{M}_{\hat{x}}(z) = \mathcal{M}_{\hat{x}}(y)$ in this illustration.

Remark 15. The covariant refitting performs an additive correction of $\hat{x}(y)$ with a fraction of the directional derivative $J\delta$ in the direction of the residual δ .

Remark 16. Observe that in Definition 14, the value of ρ varies when y varies, contrary to the map $y \mapsto \mathcal{D}_{\hat{x},z}(y)$ for which ρ is constant. Note that the mapping $y \mapsto \mathcal{D}_{\hat{x},z}(y)$ is affine, but the map $y \mapsto \mathcal{R}_{\hat{x}}(y)$ is *not*. Note that, as a consequence, the statistical interpretations given in the previous section do not hold for $\mathcal{R}_{\hat{x}}(y)$ even though they shed some light on its behavior.

Why not iterate the procedure? The refitting procedure computes $\tilde{x}^2 = \tilde{x}^1 + \rho J(y - \Phi\tilde{x}^1)$, with $\tilde{x}^1 = \hat{x}(y)$. One may wonder if it is beneficial to iterate the process as $\tilde{x}^{k+1} = \tilde{x}^k + \rho J(y - \Phi\tilde{x}^k)$ (in the same vein as [48, 4, 34, 44, 38]). Consider a denoising problem $\Phi = \text{Id}$ with Tikhonov or iso-TV, for which J is symmetrical and $\hat{x}(y) \in \text{Im}[J]$ (see examples in

subsection 3.2). The sequence will converge if and only if $J(y - \tilde{x}^k)$ vanishes; i.e., \tilde{x}^k must converge to $J^+Jy + \zeta$ with $\zeta \in \text{Ker}[J]$. By construction, $\tilde{x}^k \in \text{Im}[J] = \text{Ker}[J^\top]^\perp = \text{Ker}[J]^\perp$; hence $\zeta = 0$. Moreover, as J is symmetrical and $\hat{x}(y) \in \text{Im}[J]$, the quantity J^+Jy coincides with $JJ^+y = \mathcal{R}_{\hat{x}}^{\text{inv}}(y)$ (by virtue of Remark 4), i.e., the invariant refitting. Remembering the artifacts illustrated in Figure 3(e), this is not satisfying.

An interesting property of $\mathcal{R}_{\hat{x}}$ is the fact that it belongs to the model subspace of \hat{x} as stated in the following proposition.

Proposition 17. *Let $y \mapsto \hat{x}(y)$ be an a.e. differentiable estimator. Then for almost all $y \in \mathbb{R}^n$, one has $\mathcal{R}_{\hat{x}}(y) \in \mathcal{M}_{\hat{x}}(y)$.*

Proof. As $\mathcal{M}_{\hat{x}}(y) = \hat{x}(y) + \text{Im} J$ and $\rho J(y - \Phi\hat{x}(y)) \in \text{Im} J$, the proposition holds. ■

The case where ΦJ is an orthogonal projector leads to properties that will be of interest regarding some estimators considered in subsection 2.2.

Proposition 18. *Suppose that ΦJ is an orthogonal projector. Then, $\mathcal{R}_{\hat{x}}(y) = \hat{x}(y) + J(y - \Phi\hat{x}(y))$, and $\Phi\mathcal{R}_{\hat{x}}(y) = \Phi\mathcal{R}_{\hat{x}}^{\text{inv}}(y)$.*

Proof. By virtue of Proposition 7, $\rho = 1$, and then $\mathcal{R}_{\hat{x}}(y) = \hat{x}(y) + J(y - \Phi\hat{x}(y))$. The fact that $\Phi\mathcal{R}_{\hat{x}}(y) = \Phi\mathcal{R}_{\hat{x}}^{\text{inv}}(y)$ comes from the fact that $\Phi J(\Phi J)^+ = \Phi J$. ■

The next proposition provides, when $J\Phi$ satisfies a fixed point formulation, an expression of $\mathcal{R}_{\hat{x}}(y)$ that will be useful for efficient computations of the refitting as discussed in section 5, a notable example being iso-TV regularization.

Proposition 19. *Assume that $J\Phi\hat{x}(y) = \hat{x}(y)$. Then the covariant refitting reads $\mathcal{R}_{\hat{x}}(y) = (1 - \rho)\hat{x}(y) + \rho Jy$.*

Proof. We have $\mathcal{R}_{\hat{x}}(y) = \hat{x}(y) + \rho J(y - \Phi\hat{x}(y)) = \hat{x}(y) + \rho Jy - \rho J\Phi\hat{x}(y)$, and since $J\Phi\hat{x}(y) = \hat{x}(y)$ by assumption, this concludes the proof. ■

Interestingly, the next theorem shows that the condition $J\Phi\hat{x}(y) = \hat{x}(y)$ is met provided $\hat{x}(y)$ is solution of a variational problem with a 1-homogeneous regularizer.

Theorem 20. *Let $\hat{x}(y)$ be the unique a.e. differentiable solution of*

$$(35) \quad \hat{x}(y) = \underset{x}{\text{argmin}} F(y - \Phi x) + G(x),$$

with F, G being convex and G being 1-homogeneous. Then $J\Phi\hat{x}(y) = \hat{x}(y)$ a.e.

The proof of Theorem 20 is postponed to Appendix B.

Affine constrained least-squares, ℓ_1 synthesis, $\ell_1 - \ell_2$ analysis, aniso-TV, and iso-TV are solutions of a variational problem, with F being differentiable and G being 1-homogeneous. As a consequence, Theorem 20 shows that the aforementioned methods satisfy $J\Phi\hat{x}(y) = \hat{x}(y)$, and hence $\mathcal{R}_{\hat{x}}(y) = (1 - \rho)\hat{x}(y) + \rho Jy$.

4.3. Examples of refitting procedures. We now exemplify the previous definitions for the wide class of variational estimators introduced in subsection 2.2.

Affine constrained least-squares have Jacobian matrix $J = A(\Phi A)^+$. In this case, $\Phi J = \Phi A(\Phi A)^+$ is an orthogonal projector, $\rho = 1$, and the covariant refitting coincides with the invariant one and reads $\mathcal{R}_{\hat{x}}(y) = \mathcal{R}_{\hat{x}}^{\text{inv}}(y) = \hat{x}(y)$.

Tikhonov regularization has Jacobian matrix $J = (\Phi^\top \Phi + \lambda \Gamma^\top \Gamma)^{-1} \Phi^\top$, and in this case ρ depends on the residual δ , and the refitting reads as the weighted sum $\mathcal{R}_{\hat{x}}(y) = (1 + \rho)\hat{x}(y) - \rho J \Phi \hat{x}(y)$.

Soft- and hard-thresholding, used with $\Phi = \text{Id}$, have Jacobian matrix $J = \text{Id}_{\mathcal{I}}$. As a consequence, $\Phi J = \text{Id}_{\mathcal{I}}$ is an orthogonal projection, and the covariant refitting coincides with the invariant one, namely the hard-thresholding itself.

ℓ_1 *synthesis* has Jacobian matrix $J = \text{Id}_{\mathcal{I}}(\Phi_{\mathcal{I}})^+$, where $\Phi_{\mathcal{I}}$ has full column rank. As for the thresholding, $\Phi J = \Phi U(\Phi U)^+$ is an orthogonal projection, and the covariant refitting reads $\mathcal{R}_{\hat{x}}(y) = \mathcal{R}_{\hat{x}}^{\text{inv}}(y)$.

ℓ_1 *analysis* has Jacobian matrix $J = U(\Phi U)^+$. Again, $\Phi J = \Phi U(\Phi U)^+$ is an orthogonal projection, and the covariant refitting reads $\mathcal{R}_{\hat{x}}(y) = \mathcal{R}_{\hat{x}}^{\text{inv}}(y)$.

$\ell_1 - \ell_2$ *analysis* has the Jacobian operator given in (12), which applied to a vector $d \in \mathbb{R}^n$ is a solution of the following problem:

$$(36) \quad Jd \in \underset{x; \text{supp}(\Gamma x) \subseteq \mathcal{I}}{\text{argmin}} \quad \frac{1}{2} \|\Phi x - d\|_2^2 + \frac{\lambda}{2} \omega(\Gamma x),$$

$$\text{where } \omega : z \in \mathbb{R}^{m \times b} \mapsto \sum_{i \in \mathcal{I}} \frac{1}{\|(\Gamma \hat{x}(y))_i\|_2} \left(\|z_i\|_2^2 - \left\langle z_i, \frac{(\Gamma \hat{x}(y))_i}{\|(\Gamma \hat{x}(y))_i\|_2} \right\rangle^2 \right).$$

Note that $\omega(\Gamma Jd) = 0$ only if $(\Gamma Jd)_i$ is colinear to $(\Gamma \hat{x}(y))_i$ for all $i \in \mathcal{I}$. For iso-TV, this means that the level lines of Jd must be included in those of $\hat{x}(y)$. Moreover, by virtue of [Theorem 20](#), one has $J\Phi \hat{x}(y) = \hat{x}(y)$, and hence $\mathcal{R}_{\hat{x}}(y) = (1 - \rho)\hat{x}(y) + \rho Jy$. As a consequence, unlike the invariant refitting of $\hat{x}(y)$, the covariant refitting is constrained to be faithful to the regularity of $\hat{x}(y)$, since it enforces the discontinuities of Jd to be colinear to $(\Gamma \hat{x}(y))_{\mathcal{I}}$. This is especially important where the iso-TV solution presents transitions with high curvature. Such appealing behavior of the covariant refitting explains the results observed in [Figure 2\(e\)](#) and [Figure 3\(f\)](#).

Nonlocal means has an intricate Jacobian matrix. Nevertheless, its directional derivative has a simpler expression given, for any direction $d \in \mathbb{R}^n$, by

$$(37) \quad Jd = \frac{\sum_j \bar{w}'_{i,j} y_j + \sum \bar{w}_{i,j} d_j - \hat{x}(y)_i \sum_j \bar{w}'_{i,j}}{\sum_j \bar{w}_{i,j}} \quad \text{with} \quad \bar{w}'_{i,j} = \sum_k w'_{i-k,j-k},$$

where $\hat{x}(y)$ is defined in (7) and

$$(38) \quad w'_{i,j} = 2 \langle \mathcal{P}_i y - \mathcal{P}_j y, \mathcal{P}_i d - \mathcal{P}_j d \rangle \varphi' (\|\mathcal{P}_i y - \mathcal{P}_j y\|^2),$$

with φ' the a.e. derivative of the kernel function φ . Subsequently, the covariant refitting is obtained from its general form with two steps, by first computing $\hat{x}(y)$ and then applying the Jacobian to the direction $d = y - \hat{x}(y)$.

5. Covariant refitting in practice. This section details the computation of CLEAR for standard algorithms. We first recall some properties of two different differentiation techniques that allow computing some image of J jointly with $\hat{x}(y)$.

5.1. Algorithmic differentiation. Following [16], we consider restoration algorithms whose solutions $\hat{x}(y) = x^k$ are obtained via an iterative scheme of the form

$$(39) \quad \begin{cases} x^k &= \gamma(a^k), \\ a^{k+1} &= \psi(a^k, y). \end{cases}$$

Here, $a^k \in \mathcal{A}$ is a sequence of auxiliary variables, $\psi : \mathcal{A} \times \mathbb{R}^n \rightarrow \mathcal{A}$ is a fixed point operator in the sense that a^k converges to a^* , and $\gamma : \mathcal{A} \rightarrow \mathbb{R}^p$ is nonexpansive (i.e., $\|\gamma(a_1) - \gamma(a_2)\| \leq \|a_1 - a_2\|$ for all $a_1, a_2 \in \mathcal{A}$) entailing x^k converges to $x^* = \gamma(a^*)$.

As a result, for almost all y and for any direction $d \in \mathbb{R}^n$, the directional derivatives $\mathcal{D}_x^k = J_{\hat{x}^k}(y)d$ and $\mathcal{D}_a^k = J_{a^k}(y)d$ can be jointly obtained with x^k and a^k as

$$(40) \quad \begin{cases} x^k &= \gamma(a^k), \\ a^{k+1} &= \psi(a^k, y), \\ \mathcal{D}_x^k &= \Gamma_a \mathcal{D}_a^k, \\ \mathcal{D}_a^{k+1} &= \Psi_a \mathcal{D}_a^k + \Psi_y d, \end{cases}$$

where $\Gamma_a = \left. \frac{\partial \gamma(a)}{\partial a} \right|_{a^k}$, $\Psi_a = \left. \frac{\partial \psi(a, y)}{\partial a} \right|_{a^k}$, and $\Psi_y = \left. \frac{\partial \psi(a^k, y)}{\partial y} \right|_y$. Interestingly, in all considered cases, the cost of evaluating Γ_a , Ψ_a , and Ψ_y is about the same as the one of evaluating γ and ψ . As a result, the complexity of (40) is about twice the complexity of (39). In practice, Γ_a , Ψ_a , and Ψ_y can be implemented either thanks to their closed form expression or in a black box manner using automatic differentiation. The latter has been well studied, and we refer the reader to [25, 32] for a comprehensive study.

5.2. Finite difference based differentiation. Another strategy is to approximate directional derivatives by finite differences, for any $d \in \mathbb{R}^n$ and $\varepsilon > 0$, as

$$(41) \quad J_{\hat{x}}(y)d \approx \frac{\hat{x}(y + \varepsilon d) - \hat{x}(y)}{\varepsilon}.$$

As a result, the complexity of evaluating (41) is also twice the complexity of (39) since \hat{x} must be evaluated at both y and $y + \varepsilon d$. The main advantage of this method is that \hat{x} can be used as a black box, i.e., without any knowledge of the underlying algorithm that provides $\hat{x}(y)$. For ε small enough, it performs as well as the approach described in (40) (with $\hat{x}(y) = x^k$) that requires the knowledge of the derivatives. Indeed, if $y \mapsto \hat{x}(y)$ is Lipschitz-continuous, then (41) converges to (40) when $\varepsilon \rightarrow 0$ (by virtue of Rademacher's theorem and [22, Theorems 1–2, section 6.2]). This implies that the value ε can be chosen as small as possible (up to machine precision), yielding an accurate approximation of $J_{\hat{x}}(y)d$. This finite difference approach has been used in many fields, and notably for risk estimation; see, e.g., [54, 41, 36]. We will apply this black box strategy on state-of-the-art algorithms in [subsection 6.5](#).

5.3. Two-step computation for the general case. In the most general case, the computation of the covariant refitting, given by

$$(42) \quad \mathcal{R}_{\hat{x}}(y) = \hat{x}(y) + \rho J(y - \Phi \hat{x}(y)) \quad \text{with} \quad \rho = \frac{\langle \Phi J \delta, \delta \rangle}{\|\Phi J \delta\|_2^2} \quad \text{and} \quad \delta = y - \Phi \hat{x}(y),$$

requires us to sequentially evaluate $\hat{x}(y)$ and $J(y - \Phi\hat{x}(y))$.

With finite difference differentiation, two steps are required. First $\hat{x}(y)$ must be computed with the original algorithm, and next $J(y - \Phi\hat{x}(y))$ is obtained by finite difference (41) on the direction of the residual $d = y - \Phi\hat{x}(y)$. Once $J(y - \Phi\hat{x}(y))$ is computed, ρ can be evaluated, and subsequently (42). The overall complexity is about twice that of the original algorithm producing $\hat{x}(y)$.

With algorithmic differentiation, as $J(y - \Phi\hat{x}(y))$ depends on $\hat{x}(y)$, the original iterative scheme (39) must be run first. In the second step, $J(y - \Phi\hat{x}(y))$ is obtained with the differentiated version (40) on the direction of the residual $d = y - \Phi\hat{x}(y)$. As a result, $\hat{x}(y)$ is computed twice, first by (39), and next by (40). It leads to an overall complexity about three times that of the original algorithm. Nevertheless, in several cases, one can avoid the first step by running (40) only once.

5.4. One-step computation for specific cases. When $\hat{x}(y)$ fulfills the assumption $J\Phi\hat{x}(y) = \hat{x}(y)$ of Proposition 19, the covariant refitting reads as

$$(43) \quad \mathcal{R}_{\hat{x}}(y) = (1 - \rho)\hat{x}(y) + \rho Jy \quad \text{with} \quad \rho = \frac{\langle \Phi(Jy - \hat{x}(y)), y - \Phi\hat{x}(y) \rangle}{\|\Phi(Jy - \hat{x}(y))\|_2^2}.$$

The computations of $\hat{x}(y)$ and Jy are then sufficient to compute the refitting $\mathcal{R}_{\hat{x}}(y)$. As a result, in the case of algorithmic differentiation, (40) can be run once to get $\mathcal{R}_{\hat{x}}(y)$ since using $d = y$ directly provides $\hat{x}(y)$, Jy , and subsequently ρ . Compared to the two-step approach, the complexity of the refitting reduces to about twice that of the original step from (39). Recall that the condition $J\Phi\hat{x}(y) = \hat{x}(y)$ is met for Lasso, generalized Lasso, aniso-TV, and iso-TV; hence they can be re-enhanced with a complexity twice that of their original algorithm.

5.5. Example on a primal-dual solver for ℓ_1 analysis. In this section we instantiate algorithm (40) to the case of the primal-dual sequence of [7]. By dualizing the ℓ_1 analysis norm $x \mapsto \lambda\|\Gamma x\|_1$, the primal problem (4) can be reformulated, with $x^* = \hat{x}(y)$ as the following saddle-point problem:

$$(44) \quad (z^*, x^*) \in \arg \max_{z \in \mathbb{R}^m} \min_{x \in \mathbb{R}^p} \frac{1}{2} \|\Phi x - y\|_2^2 + \langle \Gamma x, z \rangle - \iota_{\mathcal{B}_\lambda}(z),$$

where $z^* \in \mathbb{R}^m$ is the dual variable, and $\mathcal{B}_\lambda = \{z \in \mathbb{R}^m : \|z\|_\infty \leq \lambda\}$ is the ℓ_∞ ball.

Problem (44) can be efficiently solved using the primal-dual algorithm of [7]. By taking $\sigma\tau < 1/\|\Gamma\|_2^2$, $\theta \in [0, 1]$ and initializing (for instance) $x^0 = v^0 = 0 \in \mathbb{R}^p$, $z^0 = 0 \in \mathbb{R}^m$, the algorithm reads

$$(45) \quad \begin{cases} z^{k+1} &= \Pi_{\mathcal{B}_\lambda}(z^k + \sigma\Gamma v^k), \\ x^{k+1} &= (\text{Id} + \tau\Phi^\top\Phi)^{-1}(x^k + \tau(\Phi^\top y - \Gamma^\top z^{k+1})), \\ v^{k+1} &= x^{k+1} + \theta(x^{k+1} - x^k), \end{cases}$$

where the projection of z over \mathcal{B}_λ is done componentwise as

$$(46) \quad \Pi_{\mathcal{B}_\lambda}(z)_i = \begin{cases} z_i & \text{if } |z_i| \leq \lambda, \\ \lambda \text{ sign}(z_i) & \text{otherwise.} \end{cases}$$

The sequence x^k converges to a solution x^* of the ℓ_1 analysis problem [7].

It is easy to check that the primal-dual sequence defined in (45) can be written in the general form considered in (39); see, for instance, [16]. As a result, we can use the algorithmic differentiation based strategy described by (40) as follows: for the initialization $\tilde{x}^0 = \tilde{v}^0 = 0 \in \mathbb{R}^p$, $\tilde{z}^0 = 0 \in \mathbb{R}^m$, and for $\beta = 0$, as

$$(47) \quad \begin{cases} z^{k+1} &= \Pi_{\mathcal{B}_\lambda}(z^k + \sigma\Gamma v^k), \\ x^{k+1} &= (\text{Id} + \tau\Phi^\top\Phi)^{-1}(x^k + \tau(\Phi^\top y - \Gamma^\top z^{k+1})), \\ v^{k+1} &= x^{k+1} + \theta(x^{k+1} - x^k), \\ \tilde{z}^{k+1} &= \Pi_{z^k + \sigma\Gamma v^k}(\tilde{z}^k + \sigma\Gamma\tilde{v}^k), \\ \tilde{x}^{k+1} &= (\text{Id} + \tau\Phi^\top\Phi)^{-1}(\tilde{x}^k + \tau(\Phi^\top y - \Gamma^\top \tilde{z}^{k+1})), \\ \tilde{v}^{k+1} &= \tilde{x}^{k+1} + \theta(\tilde{x}^{k+1} - \tilde{x}^k), \end{cases}$$

$$\text{where } \Pi_z(\tilde{z})_i = \begin{cases} \tilde{z}_i & \text{if } |z_i| \leq \lambda + \beta, \\ 0 & \text{otherwise.} \end{cases}$$

Recall that the refitting is $\mathcal{R}_{x^k}(y) = \tilde{x}^k$, since $J\Phi$ is an orthogonal projector.

Remark that the algorithmic differentiation of (45) is exactly (47) for $\beta = 0$; hence, $\tilde{x}^k = \mathcal{R}_{x^k}(y)$. However, if one wants to guarantee the convergence of the sequence \tilde{x}^k toward $\mathcal{R}_{\hat{x}}(y)$, one needs a small $\beta > 0$ as shown in the next theorem. In practice, β can be chosen as the smallest available positive floating number.

Theorem 21. *Assume that x^* satisfies (5) with ΦU full-column rank.² Let $\alpha > 0$ be the minimum nonzero value³ of $|\Gamma x^*|_i$ for all $i \in [m]$. Choose β such that $\alpha\sigma > \beta > 0$. Then, the sequence $\tilde{x}^k = \mathcal{R}_{x^k}(y)$ defined in (47) converges to the refitting $\mathcal{R}_{\hat{x}}(y)$ of $\hat{x}(y) = x^*$.*

The proof of this theorem is postponed to [Appendix C](#).

A similar result was obtained in [14] when solving the ℓ_1 analysis problem (4) with the Douglas–Rachford splitting algorithm described in [18, 9].

5.6. Example for a primal-dual solver for $\ell_1 - \ell_2$ analysis. The algorithm for $\ell_1 - \ell_2$ analysis regularization can be derived with the exact same considerations as for the ℓ_1 analysis case. The only difference in the application of the primal-dual algorithm comes from the nonlinear operation (46) that now reads, for $z \in \mathbb{R}^{m \times b}$, as

$$(48) \quad \Pi_{\mathcal{B}_\lambda}^{\text{iso}}(z)_i = \begin{cases} z_i & \text{if } \|z_i\|_2 \leq \lambda, \\ \lambda \frac{z_i}{\|z_i\|_2} & \text{otherwise.} \end{cases}$$

²This could be enforced as shown in [49].

³If $|\Gamma x^*|_i = 0$ for all $i \in [m]$, the result remains true for any $\alpha > 0$.

It follows that the algorithmic differentiation strategy reads as

$$(49) \quad \begin{cases} z^{k+1} &= \Pi_{\mathcal{B}_\lambda}^{\text{iso}}(z^k + \sigma\Gamma v^k), \\ x^{k+1} &= (\text{Id} + \tau\Phi^\top\Phi)^{-1}(x^k + \tau(\Phi^\top y - \Gamma^\top(z^{k+1}))), \\ v^{k+1} &= x^{k+1} + \theta(x^{k+1} - x^k), \\ \tilde{z}^{k+1} &= \Pi_{z^k + \sigma\Gamma v^k}^{\text{iso}}(\tilde{z}^k + \sigma\Gamma\tilde{v}^k), \\ \tilde{x}^{k+1} &= (\text{Id} + \tau\Phi^\top\Phi)^{-1}(\tilde{x}^k + \tau(\Phi^\top y - \Gamma^\top\tilde{z}^{k+1})), \\ \tilde{v}^{k+1} &= \tilde{x}^{k+1} + \theta(\tilde{x}^{k+1} - \tilde{x}^k), \end{cases}$$

$$\text{where } \Pi_z^{\text{iso}}(\tilde{z})_i = \begin{cases} \tilde{z}_i & \text{if } \|z_i\|_2 \leq \lambda + \beta, \\ \frac{\lambda}{\|z_i\|_2} \left(\tilde{z}_i - \left\langle \tilde{z}_i, \frac{z_i}{\|z_i\|_2} \right\rangle \frac{z_i}{\|z_i\|_2} \right) & \text{otherwise.} \end{cases}$$

Unlike the ℓ_1 case, the refitted solution is not \tilde{x}^k itself, but, following [subsection 5.4](#), it can be obtained at iteration k as

$$(50) \quad \mathcal{R}_{x^k}(y) = (1 - \rho)x^k + \rho\tilde{x}^k \quad \text{with} \quad \rho = \frac{\langle \Phi(\tilde{x}^k - x^k), y - \Phi x^k \rangle}{\|\Phi(\tilde{x}^k - x^k)\|_2^2}.$$

5.7. Example for the nonlocal means. In this section, we specify the update rule (40) to an acceleration of the nonlocal means inspired from [12, 13]. We use the procedure of [40] to correctly handle the central pixel. Again, one can check that this implementation can be written in the general form considered in the update rule (39), where the fixed point solution is obtained directly at the first iteration.

The pseudocode obtained with the algorithmic differentiation scheme, as described in (40), is given in [Figure 6](#). All variables with the suffix $'$ correspond to the directional derivative obtained by using the chain rule on the original variables. This fast computation relies on the fact that all convolutions can be computed with integral tables, leading to a global complexity in $O(s^2n)$, for both the computation of the estimator $\hat{x}(y)$ and its directional derivative Jd . Recall that the covariant refitting is obtained from its general form in two steps, by first computing first $\hat{x}(y)$ and then applying the proposed pseudocode in the direction $d = y - \hat{x}(y)$.

6. Numerical experiments and comparisons with related approaches. Here, we first give illustrations of our CLEAR method on toy image restoration problems. Then, through quantitative results in terms of peak signal-to-noise ratio (PSNR)⁴ and the structural similarity (SSIM) [52], we numerically evaluate the refitting, discuss its benefit and limitations in several scenarios, and compare this method with popular approaches from the literature.

6.1. Denoising with isotropic total-variation (iso-TV). [Figure 7](#) gives an illustration of our covariant refitting of the 2D iso-TV, where λ has been chosen large enough to highlight the behavior of the refitting. We apply it to the denoising (i.e., $\Phi = \text{Id}$) of an 8-bit piecewise constant image damaged by AWGN with standard deviation $\sigma = 20$, known as the *Shepp-Logan phantom*. As discussed earlier, iso-TV introduces a significant loss of contrast [43], typically for thin detailed structures, which are re-enhanced in our result.

The residuals $\hat{x}(y) - x_0$ and $\mathcal{R}_{\hat{x}}(y) - x_0$ highlight that our refitting technique efficiently re-enhances the attenuated structure while leaving the lost structures unchanged. Nevertheless,

⁴PSNR = $10 \log_{10} 255^2 / \frac{1}{p} \|\hat{x}(y) - x_0\|_2^2$ for an image ranging on $[0, 255]$.

Algorithm. Nonlocal means [3] and its directional derivative.

Inputs: noisy image y , direction d , noise standard-deviation σ
Parameters: half search window width s , half patch width b , kernel function φ
Outputs: $x^* = \hat{x}(y)$ and $\tilde{x} = J_{\hat{x}}(y)d$

Initialize $W \leftarrow \varphi(2\sigma^2(2b+1)^2) 1_{p_1 \times p_2}$ (add weights for the central pixels [40])
Initialize $W_y \leftarrow \varphi(2\sigma^2(2b+1)^2) y$ (accumulators for the weighted sum)
Initialize $W' \leftarrow 0_{p_1 \times p_2}$
Initialize $W'_y \leftarrow \varphi(2\sigma^2(2b+1)^2) d$
for $k \in [-s, s]^2 \setminus \{0, 0\}$ **do**
 Compute $e \leftarrow [(y - S_k(y))^2] \star \kappa$ (error between each k shifted patches [12, 13])
 Compute $w \leftarrow \varphi(e) \star \kappa$ (contribution for each patch of its k shift)
 Update $W \leftarrow W + w$ (add weights at each position)
 Update $W_y \leftarrow W_y + wS_k(y)$ (add contribution of each k shifted patches)

 Compute $e' \leftarrow [2(y - S_k(y))(d - S_k(d))] \star \kappa$
 Compute $w' \leftarrow [e' \varphi'(e)] \star \kappa$
 Update $W' \leftarrow W' + w'$
 Update $W'_y \leftarrow W'_y + w'S_k(f) + wS_k(d)$
end for
Compute $x^* \leftarrow W_y/W$ (weighted mean)
Compute $\tilde{x} \leftarrow (W'_y - W'x^*)/W$

Figure 6. Pseudoalgorithm for the computation of the nonlocal means and its Jacobian in a direction d . All arithmetic operations are elementwise, S_k is the operator that shifts all pixels in the direction k , \star is the discrete convolution operator, and $\kappa \in \mathbb{R}^{p_1 \times p_2}$ is such that $\kappa_{i,j} = 1$ if $(i, j) \in [-b, b]^2$ and 0 otherwise.

after refitting, some small residuals around the edges appear. In fact, in the vicinity of edges, iso-TV finds (barely visible) discontinuities that are not in accordance with the underlying image. This creates an overload of small constant regions. When refitting is performed, such regions are refitted to the noisy data, and they become barely visible artifacts. In other words, the refitting has re-enforced the presence of a modeling problem, resulting in an increase of residual variance for which iso-TV had originally compensated by attenuating the amplitudes.

6.2. Demasking with isotropic total variation (iso-TV). Figure 8 gives another illustration of our covariant refitting of the 2D iso-TV used for the restoration of an approximate 8-bit piecewise constant image damaged by AWGN with $\sigma = 20$, known as *Boat*. The operator Φ is a random mask removing 40% of pixels. Again, iso-TV introduces a significant loss of contrast, typically for thin details such as the contours of the objects, which are re-enhanced in our refitting result.

Inspecting the map of residuals in Figure 8(e)–(f) illustrates that our refitting technique eliminates most of the bias at the price of a small variance increase. This is clearer when looking at the mast and the ropes. While the mast was preserved by iso-TV and re-enhanced by our refitting, the ropes remain lost for both methods.




Figure 7. (a) Noisy $y = x_0 + w$ and (d) noise-free x_0 . (b) Iso-TV $\hat{x}(y)$ with $\lambda = 3$ and (e) the residual $\hat{x}(y) - x_0$. (c) Our refitting $\mathcal{R}_{\hat{x}}(y)$ and (f) the residual $\mathcal{R}_{\hat{x}}(y) - x_0$.

6.3. Denoising with nonlocal means. Figure 9 gives an illustration of our refitting procedure for the nonlocal means algorithm used in a denoising problem of the 8-bit image *Pirate*, enjoying many repetitive patterns and damaged by AWGN with $\sigma = 20$. We choose a regularizing kernel $\varphi(\cdot) = \exp(\cdot/h)$, $h > 0$, that leads to strong smoothing in order to highlight the behavior of the refitting. Our refitting technique provides favorable results: many details are enhanced compared to the standard method. This reveals that the nonlocal means is actually able to well capture the repetitions of many patterns, but this information is not used properly to create a satisfying result. The refitting produces a sharper result by enforcing the correct use of all the structures identified by patch comparisons. The maps of residuals in Figure 9(e)–(f) highlight that our refitting technique efficiently suppresses this dull effect while preserving the model originally captured by patch redundancy. Again the suppression of this phenomenon is counterbalanced by an increase in the residual variance which is prominent where the local patch redundancy assumption is violated.

In these examples, the overall residual norm is clearly reduced by the refitting because the amount of reduced bias surpasses the increase of residual variance. This favorable behavior depends on the internal parameters of the original estimator acting on the bias-variance trade-off, as we investigate next.

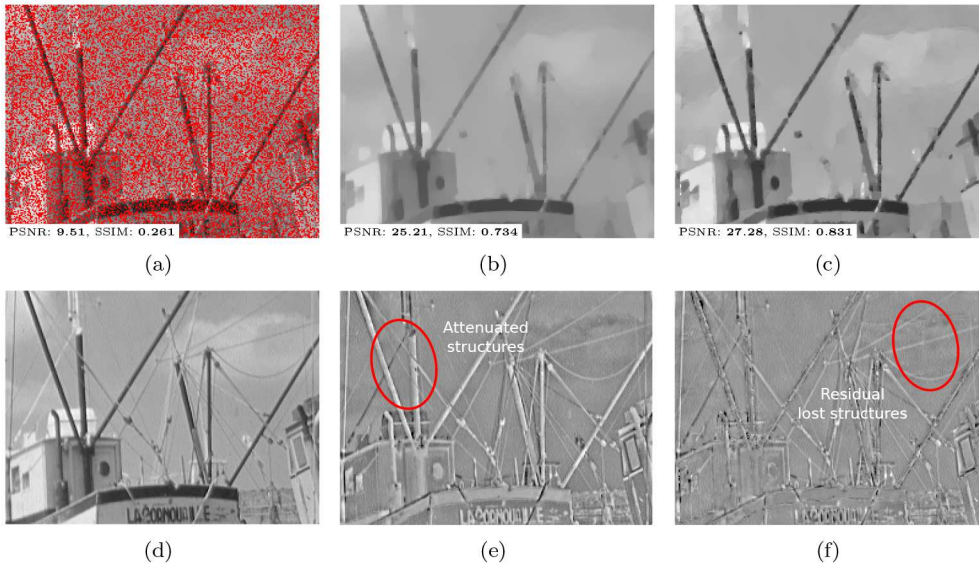


Figure 8. (a) Partial and noisy $y = \Phi x_0 + w$ (red indicates missing pixels) and (d) noise-free x_0 . (b) Iso-TV $\hat{x}(y)$ with $\lambda = 3$ and (e) the residual $\hat{x}(y) - x_0$. (c) Our refitting $\mathcal{R}_{\hat{x}}(y)$ and (f) the residual $\mathcal{R}_{\hat{x}}(y) - x_0$.

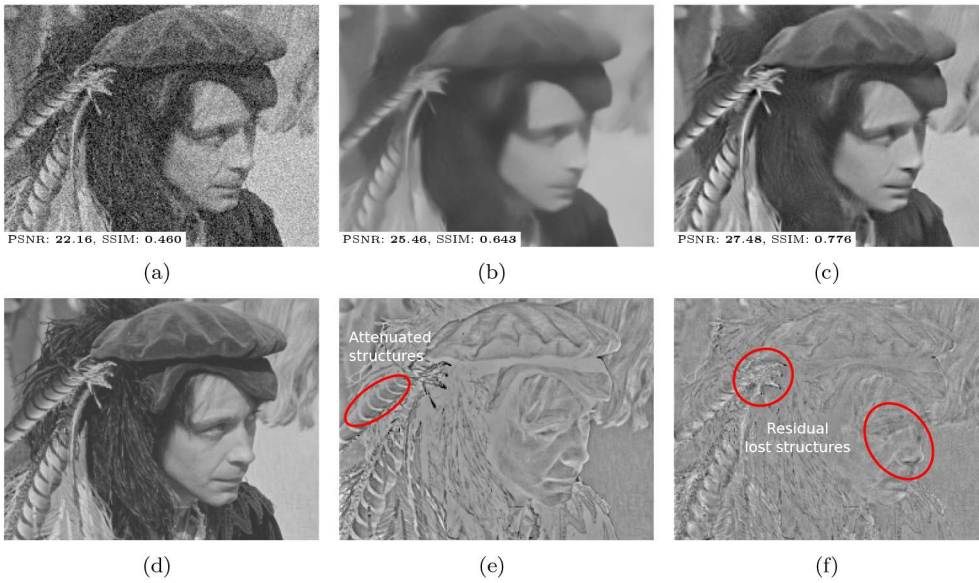


Figure 9. (a) Noisy $y = x_0 + w$ and (d) noise-free x_0 . (b) Nonlocal means $\hat{x}(y)$ with $h = 3$ and (e) the residual $\hat{x}(y) - x_0$. (c) Our refitting $\mathcal{R}_{\hat{x}}(y)$ and (f) the residual $\mathcal{R}_{\hat{x}}(y) - x_0$.

6.4. A bias-variance analysis for the covariant refitting. Previous experiments have revealed that while CLEAR tends to reduce the bias, it increases (as expected) the residual variance. It is therefore important to understand under which conditions the bias-variance trade-off is in favor of our refitting technique.

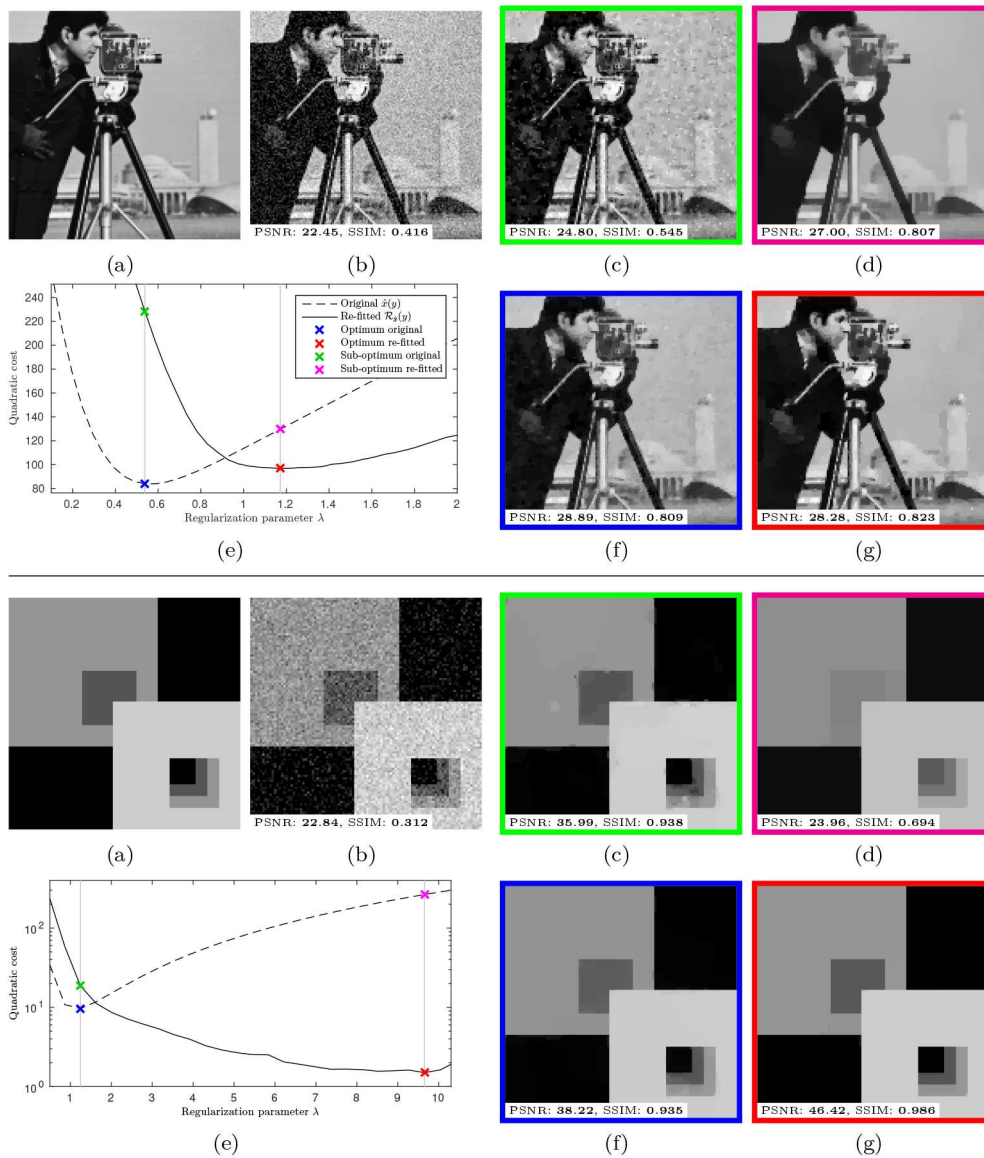


Figure 10. Experiment with aniso-TV: (Top) poorly piecewise constant case. (Bottom) pure piecewise constant case. (a) Noise-free x_0 . (b) Noisy $y = x_0 + w$. (e) MSE of $\hat{x}(y)$ and its refitting $\mathcal{R}_{\hat{x}}(y)$ w.r.t. λ . Two values of λ are selected corresponding to (c) refitting for a suboptimal λ , (d) original for a suboptimal λ , (f) original for the optimal λ , and (g) refitting for the optimal λ .

Figure 10 illustrates the evolution of performance, measured in terms of mean squared error (MSE) of both aniso-TV and its refitting version as a function of the regularization parameter λ . Two images are considered: *Cameraman*, an approximate piecewise constant image (top), and a truly piecewise constant image (bottom).

This experiment highlights that optimal results for both approaches are not reached at the same λ value. Visual inspection of the optima shows that due to the bias, the optimal solution

of aniso-TV is reached for a λ value promoting a model subspace that is not in accordance with the underlying signal: typically the presence of an overload of (barely visible) transitions in homogeneous areas. These transitions become clear when looking at the refitted version where each small region is refitted on the noisy data, revealing an excessive residual variance. Conversely, the optimal λ value for the refitting seems to retrieve the correct model, i.e., with transitions that are closely in accordance with the underlying signal. Comparing their relative performance, when both are used at their own optimal λ , reveals that our refitting brings a significant improvement if the underlying image is in fact piecewise constant.

Figure 11 provides a similar illustration of the evolution of performance for the nonlocal means and its refitted version as a function of the smoothing parameter h of the kernel function $\varphi(\cdot) = \exp(\cdot/h)$. Two images are considered: *Lady*, a crude approximation of an image with redundant patterns (top), and *Fingerprint*, an image with many redundant patterns (bottom). Similar conclusions can be made from this experiment. In particular, comparing their relative performance, when both are used at their own optimal h value, seems to demonstrate that the refitting brings an improvement when most patches of the underlying image are redundant.

While it is difficult to make a general statement, we can reasonably claim from these experiments that refitting is all the more relevant in terms of MSE than the underlying image x_0 is in agreement with the retrieved subspace model $\mathcal{M}_{\hat{x}}(y)$. In other words, refitting is safe when the original restoration technique was chosen appropriately w.r.t. the underlying image of interest. Beyond MSE performance, the refitted results at their optimal parameter choices might nevertheless be preferable (as assessed by the SSIM values): intensities and contrasts are recovered better.

6.5. Behavior on more sophisticated filters. We focus here on two filters, BM3D [11] and DDID [27], reaching state-of-the-art results in denoising. As mentioned in the previous paragraph, refitting is all the more relevant in that the first estimate is obtained with high smoothing strength. It is thus important to compare the original filter with its refitted version for varying smoothing parameters. For simplicity, we have considered only one smoothing parameter $\gamma > 0$ for these two algorithms.

For DDID, we use the authors' implementation and choose to multiply by γ the two inner parameters γ_f and γ_r [27]. For BM3D, we use the implementation of [28] and choose to multiply by γ the inner parameters $\lambda_{3D}^{\text{hard}}$ (in the hard-thresholding step) and σ (in the Wiener filtering step). Unlike previous experiments, the refitted results are obtained here by finite difference, as discussed in subsection 5.2.

Figure 12 illustrates the evolution of performance, measured in terms of MSE, of BM3D, DDID, and their refitted versions as a function of the smoothing parameter γ . BM3D is studied on the *Monarch* image and DDID on the *Lena* image. Similar conclusions can be made from this experiment; refitting reaches its optimal performance for a larger smoothing parameter. Because the original estimator is nearly unbiased, refitting becomes challenging only when the original solution is oversmoothed; otherwise the gain in terms of bias is too small to compensate for the loss in terms of variance. Given this high smoothing strength, some tiny structures have been lost, and thus the optimal refitting does not reach as good a performance as the optimal original solution.

However, even though the MSE is not necessarily improved, our refitting solutions present

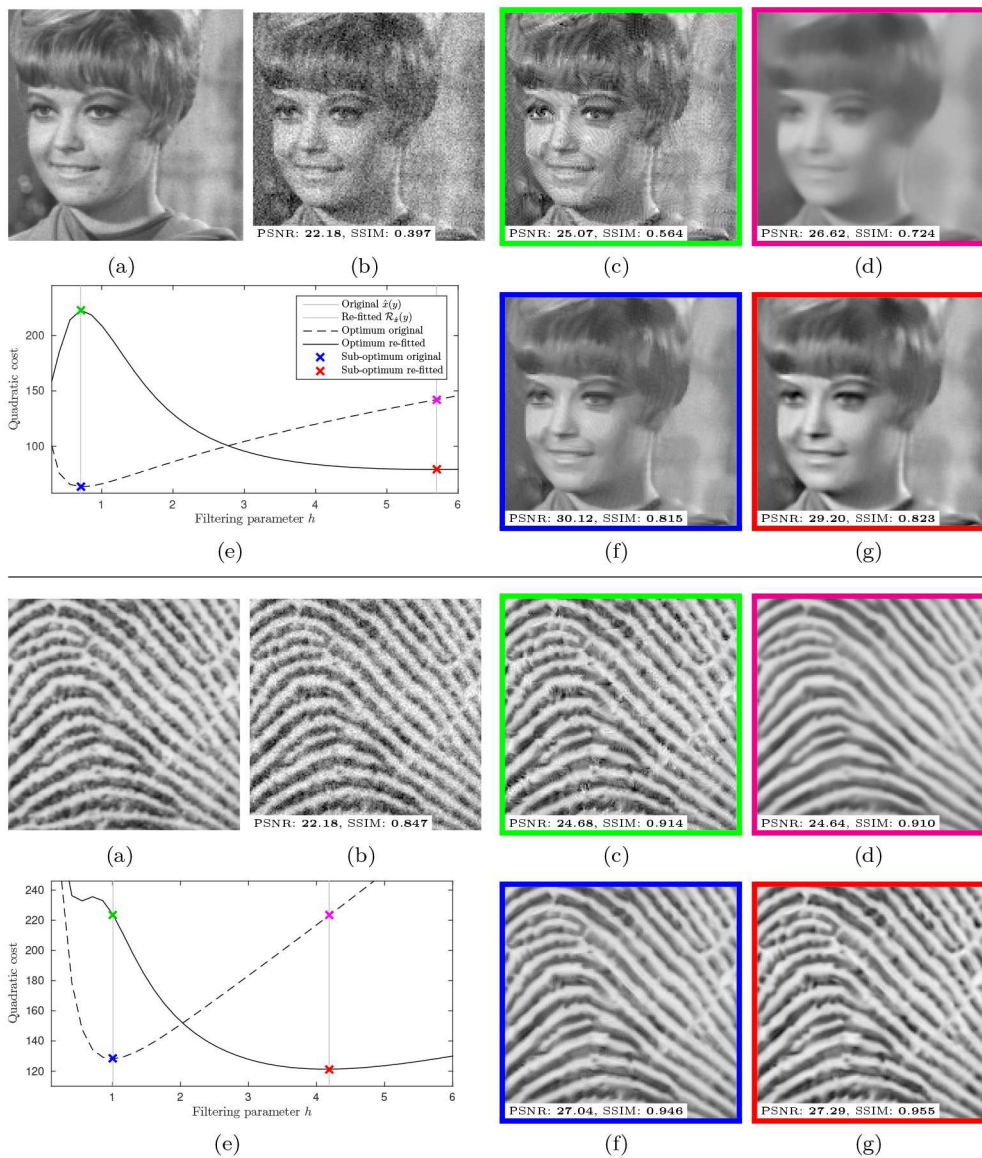


Figure 11. Experiment with nonlocal means: (Top) moderate patch redundancy case. (Bottom) high patch redundancy case. (a) Noise-free x_0 . (b) Noisy $y = x_0 + w$. (e) MSE of $\hat{x}(y)$ and its refitting $\mathcal{R}_{\hat{x}}(y)$ w.r.t. h . Two values of h are selected corresponding to (c) refitting for a suboptimal h , (d) original for a suboptimal h , (f) original for the optimal h , and (g) refitting for the optimal h .

fewer artifacts (as confirmed by the very small gap in terms of SSIM values); see, for instance, the stripes of the *Monarch* or the left cheek of *Lena*. In fact, in order to recover details with low signal-to-noise ratio, the optimal original estimators authorize the apparition of low-contrasted oscillating features. Nevertheless, a few of these oscillations tend to amplify some noise structures and hence explain these artifacts. In contrast, the optimal refitting prefers losing such details rather than taking the risk of creating arbitrary structures and is thus more

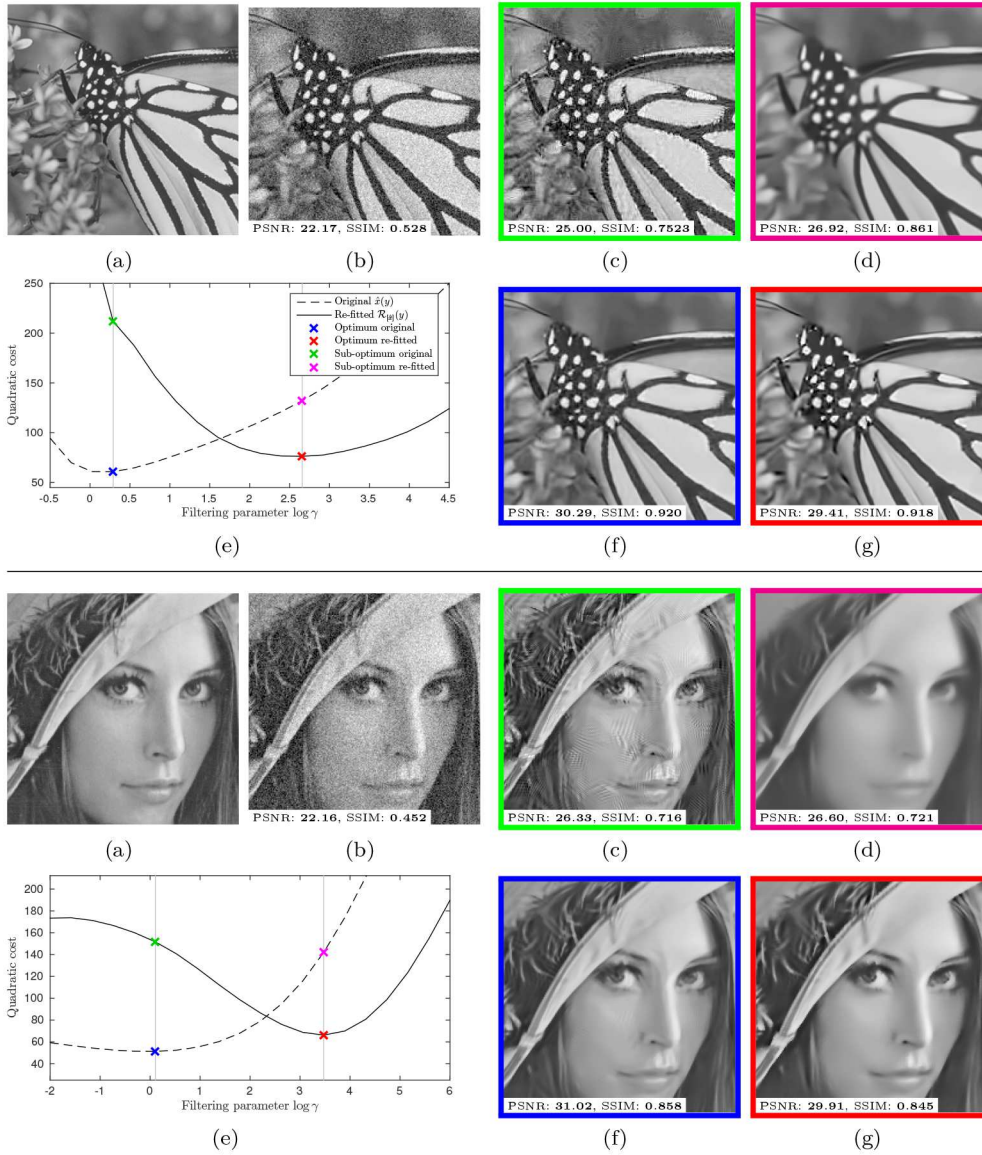


Figure 12. (Top) Experiments with BM3D [11, 28] and with DDID [27] (bottom). (a) Noise-free x_0 . (b) Noisy $y = x_0 + w$. (e) MSE of $\hat{x}(y)$ and its refitting $\mathcal{R}_{\hat{x}}(y)$ w.r.t. γ . Two values of γ are selected corresponding to (c) refitting for a suboptimal γ , (d) original for a suboptimal γ , (f) original for its optimal γ , and (g) refitting for its optimal γ .

reliable.

Regarding our previous discussion, we believe that refitting would be nevertheless beneficial in terms of MSE for the class of images that are promoted by such restoration techniques. Characterizing this class of solutions for these methods is a very challenging topic which is outside the scope of this paper.

6.6. Comparisons with other techniques devoted to the ℓ_1 case. We detail hereafter two different alternative strategies devoted to re-enhancing the solution of the ℓ_1 analysis regularization.

Iterative hard-thresholding. As shown earlier, hard-thresholding is the refitted version of soft-thresholding. Given an iterative solver $(k, y) \mapsto x^k$ composed of linear and soft-thresholding (such as primal-dual algorithms), one could consider replacing all soft-thresholding by hard-thresholding while keeping linearities unchanged—a technique often referred to as “iterative hard-thresholding” [2]. Unfortunately, such techniques only provide convergence to a local minimum of the ℓ_0 -regularized problem, and they do not converge to the sought refitting $\mathcal{R}_{\hat{x}}(y)$.

Cosupport identification based post-refitting. Another solution, referred to as post-refitting and studied in, e.g., [20, 37, 1, 29, 1], consists in identifying the (co)support $\mathcal{I} = \{i : (\Gamma\hat{x}(y))_i \neq 0\}$ and solving a least-square problem constrained to $\{x : (\Gamma x)_{\mathcal{I}^c} = 0\}$, typically with conjugate gradient descent. However, $\hat{x}(y)$ is usually obtained thanks to a converging sequence x^k , and unfortunately $\text{supp}(\Gamma x^k)$ can be far from $\text{supp}(\Gamma\hat{x}(y))$ even though x^k can be made arbitrarily close to $\hat{x}(y)$. Such erroneous support identifications lead to results that strongly deviate from $\mathcal{R}_{\hat{x}}(y)$.

Figure 13 provides a comparison of our refitting method with the two other approaches mentioned earlier for aniso-TV used on an 8-bit image (*peppers*) damaged by AWGN with $\sigma = 20$. The iterative hard-thresholding approach does not preserve the model space: transitions are not localized at the same positions as in the original solution, and suspicious oscillations are created. Post-refitting and our approach have both improved $\hat{x}(y)$ by enhancing each piece and preserving the location of transitions. Our method is nevertheless more stable than support identification, which produces many errors due to wrong cosupport identification.

Figure 14 provides another illustration highlighting the problem of support identification in an ill-posed problem. It consists of an 8-bit image (*Cameraman*) damaged by a Gaussian blur of 2px and AWGN with $\sigma = 20$. Again, while aniso-TV reduces the contrast, the refitting recovers the original amplitudes and leaves discontinuities unchanged. Post-refitting offers results comparable to ours except for suspicious oscillations due to wrong cosupport identification.

In contrast to the support identification, CLEAR requires neither the identification of the cosupport nor the identification of the model subspace $\mathcal{M}_{\hat{x}}(y)$. This is appealing since the cosupport of $\hat{x}(y)$ is difficult to identify, particularly in the analysis context. Being computed along the iterations of the original algorithm, i.e., jointly with $\hat{x}(y)$, our refitting strategy also provides more stable solutions.

6.7. Comparisons with boosting strategies. We detail hereafter other popular alternatives designed to re-enhance results of an arbitrary estimator.

Twicing and boosting. Boosting iterations, introduced in [4], is a simple approach that consists in reinjecting into the current solution \tilde{x}^k a filtered version of its residual $y - \Phi\tilde{x}^k$. The idea is that if parts of the signal were lost at iteration k , they might be retrieved in the residual. Given $\tilde{x}^0 = 0$, the iterations reads

$$(51) \quad \tilde{x}^{k+1} = \tilde{x}^k + \hat{x}(y - \Phi\tilde{x}^k).$$




Figure 13. (a) Noise-free x_0 . (b) Noisy $y = x_0 + w$. (c) Original aniso-TV $\hat{x}(y)$ with $\lambda = 1.2$. Enhanced results by (d) iterative hard-thresholding, (e) post-refitting with support identification, and (f) our proposed covariant refitting.




Figure 14. (a) Original x_0 . (b) Blurred and noisy $y = \Phi x_0 + w$. (c) Original aniso-TV $\hat{x}(y)$ with $\lambda = 0.7$. Enhanced results by (d) iterative hard-thresholding, (e) post-refitting with support identification, and (f) our proposed covariant refitting.

The first iterate is $\tilde{x}^1 = \hat{x}(y)$, and \tilde{x}^2 is known as the twicing estimate [48]. Such approaches are popular in nonparametric statistics, e.g., for kernel smoothing [10].

When k increases, its bias tends to decrease, while its variance increases; see [44]. In denoising (i.e., when $\Phi = \text{Id}$) with a linear estimator $\hat{x}(y) = Wy$ (e.g., the Tikhonov regularization), we get $\tilde{x}^k = (\text{Id} - (\text{Id} - W)^k)y$ for $k > 0$. In particular, the twicing reads as $\tilde{x}^2 = (2W - W^2)y$. In this linear case, the covariant refitting reads as $\mathcal{R}_{\hat{x}}(y) = (W + \rho W - \rho W^2)y$ and coincides with the twicing when $\rho = 1$.

Unsharp residual iteration (URI) is a related approach studied in [8] that reads as $\tilde{x}^{k+1} =$

$\tilde{x}^k + \hat{x}(y) - \hat{x}(\Phi\tilde{x}^k)$ and that is also equivalent in the linear case.

Iterative Bregman refinement. In [34], the authors proposed an iterative procedure, originally designed to improve iso-TV results, given by

$$(52) \quad \tilde{x}^{k+1} = \hat{x} \left(y + \sum_{i=1}^k (y - \Phi\tilde{x}^i) \right).$$

Unlike boosting, which iteratively filters the residual, the idea is to filter a modified version of the input y amplified by adding the sum of the residuals. When $\Phi = \text{Id}$ and $\hat{x}(y) = Wy$, the iterative Bregman refinement reads as $\tilde{x}^k = (\text{Id} - (\text{Id} - W)^k)y$ and coincides with boosting (we refer the reader to [5, 53, 24, 35] for related approaches).

SOS-boosting. In [38], the authors follow a similar idea by iteratively filtering a strengthened version of the input y . Their method, named Strengthen Operate Subtract boosting (SOS-boosting), originally proposed for $\Phi = \text{Id}$, iteratively performs the following update:

$$(53) \quad \tilde{x}^{k+1} = \tau\hat{x}(y + \alpha\Phi\tilde{x}^k) - (\tau\alpha + \tau - 1)\tilde{x}^k,$$

where α and τ are two real parameters. The first controls the emphasis of the solution (and the convergence of the sequence), while the second controls the rate of convergence. When $\Phi = \text{Id}$ and $\hat{x}(y) = Wy$, the SOS refinement with $\tau = 1$ reads as $\tilde{x}^k = Wy + \alpha(W - \text{Id})\tilde{x}^{k-1}$, and in particular, for $k = 2$, we get $\tilde{x}^2 = (W - \alpha W + \alpha W^2)y$, which coincides with our covariant refitting for the choice $\alpha = -\rho$. For all considered estimators, we have always observed $\rho > 0$, contrarily to [38], where $\alpha > 0$ is implicitly assumed. Hence, we cannot conclude that the two models match in a specific linear setting. Another difference is that while we provide an automatic way to compute ρ (see (43)), the α parameter of the SOS-boosting must be tuned by the practitioner, a possibly cumbersome task, e.g., when using cross-validation on a fixed dataset of images and/or for varying noise levels.

SAIF-boosting. As described in [31], the diffusion of a filter consists in iteratively reapplying the filter to the current estimate $\tilde{x}^{k+1} = \hat{x}(\tilde{x}^k)$. The authors of [44] noticed that, unlike the boosting method of [4], the bias of this estimator increases and its variance decreases with k . As a consequence, the authors suggest mixing the two approaches by deciding at each iteration between performing a boosting or a diffusion step. To that end, they proposed a plug-in risk estimator that crudely estimates the MSE from a prefiltered version of y . This approach is in fact applied locally on image patches and is referred to as spatially adapted iterative filter boosting (SAIF-boosting). Unlike the other techniques, SAIF-boosting cannot be used as a black-box. Indeed, it requires performing an eigendecomposition of \hat{x} locally for each patch of y . This can be efficiently done for some kernel-based averaging filters but can be very challenging for arbitrary estimators, such as, for instance, iso-TV.

Though we have compared the expressions of boosting approaches with our refitting in the case of linear estimators, it is worth mentioning that boosting approaches are scarcely used in this case. Boosting appears more relevant in the nonlinear case, since the successive reapplication of a nonlinear estimator \hat{x} allows us to recover parts of the signal that were lost at former iterations. Nevertheless, to boost the solution, the internal parameters of \hat{x} may need to be readapted at each iteration, leading to cumbersome tuning of parameters in

practice. Unlike boosting methods, a refitting approach should not modify the regularity and the structure of the first estimate. This is why CLEAR only considers the linearization of \hat{x} at y through the Jacobian.

It is important to have in mind that most of the theoretical results regarding boosting methods are often well grounded in the case where, even though \hat{x} is nonlinear, it acts locally as an averaging filter. In other words, locally, there exists a row stochastic linear operator W , i.e., $W1_n = 1_n$ (or even bistochastic, symmetric or independent of y), such that $\hat{x}(y) = Wy$. Such theoretical results could not be applied to soft-thresholding (or to more advanced methods we have considered). Indeed, for soft-thresholding, a candidate for W is the diagonal matrix defined as

$$(54) \quad W_{ii} = \begin{cases} 1 - \frac{\lambda}{|y_i|} & \text{if } |y_i| > \lambda, \\ 0 & \text{otherwise,} \end{cases}$$

which is not row stochastic. In this context, the matrix W is not the Jacobian, which is given in this case by the diagonal matrix

$$(55) \quad J_{ii} = \begin{cases} 1 & \text{if } |y_i| > \lambda, \\ 0 & \text{otherwise,} \end{cases}$$

which, unlike W , is row stochastic and locally a projector. A second limitation is that even though W is row stochastic, it might still encode a bias part. Typically, for the ℓ_1 analysis described in (5), a candidate for W is

$$(56) \quad W = U(\Phi U)^+ - U(U^\top \Phi^\top \Phi U)^{-1} U^\top (\Gamma^\top)_{\mathcal{I}} R,$$

where R is a diagonal matrix with diagonal elements $R_{ii} = \lambda/|y_i|$ for $i \in \mathcal{I}$ and 0 otherwise. One can check that if $1_n \in \text{Ker}[\Gamma]$ (which holds true for aniso-TV), then W is row stochastic. However, as seen in (5), the quantity $U(U^\top \Phi^\top \Phi U)^{-1} U^\top (\Gamma^\top)_{\mathcal{I}} R$ is the term responsible for the systematic contraction of the ℓ_1 analysis regularization (this simplifies to $\lambda/|y_i|$ for soft-thresholding). As a consequence, the bias cannot be corrected by a single application of W . The Jacobian $J = U(\Phi U)^+$, which is again row stochastic, is free of this contraction term. Therefore, our covariant refitting gets rid of this bias term after one single application of the Jacobian J .

Figures 15 and 16 provide a comparison of our covariant refitting with the three first aforementioned boosting approaches on two 8-bit images (*SitaHanuBanana* and *Barbara*) damaged by AWGN with $\sigma = 20$, respectively, for iso-TV and for nonlocal means. Note that the inner parameters of iso-TV and nonlocal means are fixed for all approaches. The α and τ parameters of the SOS-boosting approach have been tuned to offer the most satisfying results, even though we did not observe a significant impact in the iso-TV case. As expected, our covariant refitting provides results re-enhanced toward the amplitudes of the noisy inputs. In contrast, boosting approaches do not systematically refit toward the original amplitudes. While CLEAR preserves the structural content and smoothness of the original solution, the boosting approaches reinject structural contents that were not originally preserved and can present a large amount of residual noise. In these experiments, the original estimators were significantly biased, but

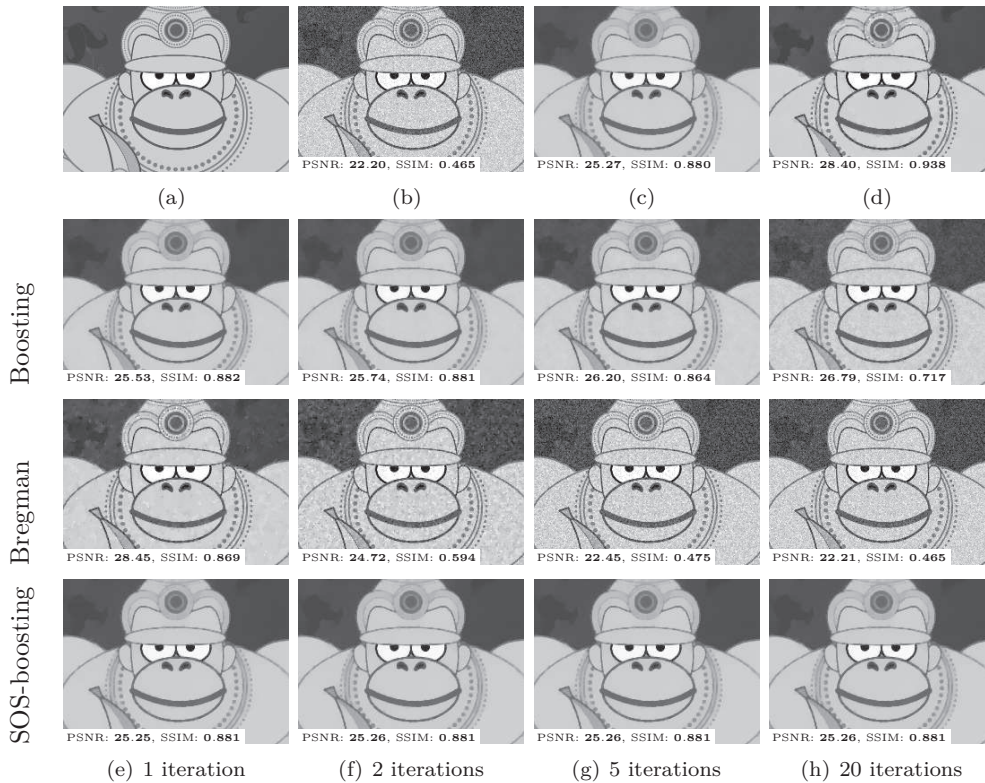


Figure 15. (a) Cartoon image x_0 . (b) Noisy version $y = x_0 + w$. (c) Iso-TV $\hat{x}(y)$ with $\lambda = 1.5$. (d) Our covariant refitting $\mathcal{R}_{\hat{x}}(y)$. (e)–(h) From top to bottom, boosting [48, 4], Bregman iterations [5], and SOS-boosting [38] at, respectively, 1, 2, 5, and 20 iterations.

for smaller parameter choices, reinjecting structural contents would have improved the PSNR more than refitting the amplitudes only. In fact, boosting techniques are more relevant to improving the quality of near unbiased estimates, while refitting techniques correct only for the inaccuracy of biased but precise estimates (i.e., with low estimation variance).

Figure 17 illustrates this point on an 8-bit image (*House*) damaged by AWGN with $\sigma = 20$ and $\sigma = 60$ for iso-TV. Unlike the previous comparison, the approaches are compared for different values of the regularization parameter λ . We have chosen to perform one single iteration for all boosting methods such that they share the same computation time with our refitting. Studying the MSE as a function of λ reveals that SOS-boosting and twicing indeed improve on the MSE for lower values of the regularization parameter (i.e., for nearly unbiased estimators). The Bregman iteration and our refitting improve on the biased estimate for higher values of the regularization parameter λ . The optimal refitting has higher MSE than the biased estimate and all other boosting approaches (tuned at their own optimal λ value). This was expected, as Figures 10, 11, and 12 have already shown that the MSE of the optimal original estimator is better than that of the optimal CLEAR. As a consequence, any “boosting” approach that barely modifies the original estimator will necessarily present lower MSE than CLEAR, in particular for twicing and SOS-boosting in this context. However, the



Figure 16. (a) Image x_0 with moderate patch redundancy. (b) Noisy version $y = x_0 + w$. (c) Nonlocal means $\hat{x}(y)$ with $h = 5.5$. (d) Covariant refitting $\mathcal{R}_{\hat{x}}(y)$. (e)–(h) From top to bottom, boosting [48, 4], Bregman iterations [5], and SOS-boosting [38] at, respectively, 1, 2, 5, and 20 iterations.

SSIM values are very close for all compared methods.

Quantifying the quality of a debiasing is a difficult task. The main issue is that the noisy image itself has a null bias, and any method that overfits to the noise tends to have a small bias too. Then debiasing cannot be defined only in terms of bias reduction; otherwise $\hat{x}(y) = \Phi^+ y$ would be an optimal rule. Debiasing should be defined as a technique minimizing the bias while preserving properties of the original solution, e.g., the model subspace (invariant refitting), or the cross-correlations (covariant refitting). Because of the lack of such a quantitative criterion, we have displayed the results obtained by each method at their optimal λ value for visual comparisons. At a relatively low noise level $\sigma = 20$, all approaches are on par. For $\sigma = 60$, fundamental differences are revealed. SOS-boosting and twicing appear close to the original biased estimate, while the Bregman iteration and our refitting lead to a significant contrast re-enhancement with larger homogeneous regions.

While in practice the results obtained by the Bregman iteration may look preferable, our refitting has a methodological benefit. If the user chooses a biased estimator for which theoretical guarantees on the model space have been proven (e.g., Lasso, total variation, etc.), our CLEAR solution enjoys the same properties, whereas they can be lost with boosting. For instance, for total variation, one can characterize the preserved/lost objects w.r.t. their

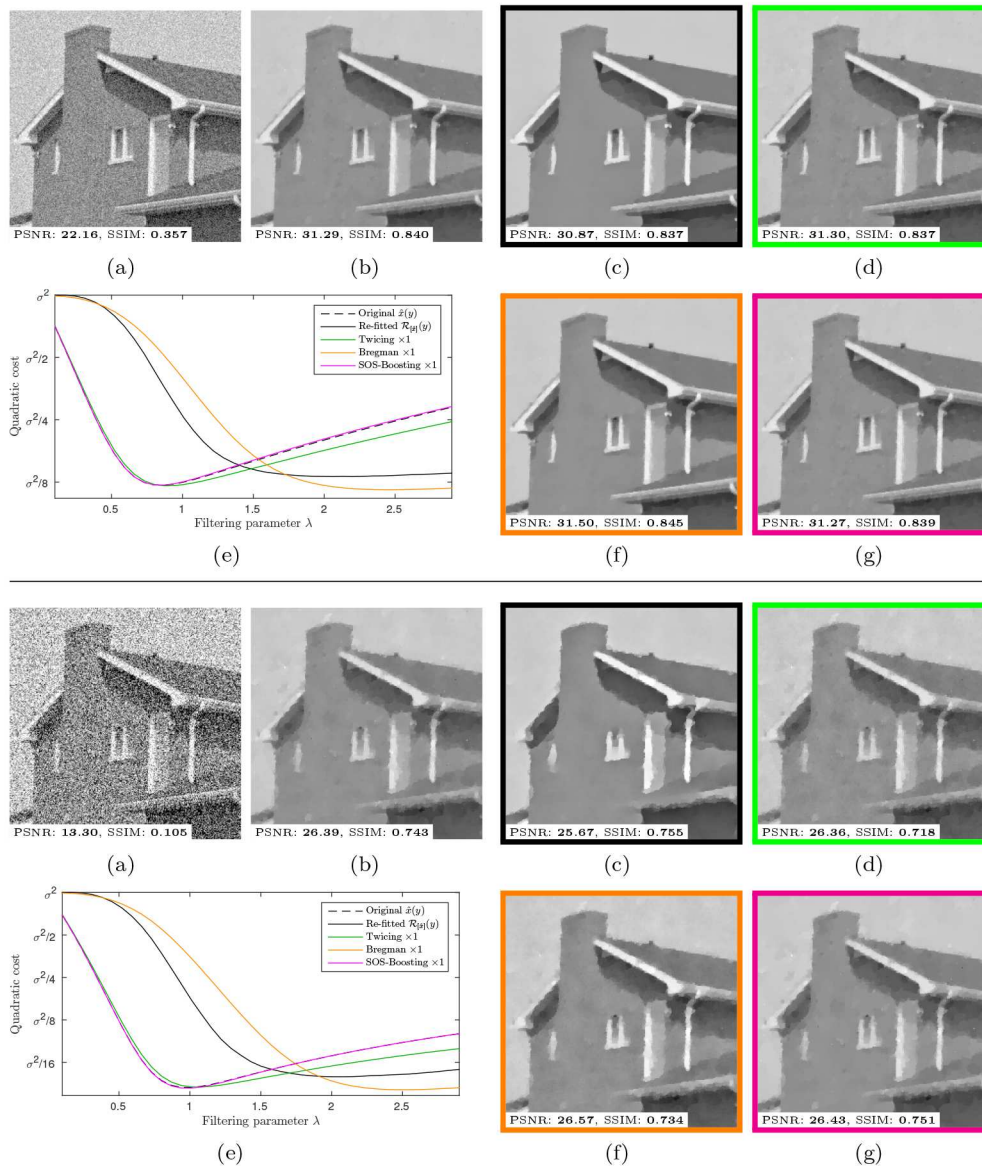


Figure 17. Comparison of refitting and boosting with iso-TV for varying parameter value λ in AWGN contexts with (top) standard deviation $\sigma = 20$ and (bottom) $\sigma = 60$. (a) Noisy image $y = x_0 + w$. (b) Original estimate $x(y)$. (c) Our refitting $\hat{x}(y)$. (d) Twicing. (f) One Bregman iteration. (g) SOS-boosting with one iteration. (b)–(d), (f)–(g) All of these results are obtained at their own respective optimal λ values. (e) MSE of $x(y)$, refitting $\hat{x}(y)$, and the three boosting approaches.

perimeter, area, and the chosen λ . Because the model subspace is preserved, the same objects can be preserved/lost after applying CLEAR. This cannot be guaranteed for boosting approaches even though they can offer a lower approximation error. This behavior in terms of MSE is well known for Lasso (for instance, soft-thresholding is known to achieve lower MSE

than hard-thresholding), but even still refitting is commonly performed because it preserves the guarantees proved for the original estimator while reaching lower bias. Given the significant efforts made by the community to develop theoretical results for solutions of variational problems, we believe it is important that our scheme inherits such guarantees.

7. Conclusion. We have introduced a generalization of the popular least-square refitting, originally aimed at reducing the systematic contraction of Lasso.

Together with this generalization, a generic implementation has been given for a wide class of ill-posed linear problem solvers. This implementation requires a computational overload of at most a factor three compared to the original solver—a factor that can even be reduced to two for most popular estimators in image processing.

While classical refitting is inspired by the standard Lasso debiasing step (i.e., least-square refitting on the estimated support), our generalization leverages the Jacobian of the estimator and does not rely on the notion of support, and its unstable identification. In particular, the proposed implementation only requires chain rules and differentiating the considered solver, and in practice it also has the benefit of increased stability compared to naive implementations.

For estimators such as Tikhonov regularization, total variation, nonlocal means, BM3D, or DDID, numerical experiments have demonstrated the efficiency of the CLEAR technique in retrieving correct intensities while respecting the structure of the original biased estimator. Moreover, it has been shown in practice that refitting is beneficial when the underlying signal structure is well captured by the original estimator. Otherwise, refitting leads to too-simplistic approximations, typically reflecting an inaccurate prior model. In other words, if the considered estimator is adequate w.r.t. the application context, then refitting is recommended.

We have highlighted the importance in distinguishing boosting approaches from the refitting approach. In particular, refitting should be preferred in applications where the content of the original solution must be preserved. While boosting approaches are mostly used to enhance near unbiased estimators (typically coming from combinatorial or nonconvex problems), refitting is all the more relevant for estimators that present biases. For instance, refitting is essential for estimators' solution of a convex problem that requires a large bias correction to accurately retrieve the content of the signal of interest, a canonical example being iso-TV. Nevertheless, we believe that the notion of Jacobian-based refitting could be of interest for boosting applications, and we leave this to future work.

Appendix A. Sketch of proofs. This section details how to retrieve closed form expressions of some of the estimators studied in the paper.

A.1. Retrieving the least-square solution. We aim at retrieving here a minimizer of $\|\Phi x - y\|_2^2 + \iota_C(x)$, where $C = b + \text{Im}[A]$, $b \in \mathbb{R}^p$, and $A \in \mathbb{R}^{p \times n}$. The initial problem can be recast as

$$(57) \quad \underset{x \in C}{\text{argmin}} \|\Phi x - y\|_2^2 = b + A \cdot \underset{t \in \mathbb{R}^n}{\text{argmin}} \|\Phi A t - (y - \Phi b)\|_2^2.$$

The right-hand side problem being differentiable and convex, its first-order optimality conditions give $A^\top \Phi^\top \Phi A t = A^\top \Phi^\top (y - \Phi b)$. In particular, $t = (\Phi A)^+(y - \Phi b)$ is a solution, and hence $x = b + A t$, i.e., $x = b + A(\Phi A)^+(y - \Phi b)$ is a solution.

A.2. Retrieving the Tikhonov solution. We consider the minimization problem, defined for $\Gamma \in \mathbb{R}^{m \times p}$ and $\lambda > 0$, of $\frac{1}{2}\|\Phi x - y\|_2^2 + \frac{\lambda}{2}\|\Gamma x\|_2^2$. The objective being differentiable and convex, its first-order optimality conditions give

$$(58) \quad \Phi^\top(\Phi x - y) + \lambda\Gamma^\top\Gamma x = 0 \Leftrightarrow (\Phi^\top\Phi + \lambda\Gamma^\top\Gamma)x = \Phi^\top y.$$

Provided $\text{Ker } \Phi \cap \text{Ker } \Gamma = \{0\}$ and $\lambda > 0$, the quantity $\Phi^\top\Phi + \lambda\Gamma^\top\Gamma$ is invertible and $x = (\Phi^\top\Phi + \lambda\Gamma^\top\Gamma)^{-1}\Phi^\top y$ is the unique solution.

A.3. Retrieving the hard-thresholding solution. We consider the minimization problem, defined for $\lambda > 0$, of $E(x, y) = \frac{1}{2}\|x - y\|_2^2 + \frac{\lambda^2}{2}\|x\|_0$. The problem is separable, meaning that $[\text{argmin}_{x \in \mathbb{R}^n} E(x, y)]_i = \text{argmin}_{x_i \in \mathbb{R}} E_i(x_i, y_i)$ with

$$(59) \quad E_i(x_i, y_i) = \frac{1}{2} \begin{cases} y_i^2 & \text{if } x_i = 0, \\ (x_i - y_i)^2 + \lambda^2 & \text{otherwise.} \end{cases}$$

Since $y_i^2 \leq \min_{x_i} [(x_i - y_i)^2 + \lambda^2] \Leftrightarrow |y_i| \leq \lambda$, we get

$$(60) \quad \min_{x_i} E_i(x_i, y_i) = \frac{1}{2} \begin{cases} y_i^2 & \text{if } |y_i| \leq \lambda, \\ \lambda^2 & \text{otherwise,} \end{cases}$$

which is reached by setting $x_i = 0$ when $|y_i| \leq \lambda$ and $x_i = y_i$ otherwise.

A.4. Retrieving the soft-thresholding solution. We consider the minimization problem, defined for $\lambda > 0$, of $E(x, y) = \frac{1}{2}\|x - y\|_2^2 + \lambda\|x\|_1$, which is, as for hard-thresholding, separable with $E_i(x_i, y_i) = \frac{1}{2}(x_i - y_i)^2 + \lambda|x_i|$. By convexity, a minimum is reached when zero belongs to its subdifferential. Hence, x_i is solution if

$$(61) \quad 0 \in \partial E_i(x_i, y_i) \Leftrightarrow x_i \in y_i - \lambda \begin{cases} \text{sign}(x_i) & \text{if } |x_i| > 0, \\ [-1, 1] & \text{otherwise,} \end{cases}$$

which holds true by setting $x_i = 0$ when $|y_i| \leq \lambda$ and $x_i = y_i - \lambda \text{sign}(y_i)$ otherwise.

A.5. Retrieving the (blockwise) nonlocal means. We assume periodical boundary conditions such that all quantities q indexed by $i = (i_1, i_2) \in \mathbb{Z}^2$ satisfy $q_i = q_{(i_1, i_2)} = q_{(i_1+k_1p_1, i_2+k_2p_2)}$ for all $(k_1, k_2) \in \mathbb{Z}^2$. This leads, for $l \in [-b, b]^2$, to the following relationship:

$$(62) \quad \begin{aligned} F(x, y) &= \frac{1}{2} \sum_{i,j} w_{i,j} \|\mathcal{P}_i x - \mathcal{P}_j y\|_2^2 = \frac{1}{2} \sum_{i,j} w_{i,j} \sum_l (x_{i+l} - y_{j+l})^2 \\ &= \frac{1}{2} \sum_{i,j} \left[\sum_l w_{i,j} (x_{i+l} - y_{j+l})^2 \right] = \frac{1}{2} \sum_{i,j} \left[\sum_l w_{i-l, j-l} (x_i - y_j)^2 \right] \\ &= \frac{1}{2} \sum_{i,j} \left[\sum_l w_{i-l, j-l} \right] (x_i - y_j)^2 = \frac{1}{2} \sum_{i,j} \bar{w}_{i,j} (x_i - y_j)^2. \end{aligned}$$

For all $i \in [p_1] \times [p_2]$, studying the first optimality conditions gives

$$(63) \quad \frac{\partial F(x, y)}{\partial x_i} = 0 \Leftrightarrow \sum_j \bar{w}_{i,j} (x_i - y_j) = 0 \Leftrightarrow x_i = \frac{\sum_j \bar{w}_{i,j} y_j}{\sum_j \bar{w}_{i,j}}.$$

Appendix B. Proof of Theorem 20. Before turning to the proof of this theorem, let us introduce a lemma.

Lemma 22. *For all y , let $\hat{x}(y)$ be a solution of*

$$(64) \quad \hat{x}(y) \in \underset{x}{\operatorname{argmin}} F(y - \Phi x) + G(x),$$

with F, G two convex functions and G being 1-homogeneous. Then for all $\varepsilon \in [0, 1]$, the following holds:

$$(65) \quad (1 - \varepsilon)\hat{x}(y) \in \underset{x}{\operatorname{argmin}} F(y - \varepsilon\Phi\hat{x}(y) - \Phi x) + G(x).$$

Proof. Note that if G is a convex and 1-homogeneous function, then G is subadditive; i.e., $G(a) + G(b) \geq G(a + b)$. Next, assume that, for some $\varepsilon \in [0, 1]$, (65) does not hold, so that there exists v such that

$$(66) \quad F(y - \varepsilon\Phi\hat{x}(y) - \Phi v) + G(v) < F(y - \varepsilon\Phi\hat{x}(y) - (1 - \varepsilon)\Phi\hat{x}(y)) + G((1 - \varepsilon)\hat{x}(y)) \\ < F(y - \Phi\hat{x}(y)) + (1 - \varepsilon)G(\hat{x}(y)).$$

It follows that

$$(67) \quad F(y - \varepsilon\Phi\hat{x}(y) - \Phi v) + G(v) + \varepsilon G(\hat{x}(y)) < F(y - \hat{x}(y)) + G(\hat{x}(y)).$$

We also have $G(v) + \varepsilon G(\hat{x}(y)) = G(v) + G(\varepsilon\hat{x}(y)) \geq G(v + \varepsilon\hat{x}(y))$, since G is 1-homogeneous and subadditive. Hence, for $w = v + \varepsilon\hat{x}(y)$, we get

$$(68) \quad F(y - \Phi w) + G(w) < F(y - \Phi\hat{x}(y)) + G(\hat{x}(y)),$$

which contradicts $\hat{x}(y) \in \operatorname{argmin} F(y - \Phi x) + G(x)$ and then concludes the proof. \blacksquare

Proof of Theorem 20. By virtue of Lemma 22 and definition of $\hat{x}(y)$, we have $(1 - \varepsilon)\hat{x}(y) = \hat{x}(y - \varepsilon\Phi\hat{x}(y))$ since $\hat{x}(y)$ is supposed to be the unique solution for all y . Now, recall that the linear Jacobian operator applied to $\Phi\hat{x}(y)$ is the directional derivative of $\hat{x}(y)$ in the direction $\Phi\hat{x}(y)$; then, for almost all y , we get

$$(69) \quad J\Phi\hat{x}(y) \triangleq \lim_{\varepsilon \rightarrow 0} \frac{\hat{x}(y) - \hat{x}(y - \varepsilon\Phi\hat{x}(y))}{\varepsilon} = \lim_{\varepsilon \rightarrow 0} \frac{\hat{x}(y) - (1 - \varepsilon)\hat{x}(y)}{\varepsilon} = \hat{x}(y),$$

which concludes the proof. \blacksquare

Appendix C. Proof of Theorem 21. Before turning to the proof of this theorem, let us introduce a lemma.

Lemma 23. *The refitting $\mathcal{R}_{\hat{x}}(y)$ of the ℓ_1 analysis regularization $x^* = \hat{x}(y)$ is the solution of the saddle-point problem*

$$(70) \quad \min_{\hat{x} \in \mathbb{R}^p} \max_{\tilde{z} \in \mathbb{R}^m} \|\Phi\tilde{x} - y\|_2^2 + \langle \Gamma\tilde{x}, \tilde{z} \rangle - \iota_{S_{\mathcal{I}}}(\tilde{z}),$$

where $\iota_{S_{\mathcal{I}}}$ is the indicator function of the convex set $S_{\mathcal{I}} = \{p \in \mathbb{R}^m : p_{\mathcal{I}} = 0\}$.

Proof. As ΦU has full column rank, the refitting of the solution (5) is the unique solution of the constrained least-square problem (see subsection 3.2)

$$(71) \quad \mathcal{R}_{\hat{x}}(y) = U(\Phi U)^+ y = \underset{\tilde{x} \in \mathcal{M}_{\hat{x}}(y)}{\operatorname{argmin}} \|\Phi \tilde{x} - y\|_2^2.$$

We remark that $\tilde{x} \in \mathcal{M}_{\hat{x}}(y) = \operatorname{Ker}[\operatorname{Id}_{\mathcal{I}^c}^t \Gamma] \Leftrightarrow (\Gamma \tilde{x})_{\mathcal{I}^c} = 0 \Leftrightarrow \iota_{S_{\mathcal{I}^c}}(\Gamma \tilde{x}) = 0$, where $S_{\mathcal{I}^c} = \{p \in \mathbb{R}^m : p_{\mathcal{I}^c} = 0\}$.

Using Fenchel transform, $\iota_{S_{\mathcal{I}^c}}(\Gamma \tilde{x}) = \max_{\tilde{z}} \langle \Gamma \tilde{x}, \tilde{z} \rangle - \iota_{S_{\mathcal{I}^c}}^*(\tilde{z})$, where $\iota_{S_{\mathcal{I}^c}}^*$ is the convex conjugate of $\iota_{S_{\mathcal{I}^c}}$. Observing that $\iota_{S_{\mathcal{I}}} = \iota_{S_{\mathcal{I}^c}}^*$ concludes the proof. ■

Given Lemma 23, replacing $\Pi_{z^k + \sigma \Gamma v^k}$ in (47) by the projection onto $S_{\mathcal{I}}$, i.e.,

$$(72) \quad \Pi_{S_{\mathcal{I}}}(\tilde{z})_{\mathcal{I}^c} = \tilde{z}_{\mathcal{I}^c} \quad \text{and} \quad \Pi_{S_{\mathcal{I}}}(\tilde{z})_{\mathcal{I}} = 0,$$

leads to the primal-dual algorithm of [7] applied to problem (70) which converges to the refitted estimator $\mathcal{R}_{\hat{x}}(y)$. It remains to prove that the projection $\Pi_{z^k + \sigma \Gamma v^k}$ defined in (47) converges to $\Pi_{S_{\mathcal{I}}}$ in finite time.

Proof of Theorem 21. First consider $i \in \mathcal{I}$, i.e., $|\Gamma x^*|_i > 0$. By assumption on α , $|\Gamma x^*|_i \geq \alpha > 0$. Necessarily, $z_i^* = \lambda \operatorname{sign}(\Gamma x^*)_i$ in order to maximize (44). Hence, $|z^* + \sigma \Gamma x^*|_i \geq \lambda + \sigma \alpha$. Using the triangle inequality shows that

$$(73) \quad \lambda + \sigma \alpha \leq |z^* + \sigma \Gamma x^*|_i \leq |z^* - z^k|_i + \sigma |\Gamma x^* - \Gamma v^k|_i + |z^k + \sigma \Gamma v^k|_i.$$

Choose $\varepsilon > 0$ sufficiently small such that $\sigma \alpha - \varepsilon(1 + \sigma) > \beta$. From the convergence of the primal-dual algorithm of [7], the sequence (z^k, x^k, v^k) converges to (z^*, x^*, v^*) . Therefore, for k large enough, $|z^* - z^k|_i < \varepsilon$, $|\Gamma x^* - \Gamma v^k|_i < \varepsilon$, and

$$(74) \quad |z^k + \sigma \Gamma v^k|_i \geq \lambda + \sigma \alpha - \varepsilon(1 + \sigma) > \lambda + \beta.$$

Next consider $i \in \mathcal{I}^c$, i.e., $|\Gamma x^*|_i = 0$, where by definition $|z^*|_i \leq \lambda$. Using again the triangle inequality shows that

$$(75) \quad |z^k + \sigma \Gamma v^k|_i \leq |z^k - z^*|_i + \sigma |\Gamma v^k - \Gamma x^*|_i + |z^*|_i.$$

Choose $\varepsilon > 0$ sufficiently small such that $\varepsilon(1 + \sigma) < \beta$. As $(z^k, x^k, v^k) \rightarrow (z^*, x^*, v^*)$, for k large enough, $|z^k - z^*|_i < \varepsilon$, $|\Gamma v^k - \Gamma x^*|_i < \varepsilon$, and

$$(76) \quad |z^k + \sigma \Gamma v^k|_i < \lambda + \varepsilon(1 + \sigma) \leq \lambda + \beta.$$

It follows that for k sufficiently large $|z^k + \sigma \Gamma v^k|_i \leq \lambda + \beta$ if and only if $i \in \mathcal{I}^c$, and hence $\Pi_{z^k + \sigma \Gamma v^k}(\tilde{z}) = \Pi_{S_{\mathcal{I}}}(\tilde{z})$. As a result, all subsequent iterations of (47) will solve (70), and hence from Lemma 23 this concludes the proof of the theorem. ■

REFERENCES

- [1] A. BELLONI AND V. CHERNOZHUKOV, *Least squares after model selection in high-dimensional sparse models*, *Bernoulli*, 19 (2013), pp. 521–547.
- [2] T. BLUMENSATH AND M. E. DAVIES, *Iterative thresholding for sparse approximations*, *J. Fourier Anal. Appl.*, 14 (2008), pp. 629–654.
- [3] A. BUADES, B. COLL, AND J.-M. MOREL, *A review of image denoising algorithms, with a new one*, *Multiscale Model. Simul.*, 4 (2005), pp. 490–530, <https://doi.org/10.1137/040616024>.
- [4] P. BÜHLMANN AND B. YU, *Boosting with the L2 loss: Regression and classification*, *J. Amer. Statist. Assoc.*, 98 (2003), pp. 324–339.
- [5] M. BURGER, G. GILBOA, S. OSHER, AND J. XU, *Nonlinear inverse scale space methods*, *Commun. Math. Sci.*, 4 (2006), pp. 179–212.
- [6] A. CHAMBOLLE, V. DUVAL, G. PEYRÉ, AND C. POON, *Geometric Properties of Solutions to the Total Variation Denoising Problem*, preprint, <https://arxiv.org/abs/1602.00087>, 2016.
- [7] A. CHAMBOLLE AND T. POCK, *A first-order primal-dual algorithm for convex problems with applications to imaging*, *J. Math. Imaging Vis.*, 40 (2011), pp. 120–145.
- [8] M. R. CHAREST AND P. MILANFAR, *On iterative regularization and its application*, *IEEE Trans. Circuits Syst. Video Tech.*, 18 (2008), pp. 406–411.
- [9] P. L. COMBETTES AND J.-C. PESQUET, *A Douglas-Rachford splitting approach to nonsmooth convex variational signal recovery*, *IEEE J. Selected Topics Signal Process.*, 1 (2007), pp. 564–574.
- [10] P.-A. CORNILLON, N. W. HENGARTNER, AND E. MATZNER-LÖBER, *Recursive bias estimation for multivariate regression smoothers*, *ESAIM Probab. Stat.*, 18 (2014), pp. 483–502.
- [11] 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.
- [12] J. DARBON, A. CUNHA, T. F. CHAN, S. OSHER, AND G. J. JENSEN, *Fast nonlocal filtering applied to electron cryomicroscopy*, in *Proceedings of ISBI 2008*, Paris, France, 2008, pp. 1331–1334.
- [13] C.-A. DELEDALLE, V. DUVAL, AND J. SALMON, *Non-local methods with shape-adaptive patches (NLM-SAP)*, *J. Math. Imaging Vis.*, 43 (2012), pp. 103–120.
- [14] C.-A. DELEDALLE, N. PAPADAKIS, AND J. SALMON, *Contrast re-enhancement of total-variation regularization jointly with the Douglas-Rachford iterations*, in *SPARS 2015*, Cambridge, UK, 2015.
- [15] C.-A. DELEDALLE, N. PAPADAKIS, AND J. SALMON, *On debiasing restoration algorithms: Applications to total-variation and nonlocal-means*, in *Scale Space and Variational Methods in Computer Vision*, Lecture Notes in Comput. Sci. 9087, Springer, Berlin, 2015, pp. 129–141.
- [16] C.-A. DELEDALLE, S. VAITER, J. FADILI, AND G. PEYRÉ, *Stein unbiased gradient estimator of the risk (SUGAR) for multiple parameter selection*, *SIAM J. Imaging Sci.*, 7 (2014), pp. 2448–2487, <https://doi.org/10.1137/140968045>.
- [17] D. L. DONOHO AND J. M. JOHNSTONE, *Ideal spatial adaptation by wavelet shrinkage*, *Biometrika*, 81 (1994), pp. 425–455.
- [18] J. DOUGLAS AND H. RACHFORD, *On the numerical solution of heat conduction problems in two and three space variables*, *Trans. Amer. Math. Soc.*, 82 (1956), pp. 421–439.
- [19] V. DUVAL, J.-F. AUJOL, AND Y. GOUSSEAU, *A bias-variance approach for the nonlocal means*, *SIAM J. Imaging Sci.*, 4 (2011), pp. 760–788, <https://doi.org/10.1137/100790902>.
- [20] B. EFRON, T. HASTIE, I. JOHNSTONE, AND R. TIBSHIRANI, *Least angle regression*, *Ann. Statist.*, 32 (2004), pp. 407–499.
- [21] M. ELAD, P. MILANFAR, AND R. RUBINSTEIN, *Analysis versus synthesis in signal priors*, *Inverse Problems*, 23 (2007), pp. 947–968.
- [22] L. C. EVANS AND R. F. GARIEPY, *Measure Theory and Fine Properties of Functions*, CRC Press, Boca Raton, FL, 1992.
- [23] E. I. GEORGE, *Minimax multiple shrinkage estimation*, *Ann. Statist.*, 14 (1986), pp. 188–205.
- [24] G. GILBOA, *A total variation spectral framework for scale and texture analysis*, *SIAM J. Imaging Sci.*, 7 (2014), pp. 1937–1961, <https://doi.org/10.1137/130930704>.
- [25] A. GRIEWANK AND A. WALTHER, *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*, 2nd ed., SIAM, Philadelphia, 2008.

- [26] A. E. HOERL AND R. W. KENNARD, *Ridge regression: Biased estimation for nonorthogonal problems*, *Technometrics*, 12 (1970), pp. 55–67.
- [27] C. KNAUS AND M. ZWICKER, *Dual-domain image denoising*, in *ICIP 2013*, IEEE, Washington, DC, pp. 440–444.
- [28] M. LEBRUN, *An analysis and implementation of the BM3D image denoising method*, *Image Processing On Line*, 2 (2012), pp. 175–213.
- [29] J. LEDERER, *Trust, but Verify: Benefits and Pitfalls of Least-Squares Refitting in High Dimensions*, preprint, <https://arxiv.org/abs/1306.0113>, 2013.
- [30] Y. LIN AND H. H. ZHANG, *Component selection and smoothing in multivariate nonparametric regression*, *Ann. Statist.*, 34 (2006), pp. 2272–2297.
- [31] 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.
- [32] U. NAUMANN, *Optimal Jacobian accumulation is NP-complete*, *Math. Program.*, 112 (2008), pp. 427–441.
- [33] G. R. OBOZINSKI, M. J. WAINWRIGHT, AND M. I. JORDAN, *High-dimensional support union recovery in multivariate regression*, in *NIPS 2008*, Vancouver, Canada, 2008, pp. 1217–1224.
- [34] S. OSHER, M. BURGER, D. GOLDFARB, J. XU, AND W. YIN, *An iterative regularization method for total variation-based image restoration*, *Multiscale Model. Simul.*, 4 (2005), pp. 460–489, <https://doi.org/10.1137/040605412>.
- [35] S. OSHER, F. RUAN, J. XIONG, Y. YAO, AND W. YIN, *Sparse recovery via differential inclusions*, *Appl. Comput. Harmon. Anal.*, 41 (2016), pp. 436–469.
- [36] S. RAMANI, T. BLU, AND M. UNSER, *Monte-Carlo SURE: A black-box optimization of regularization parameters for general denoising algorithms*, *IEEE Trans. Image Process.*, 17 (2008), pp. 1540–1554.
- [37] P. RIGOLLET AND A. B. TSYBAKOV, *Exponential screening and optimal rates of sparse estimation*, *Ann. Statist.*, 39 (2011), pp. 731–471.
- [38] 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>.
- [39] L. I. RUDIN, S. OSHER, AND E. FATEMI, *Nonlinear total variation based noise removal algorithms*, *Phys. D*, 60 (1992), pp. 259–268.
- [40] J. SALMON, *On two parameters for denoising with non-local means*, *IEEE Signal Process. Lett.*, 17 (2010), pp. 269–272.
- [41] X. SHEN AND J. YE, *Adaptive model selection*, *J. Amer. Statist. Assoc.*, 97 (2002), pp. 210–221.
- [42] C. STEIN, *Inadmissibility of the usual estimator for the mean of a multivariate normal distribution*, in *Proceedings of the 3rd Berkeley Symposium on Mathematical Statistics and Probability*, Vol. 1, 1956, pp. 197–206.
- [43] D. STRONG AND T. CHAN, *Edge-preserving and scale-dependent properties of total variation regularization*, *Inverse Problems*, 19 (2003), pp. 165–187.
- [44] H. TALEBI, X. ZHU, AND P. MILANFAR, *How to SAIF-ly boost denoising performance*, *IEEE Trans. Image Process.*, 22 (2013), pp. 1470–1485.
- [45] R. TIBSHIRANI, *Regression shrinkage and selection via the lasso*, *J. Roy Statist. Soc. Ser. B*, 58 (1996), pp. 267–288.
- [46] R. J. TIBSHIRANI, *The lasso problem and uniqueness*, *Electron. J. Statist.*, 7 (2013), pp. 1456–1490.
- [47] A. N. TIKHONOV, *On the stability of inverse problems*, *Dokl. Akad. Nauk SSSR*, 39 (1943), pp. 176–179.
- [48] J. W. TUKEY, *Exploratory Data Analysis*, Pearson, Reading, MA, 1977.
- [49] S. VAITER, C.-A. DELEDALLE, G. PEYRÉ, C. DOSSAL, AND J. FADILI, *Local behavior of sparse analysis regularization: Applications to risk estimation*, *Appl. Comput. Harmon. Anal.*, 35 (2013), pp. 433–451.
- [50] S. VAITER, C.-A. DELEDALLE, J. M. FADILI, G. PEYRÉ, AND C. DOSSAL, *The degrees of freedom of partly smooth regularizers*, *Ann. Inst. Stat. Math.* (2016), <https://link.springer.com/article/10.1007/s10463-016-0563-z>.
- [51] D. VAN DE VILLE AND M. KOCHER, *Nonlocal means with dimensionality reduction and sure-based parameter selection*, *IEEE Trans. Image Process.*, 20 (2011), pp. 2683–2690.
- [52] Z. WANG, A. C. BOVIK, H. R. SHEIKH, AND E. P. SIMONCELLI, *Image quality assessment: From error visibility to structural similarity*, *IEEE Trans. Image Process.*, 13 (2004), pp. 600–612.

- [53] J. XU AND S. OSHER, *Iterative regularization and nonlinear inverse scale space applied to wavelet-based denoising*, IEEE Trans. Image Process., 16 (2007), pp. 534–544.
- [54] J. YE, *On measuring and correcting the effects of data mining and model selection*, J. Amer. Statist. Assoc., 93 (1998), pp. 120–131.
- [55] M. YUAN AND Y. LIN, *Model selection and estimation in regression with grouped variables*, J. R. Stat. Soc. Ser. B Stat. Methodol., 68 (2006), pp. 49–67.