

Total Variation Projection with First Order Schemes

Jalal M. Fadili and Gabriel Peyré

Abstract—This article proposes a new algorithm to compute the projection on the set of images whose total variation is bounded by a constant. The projection is computed through a dual formulation that is solved by first order non-smooth optimization methods. This yields an iterative algorithm that computes iterative soft thresholding of the dual vector fields. This projection algorithm can then be used as a building block in a variety of applications such as solving inverse problems under a total variation constraint, or for texture synthesis. Numerical results show that our algorithm competes favorably with state-of-the-art TV projection methods to solve denoising, texture synthesis, inpainting and deconvolution problems.

Index Terms—Total variation, projection, duality, proximal operator, forward-backward splitting, Nesterov scheme, inverse problems.

I. INTRODUCTION

TOTAL variation is a well known image prior introduced by Rudin, Osher and Fatemi [1]. For a differentiable function $f : \Omega = [0, 1]^2 \rightarrow \mathbb{R}$, it is computed as $\|f\|_{\text{TV}} = \int_{\Omega} |\nabla f|$, and can be extended to the space $\text{BV}([0, 1]^2)$ that contains functions with discontinuities.

The total variation is used as a regularization to denoise an image f_0 by solving the strictly convex problem

$$\min_{f \in \text{BV}([0,1]^2)} \frac{1}{2} \|f - f_0\|^2 + \lambda \|f\|_{\text{TV}}, \quad (1)$$

as originally proposed in [1]. The regularization weight λ should be tuned to match the noise level contaminating f_0 . Several algorithms have been proposed to solve this problem, see for instance [2], [3], [4], [5], [6], [7], [8], [9]. Such primal, dual, or primal-dual schemes for denoising are often a building block for solving more complex inverse problems; see e.g. [10], [11].

TV projection for denoising. Much less work has focused on denoising an image f_0 by projecting it on a total variation ball of radius $\tau < \|f_0\|_{\text{TV}}$, which requires to solve¹

$$\min_{\|f\|_{\text{TV}} \leq \tau} \|f - f_0\|. \quad (2)$$

Such a formulation might be preferable over (1) when little is known about the noise level perturbing f_0 , but when an estimate τ of the total variation of the clean image is known. Computing the solution of (2) with a fast algorithm is thus important for denoising application as has been advocated in [12].

Jalal M. Fadili is with GREYC CNRS-ENSICAEN-Université de Caen, 14050 Caen Cedex, France., email: jalal.fadili@greyc.ensicaen.fr

Gabriel Peyré is with the CEREMADE CNRS-Université Paris-Dauphine, 75775 Paris Cedex 16 France, email: gabriel.peyre@ceremade.dauphine.fr

¹This constraint set is obviously a closed convex set.

An iterative projected sub-gradient method was introduced in [13], [12]. We propose in this paper a different algorithm that is based on a dual regularization of the primal projection problem. This bears similarities with Chambolle's algorithm [6] that solves the primal TV regularization (1) using a dual projection. Our dual problem is then solved using two first-order iterative schemes: *one-step* forward-backward splitting which dates back to [14], [15] and recently revitalized in [16], and accelerated *multi-step* Nesterov scheme [17]. Both of these algorithms require only the computation of a soft thresholding applied to the dual vector fields at each iteration.

Total variation projection also have far-reaching applications beyond denoising. For example, the extraction of Cheeger sets in landslides modeling [18] can be relaxed as a TV projection problem with boundary constraints, that has been recently solved using sub-gradient projection [19].

TV projection for synthesis. Classical synthesis methods constrain wavelet coefficients [20], [21] and are suitable to model some natural phenomena. Computer graphics methods do not rely on statistical modeling and generate a texture through a consistent copy of pixels from an example image [22], [23]. Higher order statistical models [24], [25] improve the visual quality of synthesis by capturing geometric singularities. Section VI-B shows that a total variation constraint can also be used to enhance the sharpness of edges in texture synthesis.

TV projection for inverse problems. Total variation projection might be useful as a proxy to solve more challenging inverse problems. Popular linear inverse problems such as inpainting and deconvolution have been the subject of a flurry of research activity where TV has been extensively used to regularize them.

Classical methods for inpainting use partial differential equations that propagate the information from the boundary of the missing region to its interior, see for instance [26], [27], [28]. Tensor diffusion makes use of a geometric layer to drive the diffusion, see [29]. TV regularization was proposed in [30] for inpainting. Sparsity over a redundant frame has also been successfully used to regularize the inpainting problem [31], [32].

There is an extensive literature on the deconvolution problem image processing, and many algorithms have been developed to tackle it. For instance, TV penalty has been used as a regularization term within a variational framework, see e.g. [33], [10], [11], [34], [9]. Wavelet-based or more generally sparsity-based deconvolution methods have also received considerable attention over the last decade, see e.g. [35], [36], [37], [38].

Inverse problems such as inpainting or deconvolution can be regularized with a total variation constraint [12]. Sections VI-C and VI-D are devoted to show how these TV-constrained

inverse problems can be solved efficiently using a projected gradient descent iteration, whose projector is computed with our algorithm.

II. NOTATION AND ELEMENTS OF CONVEX ANALYSIS

Throughout the paper, an image f of $N = n \times n$ pixels is a vector in \mathbb{R}^N . We denote by $\|\cdot\|$ the norm induced by the inner product $\langle \cdot, \cdot \rangle$ in \mathbb{R}^N . We let $\mathcal{U} = \mathbb{R}^N \times \mathbb{R}^N$ be the space of vector fields with associated inner product $\langle u, v \rangle_{\mathcal{U}} = \sum_{0 \leq i, j \leq n-1} (u_1[i, j]v_1[i, j] + u_2[i, j]v_2[i, j])$, $\forall u = (u_1, u_2)$ and $v = (v_1, v_2) \in \mathcal{U}$. The ℓ^1 and ℓ^∞ norms of a vector field $u = (u_1, u_2) \in \mathcal{U}$ are respectively $\|u\|_1 = \sum_{0 \leq i, j \leq n-1} |u[i, j]|$ and $\|u\|_\infty = \max_{i, j} |u[i, j]|$, where $|u[i, j]| := \sqrt{u_1[i, j]^2 + u_2[i, j]^2}$.

Let C a nonempty convex set. The indicator function ι_C of C is

$$\iota_C(z) = \begin{cases} 0, & \text{if } z \in C, \\ +\infty, & \text{otherwise.} \end{cases}$$

The domain of a function φ is defined by $\text{dom}(\varphi) = \{z : \varphi(z) < +\infty\}$ and φ is proper if $\text{dom}(\varphi) \neq \emptyset$. The conjugate of a proper, lower-semicontinuous and convex function φ is the proper closed convex function φ^* defined by

$$\varphi^*(z) = \sup_{z \in \text{dom}(\varphi)} \langle w, z \rangle - \varphi(z), \quad (3)$$

and we have the bi-conjugate $\varphi^{**} = \varphi$.

The subdifferential of a proper, lower-semicontinuous and convex function φ at z is the set-valued map $\partial\varphi$

$$\partial\varphi(z) = \{w | \forall v, \varphi(v) \geq \varphi(z) + \langle w, v - z \rangle\}. \quad (4)$$

An element w of $\partial\varphi$ is called a subgradient. If φ is Gâteaux-differentiable at z , its only subgradient is its gradient. A function φ is strongly convex with modulus $c > 0$ if and only if

$$\varphi(v) \geq \varphi(z) + \langle w, v - z \rangle + \frac{c}{2} \|v - z\|^2, \forall v. \quad (5)$$

See e.g. [39] for comprehensive account on convex analysis and subdifferential calculus.

We also define the notion of a proximity operator, which was introduced in [40] as a generalization of convex projection operator. For every z , the function $w \mapsto \frac{1}{2} \|w - z\|^2 + \varphi(w)$ achieves its infimum at a unique point denoted by $\text{prox}_\varphi z$. The uniquely-valued operator prox_φ thus defined is the *proximity operator* of φ .

III. DISCRETE DUAL FORMULATION

The optimization (2) is computed numerically by discretizing the gradient operator and the total variation to project an image of $N = n \times n$ pixels.

A. Discrete Total Variation

A discretized gradient for an image $f \in \mathbb{R}^N$ is defined as $\nabla f[i, j] = (\partial_x f[i, j], \partial_y f[i, j])$, where

$$\partial_x f[i, j] = \begin{cases} f[i+1, j] - f[i, j] & \text{if } 0 \leq i < n-1, \\ 0 & \text{otherwise,} \end{cases}$$

$$\partial_y f[i, j] = \begin{cases} f[i, j+1] - f[i, j] & \text{if } 0 \leq j < n-1, \\ 0 & \text{otherwise.} \end{cases}$$

The gradient is thus a vector field $\nabla f \in \mathcal{U}$. The discrete total variation is $\|f\|_{\text{TV}} = \|\nabla f\|_1$ where the ℓ^1 norm of a vector field in \mathcal{U} is defined in Section II.

The adjoint of the gradient is $\nabla^* = -\text{div}$, where the divergence of a vector field $u = (u_1, u_2) \in \mathcal{U}$ is $\text{div}(u) = \partial_x^* u_1 + \partial_y^* u_2$, with

$$\partial_x^* f[i, j] = \begin{cases} f[i, j] - f[i-1, j] & \text{if } 0 < i < n, \\ 0 & \text{otherwise,} \end{cases}$$

$$\partial_y^* f[i, j] = \begin{cases} f[i, j] - f[i, j-1] & \text{if } 0 < j < n, \\ 0 & \text{otherwise.} \end{cases}$$

B. Total Variation Projection

The goal of this paper is to compute the projection f^* of an image f_0 on a set of images with bounded variation

$$f^* = \underset{f \in \mathbb{R}^N, \|f\|_{\text{TV}} \leq \tau}{\text{argmin}} \|f - f_0\|^2. \quad (6)$$

where $0 < \tau < \|f_0\|_{\text{TV}}$ to avoid the trivial solution $f^* = f_0$.

The following proposition shows that the primal constrained optimization problem (6) is recast into a penalized form.

Proposition 1. *For any $f \in \mathbb{R}^N$, the primal solution is recovered as $f^* = f_0 - \text{div}(u^*)$ where u^* is the solution of the dual problem*

$$\inf_{u \in \mathbb{R}^N \times \mathbb{R}^2} \frac{1}{2} \|f_0 - \text{div}(u)\|^2 + \tau \|u\|_\infty. \quad (7)$$

Proof: Let's introduce the dual variable $u \in \mathcal{U}$. First, it is easy to show that the Fenchel conjugate of the ℓ^∞ norm is the indicator function of the ℓ^1 norm. Indeed, by the Hölder inequality we have $\langle u, v \rangle_{\mathcal{U}} \leq \|u\|_\infty \|v\|_1$, we then get

$$\begin{aligned} (\tau \|\cdot\|_\infty)^*(v) &= \sup_{u \in \mathcal{U}} \langle u, v \rangle_{\mathcal{U}} - \tau \|u\|_\infty \\ &= \sup_{\rho > 0} \sup_{\|u\|_\infty = \rho} \langle u, v \rangle_{\mathcal{U}} - \tau \rho \\ &= \sup_{\rho > 0} \rho (\|v\|_1 - \tau) = \iota_{\{\|v\|_1 \leq \tau\}}(v). \end{aligned}$$

Thus, using the bi-conjugate relation, one can write the TV-ball indicator function as

$$\iota_{\{\|v\|_{\text{TV}} \leq \tau\}}(f) = \sup_{u \in \mathcal{U}} \langle u, \nabla f \rangle_{\mathcal{U}} - \tau \|u\|_\infty.$$

This allows to rewrite (6) as

$$\begin{aligned} &\min_{f \in \mathbb{R}^N} \frac{1}{2} \|f - f_0\|^2 + \iota_{\{\|v\|_{\text{TV}} \leq \tau\}}(f) \\ &= \sup_{u \in \mathcal{U}} -\tau \|u\|_\infty + \min_{f \in \mathbb{R}^N} \langle u, \nabla f \rangle_{\mathcal{U}} + \frac{1}{2} \|f - f_0\|^2 \\ &= -\inf_{u \in \mathcal{U}} \tau \|u\|_\infty - \min_{f \in \mathbb{R}^N} -\langle \text{div}(u), f \rangle + \frac{1}{2} \|f - f_0\|^2 \end{aligned} \quad (9)$$

²An alternative proof can be found in e.g. [39, Proposition V.3.2.1].

The inner minimization involves a strongly convex function, whose unique primal solution f^* is recovered from the dual variable u as

$$f^* = \operatorname{argmin}_{f \in \mathbb{R}^N} -\langle \operatorname{div}(u), f \rangle + \frac{1}{2} \|f - f_0\|^2 = f_0 + \operatorname{div}(u), \quad (10)$$

and

$$-\langle \operatorname{div}(u), f^* \rangle + \frac{1}{2} \|f^* - f_0\|^2 = -\frac{1}{2} \|f_0 + \operatorname{div}(u)\|^2 + \frac{1}{2} \|f_0\|^2. \quad (11)$$

Combining (8) and (11) with the obvious change of sign on u leads to the optimization problem (7). ■

Remark 1. Proposition 1 can be alternatively proved using Fenchel-Rockafellar duality formula well-known in convex analysis, see e.g. [41, Section III.4]. We have chosen this line of proof for convenience and the sake of accessibility to the reader.

IV. FIRST-ORDER SCHEMES ON THE DUAL PROBLEM

In going from the primal problem (6) to the dual formulation (7), we have replaced a constrained problem with an unconstrained penalized form. The primal solution is easily recovered from the dual solution as $f^* = f_0 - \operatorname{div}(u^*)$. Furthermore, it turns out that the dual problem (7) is easier to solve with various first order non-smooth optimization schemes. Indeed, (7) involves a quadratic form which has a Lipschitz continuous gradient, and the non-differentiable ℓ^∞ term. Section IV-B presents a *one-step* forward-backward splitting iteration, as explained for instance in [15], [16], while Section IV-C proposes an accelerated *multi-step* scheme due to Nesterov, see [17] and also [7] for some applications to image processing. It turns out that both these algorithms involve the computation of the proximity operator associated to the ℓ^∞ norm that can be computed explicitly as explained in the following section.

A. ℓ^∞ Proximity Operator

Recall from Section II that the proximity operator $\operatorname{prox}_{\kappa \|\cdot\|_\infty}(u)$ associated to $\kappa \|\cdot\|_\infty$ amounts to solving

$$\operatorname{prox}_{\kappa \|\cdot\|_\infty}(u) = \operatorname{argmin}_{v \in \mathcal{U}} \frac{1}{2} \|u - v\|^2 + \kappa \|v\|_\infty. \quad (12)$$

Proposition 2 hereafter shows that $\operatorname{prox}_{\kappa \|\cdot\|_\infty}(u)$ is computed explicitly through a soft thresholding S_λ applied on the dual vector field u for a properly chosen value of λ .

To get the precise value of λ for a given vector field $u \in \mathcal{U}$, we need to compute $d[0] \leq d[1] \leq \dots \leq d[N-1]$ that orders the set of norms

$$\{d[t]\}_{t=0}^{N-1} = \{|u[i, j]|\}_{i, j=0}^{n-1}, \quad (13)$$

and also the cumulated ordered norms

$$D[s] = \sum_{t=s+1}^{N-1} d[t]. \quad (14)$$

Proposition 2. For $u \in \mathcal{U}$, with d and D as defined in (13)-(14), we have $\operatorname{prox}_{\kappa \|\cdot\|_\infty}(u) = 0$ if $\|u\|_1 \leq \kappa$ and $\operatorname{prox}_{\kappa \|\cdot\|_\infty}(u) = u - S_\lambda(u)$ otherwise, where

$$S_\lambda(u)[i, j] = \max\left(1 - \frac{\lambda}{|u[i, j]|}, 0\right) u[i, j] \quad (15)$$

and $\lambda > 0$ is given by

$$\lambda = d[t] + (d[t+1] - d[t]) \frac{D[t+1] - \kappa}{D[t+1] - D[t]} \quad (16)$$

where t is such that $D[t+1] \leq \kappa < D[t]$.

Proof: Using a similar proof to that of Proposition 1, we have the relation between the primal and dual minimization problems

$$\begin{aligned} \operatorname{prox}_{\kappa \|\cdot\|_\infty}(u) &= \operatorname{argmin}_{v \in \mathcal{U}} \frac{1}{2} \|u - v\|^2 + \kappa \|v\|_\infty \\ \iff \operatorname{Proj}_{\{\|\cdot\|_1 \leq \kappa\}}(u) &= \operatorname{argmin}_{w \in \mathcal{U}} \frac{1}{2} \|u - w\|^2 + \iota_{\{\|\cdot\|_1 \leq \kappa\}}(w). \end{aligned}$$

where $\operatorname{Proj}_{\{\|\cdot\|_1 \leq \kappa\}}(u)$ is the orthogonal projection of u onto the closed ℓ^1 ball in \mathcal{U} of radius κ . In the same vein as (10), the relation between the primal and dual solutions (both are unique here) yields

$$\operatorname{prox}_{\kappa \|\cdot\|_\infty}(u) = u - \operatorname{Proj}_{\{\|\cdot\|_1 \leq \kappa\}}(u). \quad (17)$$

If $\|u\|_1 \leq \kappa$, then $\operatorname{prox}_{\kappa \|\cdot\|_\infty}(u) = 0$. Otherwise, the projection $\operatorname{Proj}_{\{\|\cdot\|_1 \leq \kappa\}}(u)$ is computed by finding the Lagrange multiplier $\lambda(\kappa)$ such that³

$$\operatorname{Proj}_{\{\|\cdot\|_1 \leq \kappa\}}(u) = \operatorname{argmin}_v \frac{1}{2} \|u - v\|^2 + \lambda(\kappa) \|v\|_1. \quad (18)$$

As noticed for instance in [42] for wavelet thresholding, the solution of (18) has a closed-form known as soft-thresholding extended to vector fields

$$\operatorname{Proj}_{\{\|\cdot\|_1 \leq \kappa\}}(u) = S_{\lambda(\kappa)}(u).$$

Observe that $\|S_\lambda(u)\|_1 = \sum_{|u[i, j]| > \lambda} (|u[i, j]| - \lambda)$ is a piecewise affine and decreasing function of λ with slope changing at the ordered values $d[t]$. One can then check that the value of λ that meets the constraint $\|S_\lambda(u)\|_1 = \kappa$ is the one given by (16). ■

In words, Proposition 2 tells us that the proximity operator of the ℓ^∞ -norm of a vector field $u \in \mathcal{U}$ is computed by sorting its magnitudes, which of course can be done in $O(N \log N)$ expected time. A similar ingredient appear in the the projection on the ℓ^1 -ball for scalar fields as in [43], [44]. An improved approach replaces sorting with a median-search-like procedure whose expected complexity is linear in N , this has been rediscovered independently in [45] and [46].

Remark 2. Equation (17) can be proved alternatively using Moreau decomposition which allows to conclude that $\operatorname{prox}_{\varphi^*}(u) = u - \operatorname{prox}_\varphi(u)$, see for instance [47] and [16, Lemma 2.10].

³Actually, there is a bijection between λ and κ such that the constrained and Lagrangian problems share the same solution.

B. Forward-backward Projection Algorithm

The projection $f^* = f_0 - \text{div}(u^*)$ is computed by solving the dual unconstrained optimization problem (7). Owing to Lipschitz differentiability of the quadratic term, the dual problem verifies the necessary prerequisite to be solved with the one-step forward-backward splitting recursion.

The forward-backward scheme can be written in a compact form with descent step-size $\mu > 0$ ⁴

$$u^{(k+1)} = \text{prox}_{\mu\tau\|\cdot\|_\infty} \left(u^{(k)} - \mu \nabla \left(f_0 - \text{div}(u^{(k)}) \right) \right), \quad (19)$$

where the proximity operator for $\kappa = \mu\tau$ is defined in (12). The overall algorithm to minimize (6) is summarized in Algorithm 1.

Algorithm 1: Forward-backward total variation projection.

Initialization: choose some $u^{(0)} \in \mathcal{U}$, $\mu \in (0, 1/4)$, set $k = 0$.

Main iteration:

while $\|u^{(k+1)} - u^{(k)}\| > \eta$ **do**

1) *Gradient descent step:* compute

$$\tilde{u}^{(k)} = u^{(k)} - \mu \nabla \left(f_0 - \text{div}(u^{(k)}) \right).$$

2) *TV correction step:* compute d and D as defined in (13) and (14) and set λ as in (16) with $\kappa = \mu\tau$. Define

$$u^{(k+1)} = \tilde{u}^{(k)} - S_\lambda(\tilde{u}^{(k)}),$$

using the soft thresholding operator (15).

3) $k = k + 1$.

Output: $f^* = f_0 - \text{div}(u^{(k+1)})$.

1) *Convergence analysis:* Theorem 1 ensures that the sequence $(f^{(k)})_{k \in \mathbb{N}}$ obtained from Algorithm 1 converges to the solution of (6) with a convergence rate $O(1/k)$.

Theorem 1. Suppose that $\mu \in (0, 1/4)$. Let $u^{(0)} \in \mathbb{R}^{N \times 2}$. The sequence of iterates $f^{(k)} = f_0 - \text{div}(u^{(k)})$, where $u^{(k)}$ is the dual sequence in (19), converges to f^* . More precisely, there exists a $C > 0$ such that

$$\|f^{(k)} - f^*\|^2 \leq C/k.$$

Proof: By coercivity, the set of solutions of (7) is non empty. Moreover, let $\|\text{div}\|$ be the operator norm of the discrete divergence. The term $\|f_0 - \text{div}(u)\|^2/2$ is differentiable whose gradient is Lipschitz continuous with Lipschitz constant $\|\text{div}\|^2 \leq 8$ [6]. Thus, by continuity of the mapping $u \mapsto f = f_0 - \text{div}(u)$, and since the projection f^* is unique, applying [48, Corollary 6.5], we deduce that the sequence of iterates $f^{(k)}$ converges to f^* provided that the step-size satisfies $0 < \underline{\mu} \leq \mu \leq \bar{\mu} < 1/4 \leq 2/\|\text{div}\|^2$.

Let's now turn to the convergence rate. We let $J(f)$ and $\tilde{J}(u)$ be the primal and the dual objectives given in (6) and

⁴The descent step-size can be even a sequence $\mu_k > 0$ that varies from one iteration to another.

(7), namely

$$\begin{aligned} J(f) &= \underbrace{\frac{1}{2}\|f - f_0\|^2}_{H(f)} + \underbrace{\iota_{\{\|\cdot\|_{\text{TV}} \leq \tau\}}(f)}_{G(\nabla f)}, \\ \tilde{J}(u) &= \underbrace{\frac{1}{2}\|f_0 - \text{div}(u)\|^2 - \frac{1}{2}\|f_0\|^2}_{H^*(\text{div}(u))} + \underbrace{\tau\|u\|_\infty}_{G^*(u)}, \end{aligned} \quad (20)$$

where H^* and G^* are the conjugates of H and G respectively. Recall that from Proposition 1, we have

$$\begin{aligned} \min_f J(f) &= -\inf_u \tilde{J}(u) \\ \iff J(f) \geq J(f^*) &= -\tilde{J}(u^*) \geq -\tilde{J}(u), \forall (f, u) \in \mathbb{R}^N \times \mathcal{U} \end{aligned}$$

and from the primal-dual solutions relationship

$$\begin{aligned} H(f^*) + H^*(\text{div}(u^*)) &= \langle f^*, \text{div}(u^*) \rangle \\ \iff f^* &= (D_u H^*)(\text{div}(u^*)), \end{aligned} \quad (21)$$

where $D_u H^*$ is the gradient of H^* with respect to u .

We introduce the following two notions which are essential to prove the convergence rate of our scheme. First, for an optimal solution $u^* \in \mathcal{U}$, we define the Bregman-like distance as the functional $\forall v \in \mathcal{U}$

$$B(v) = G^*(v) - G^*(u^*) + \langle -\nabla(D_u H^*)(\text{div}(u^*)), v - u^* \rangle_{\mathcal{U}}, \quad (22) \quad \forall v \in \mathcal{U}.$$

This is indeed a distance-like function to u^* since $B(v)$ is non-negative and $B(u^*) = 0$. $B(v)$ is non-negative by applying the subgradient inequality (4) to G^* since the minimality condition corresponding to (7) is equivalent to the inclusion $\nabla(D_u H^*)(\text{div}(u^*)) \in \partial G^*(u^*)$. The Bregman distance is widely used to analyze properties of descent algorithms, see [49], [50] and references therein.

Additionally, as H^* is differentiable, we define the Taylor distance as the remainder of the 1st order Taylor expansion of H^* at u^*

$$T(v) = H^*(\text{div}(v)) - H^*(\text{div}(u^*)) - \langle -\nabla(D_u H^*)(\text{div}(u^*)), v - u^* \rangle_{\mathcal{U}}, \quad \forall v \in \mathcal{U}. \quad (23)$$

This is again a non-negative function (H^* is convex), and $T(u^*) = 0$.

It is not difficult to see that these functionals verify the following property

$$B(u^{(k)}) + T(u^{(k)}) = \tilde{J}(u^{(k)}) - \tilde{J}(u^*). \quad (24)$$

Since the function $H^*(f) = \|f - f_0\|^2/2$ is strongly convexity of modulus 1, and using (21), we have

$$\begin{aligned} T(u^{(k)}) &= H^*(\text{div}(u^{(k)})) - H^*(\text{div}(u^*)) - \langle -\nabla(D_u H^*)(\text{div}(u^*)), u^{(k)} - u^* \rangle_{\mathcal{U}} \\ &= H^*(\text{div}(u^{(k)})) - H^*(\text{div}(u^*)) - \langle (D_u H^*)(\text{div}(u^*)), \text{div}(u^{(k)}) - \text{div}(u^*) \rangle \\ &= \frac{1}{2}\|\text{div}(u^{(k)}) - \text{div}(u^*)\|^2 = \frac{1}{2}\|f^{(k)} - f^*\|^2. \end{aligned} \quad (25)$$

Using [17, Theorem 4], the convergence rate over \tilde{J} is such that

$$\tilde{J}(u^{(k)}) - \tilde{J}(u^*) \leq \frac{2\|\operatorname{div}\|^2 R^2}{k+2}, \quad \forall k \geq 0, \quad (26)$$

where R is the radius of the sublevel sets of \tilde{J} , i.e. $R = \max_{u: \tilde{J}(u) \leq \tilde{J}(u^{(0)})} \|u - u^*\| < +\infty$.

Thus, by positivity of $B(u^{(k)})$, and $\|\operatorname{div}\|^2 \leq 8$

$$T(u^{(k)}) \leq \tilde{J}(u^{(k)}) - \tilde{J}(u^*) \leq \frac{16R^2}{k+2}, \quad \forall k \geq 0. \quad (27)$$

Piecing together (25) and (27), we obtain

$$\begin{aligned} \|f^{(k)} - f^*\|^2 &\leq 2T(u^{(k)}) \\ &\leq \frac{32R^2}{k+2}, \quad \forall k \geq 0, \end{aligned} \quad (28)$$

which gives the desired rate. \blacksquare

Remark 3. • *It is important to note that the above proof is general and applies to any problem in the form (20) beyond the TV projection problem considered here, provided that H^* is strongly convex with Lipschitz continuous gradient, or by properties of the conjugate that H is strongly convex and has a Lipschitzian gradient [39, Theorem 4.2.1 and 4.2.2]. Strong convexity is of course important to derive the rate (28) over the primal iterates, but the rate on the dual objective in (26) remains valid anyway. Note that the equality (25) becomes a lower-bound inequality for general strongly convex H^* .*

- *We point out that another proof of the convergence rate of the forward-backward splitting in terms of the objective has recently appeared in [50, Proposition 2]. The rate of [17] is nevertheless sharper.*

This result asserts that Algorithm 1 necessitates as large as $O(1/\eta)$ iterations to reach a η convergence tolerance on the iterates.

C. Multi-step Projection Algorithm

In [17], [51], Nesterov proposes an accelerated multi-step gradient algorithm to solve functionals formed as a sum of two convex terms: a smooth one with Lipschitz-continuous gradient and a non-necessarily smooth term whose structure is simple. By simple, we intend in our context that the corresponding proximity operator is accessible. The dual problem (7) falls within the scope of the Nesterov algorithm. Unlike the forward-backward iteration (19), Nesterov accelerated version uses explicitly all previous iterates $u^{(i)}$, $i < k$ to compute $u^{(k)}$, hence the name multi-step algorithm.

Algorithm 2 details the steps of Nesterov scheme to minimize (7). It is formulated using proximal operators, as described in [7].

1) *Convergence analysis:* The following results gives the convergence rate of the primal sequence $f^{(k)} = f_0 - \operatorname{div}(u^{(k)})$ obtained from Algorithm 2.

Algorithm 2: Nesterov total variation projection.

Initialization: $u^{(0)} \in \mathcal{U}$, $A_0 = 0$, $\xi^{(0)} = 0$, $\mu \in (0, 1/4)$.

Main iteration:

while $\|u^{(k+1)} - u^{(k)}\| > \eta$ **do**

1) *First proximal computation:*

$$v^{(k)} = \operatorname{prox}_{A_k \tau \|\cdot\|_\infty}(u^{(0)} - \xi^{(k)}),$$

where the proximal operator is computed as defined in Proposition 2 with $\kappa = A_k \tau$.

2) Set $a_k = (\mu + \sqrt{\mu^2 + 4\mu A_k})/2$ and

$$\omega^{(k)} = \frac{A_k u^{(k)} + a_k v^{(k)}}{A_k + a_k}.$$

3) *Second proximal computation:*

$$\tilde{\omega}^{(k)} = \omega^{(k)} - \frac{\mu}{2} \nabla (f_0 - \operatorname{div}(\omega^{(k)})),$$

$$u^{(k+1)} = \operatorname{prox}_{\mu\tau/2\|\cdot\|_\infty}(\tilde{\omega}^{(k)}),$$

where the proximal operator is computed as defined in Proposition 2 with $\kappa = \mu\tau/2$.

4) Update $A_{k+1} = A_k + a_k$ and

$$\xi^{(k+1)} = \xi^{(k)} + a_k \nabla (f_0 - \operatorname{div}(u^{(k+1)})).$$

5) $k = k + 1$.

Output: $f^* = f_0 - \operatorname{div}(u^{(k+1)})$.

Theorem 2. *Suppose that $\mu \in (0, 1/4)$. Let $u^{(0)} \in \mathbb{R}^{N \times 2}$. Then, after k iterations, the sequence of iterates $(f^{(k)})_{k \geq 1}$ is such that $\exists C > 0$*

$$\|f^{(k)} - f^*\|^2 \leq C/k^2.$$

Proof: The proof of this result is patterned after that of Theorem 1 starting at (26). Indeed, by virtue of [51, Theorem 3], and $\|\operatorname{div}\|^2 \leq 8$, we arrive at

$$\tilde{J}(u^{(k)}) - \tilde{J}(u^*) \leq \frac{32R^2}{k^2}, \quad k \geq 1, \quad (29)$$

where now $R = \|u^* - u^{(0)}\|^2$. Finally, from (25), we deduce the desired rate

$$\|f^{(k)} - f^*\|^2 \leq \frac{64R^2}{k^2}, \quad k \geq 1. \quad (30)$$

Remark 4. *Again, in the same vein as the proof of Theorem 1, the above proof is general and applies straightforwardly to any problem in the form (20) beyond TV projection, with the proviso that H^* is strongly convex to get the rate on the iterates. For the case where H^* is not necessarily strongly convex, but has a bounded domain, the convergence rate over the primal objective was established in [51, Theorem 3] and [52]. The latter followed the same lines of proof as the former.*

This result asserts that Algorithm 2 necessitates only $O(1/\sqrt{\eta})$ iterations to reach a η -convergence tolerance both on the primal iterates and objective. This is much faster than $O(1/\eta)$ iterations required by the forward-backward scheme.

V. INVERSE PROBLEMS

Image acquisition devices compute $P \leq N$ noisy measurements

$$y = \Phi f_0 + \varepsilon \in \mathbb{R}^P. \quad (31)$$

of a high resolution image $f_0 \in \mathbb{R}^N$, where ε is an additive noise. The linear bounded operator Φ typically accounts for blurring, sub-sampling or missing pixels so that the measured data y only captures a small fraction of the original image f one wishes to recover.

A. Regularization with TV Constraint

A total variation prior allows to regularize the solution by reducing the space of candidate solutions of the inverse problem to those belonging to a total variation ball of finite radius. Thus solving such a constrained inverse problem provides an image that both matches approximately the forward measurements y and that has a low total variation. If the noise is assumed to be zero-mean of finite variance, it is devised to solve the constrained optimization problem

$$\min_{f \in \mathbb{R}^N} \|f\|_{\text{TV}} \quad \text{subject to} \quad \|\Phi f - y\| \leq \gamma \quad (32)$$

where γ is related to the noise standard deviation supposedly known a priori.

On the contrary, if a little is known about the noise ε , but one has some guess τ on the total variation of the image, it is better to consider the following problem

$$f^* = \operatorname{argmin}_{f \in \mathbb{R}^N} \frac{1}{2} \|\Phi f - y\|^2 \quad \text{subject to} \quad \|f\|_{\text{TV}} \leq \tau, \quad (33)$$

where the minimum is not necessarily unique depending on the kernel of the operator Φ .

B. Projected Gradient Descent

The solution to (33) requires the minimization of the gradient Lipschitz functional $\frac{1}{2} \|\Phi f - y\|^2$ under a convex constraint. It is thus possible to use a projected gradient descent

$$f^{(\ell+1)} = \operatorname{Proj}_{\{\|\cdot\|_{\text{TV}} \leq \tau\}} \left(f^{(\ell)} + \nu \Phi^*(y - \Phi f^{(\ell)}) \right), \quad (34)$$

where $\operatorname{Proj}_{\{\|\cdot\|_{\text{TV}} \leq \tau\}}$ is the projector on the TV ball defined in (6).

The following Theorem ensures the convergence of the iteration.

Theorem 3. *If $\nu \in (0, 2/\|\Phi\|^2)$, where $\|\cdot\|$ is the operator spectral norm, then $f^{(\ell)}$ converges to a minimizer f^* of (33) with the rate $O(1/\ell)$ on the objective. If moreover, Φ is injective, then the sequence $f^{(\ell)}$ also converges at the rate $O(1/\ell)$ to the unique minimizer of (33).*

Proof: We use the same arguments as in the first part in the proof of Theorem 1. The convergence rate on the objective is a classical result that can be found in [53]. The convergence rate on the iterates is a consequence of strong convexity when Φ is an injective operator. ■

The projector in (34) cannot be computed exactly but is rather computed via a nested inner iteration using either the forward-backward projection detailed in Algorithm 1 or the multi-step Nesterov projection detailed in Algorithm 2. As these projection algorithms are ran a finite number of inner iterations N_ℓ at each outer iteration ℓ , they yield an estimate

of the TV-ball projector up to an error term a_ℓ . Put formally, we have

$$f^{(\ell+1)} = \tilde{f}^{(\ell)} - \operatorname{div} \left(u^{[\ell+1]} \right) = \operatorname{Proj}_{\{\|\cdot\|_{\text{TV}} \leq \tau\}}(\tilde{f}^{(\ell)}) + a_\ell, \quad (35)$$

where $\tilde{f}^{(\ell)} = f^{(\ell)} + \nu \Phi^*(y - \Phi f^{(\ell)})$, and $u^{[\ell+1]}$ is the dual vector solution provided by applying Algorithm 1 or 2 to $\tilde{f}^{(\ell)}$ with a convergence tolerance $\eta_{\text{proj},\ell}$. At this stage, we advise to use $u^{[\ell]}$ as an initialization in the next call of Algorithm 1 or 2 at the outer iteration $\ell + 1$. This initialization is intended to make the constant smaller in the convergence rate of the projection algorithms; see the constants R in the proofs of Theorem 1 and 2. This leads to the TV projection algorithm to solve inverse problems summarized in Algorithm 3.

Algorithm 3: TV-constrained inverse problem resolution algorithm.

Initialization: set $f^{(0)} = 0$, $u^{[0]} = 0$, $\ell = 0$ and $\nu \in (0, 2/\|\Phi\|^2)$.

Main iteration:

while $\|f^{(\ell+1)} - f^{(\ell)}\| > \eta$ **do**

1) *Gradient descent step:*

$$\tilde{f}^{(\ell)} = f^{(\ell)} + \nu \Phi^*(y - \Phi f^{(\ell)}).$$

2) *Projection step:* use $u^{[\ell]}$ as an initialization in Algorithm 1 or Algorithm 2, and apply them to $\tilde{f}^{(\ell)}$ with convergence tolerance $\eta_{\text{proj},\ell}$. Get the new dual solution $u^{[\ell+1]}$ given by these algorithms.

3) *Update:* Set $f^{(\ell+1)} = \tilde{f}^{(\ell)} - \operatorname{div} \left(u^{[\ell+1]} \right)$.

4) $\ell = \ell + 1$.

The errors a_ℓ are inevitable and may prevent the above algorithm from converging. But fortunately, a distinctive property of the projected gradient (which is actually a special instance of the forward-backward scheme), is its robustness to these errors under appropriate sufficient conditions to be made precise in the next proposition; see also [48], [16]. Denote C as the constant appearing either in the rates of Theorem 1 or 2.

Proposition 3. *Suppose that the sequence $(\eta_{\text{proj},\ell})_{\ell \in \mathbb{N}}$ is summable. If Algorithm 1 or 2 are ran $N_\ell \geq (2C)^{1/\alpha} \eta_{\text{proj},\ell}^{-1/\alpha}$ iterations, where $\alpha = 1/2$ for Algorithm 1 and $\alpha = 1$ for Algorithm 2, then Algorithm 3 converges.*

Proof: In view of Theorem 1 and 2, as well as the triangle inequality, it is sufficient to take $N_\ell \geq (2C)^{1/\alpha} \eta_{\text{proj},\ell}^{-1/\alpha}$ in Algorithm 1 and 2 to reach a $\eta_{\text{proj},\ell}$ -tolerance on the successive primal iterates. From the definition of a_ℓ in (35), this implies that $\|a_\ell\| \leq CN_\ell^{-\alpha} \leq \eta_{\text{proj},\ell}/2$. For the projected gradient outer iteration to converge, the error sequence a_ℓ must obey $\sum_{\ell \in \mathbb{N}} \|a_\ell\| < +\infty$ [48], [16]. It is then sufficient to require that the sequence $(\eta_{\text{proj},\ell})_{\ell \in \mathbb{N}}$ is summable. ■

This result provides a useful guideline on the way the sequence of tolerances $\eta_{\text{proj},\ell}$ should be chosen. But putting exactly this choice into practice is somewhat delicate as we have to circumvent two main difficulties: (i) the estimation of

the constant C , more precisely the abstract constants R in (28)-(30), and (ii) the choice of a summable sequence $\eta_{\text{proj},\ell}$ such that the dual projection algorithms converge in a reasonable number of iterations given their predicted convergence rates. We here propose to take $\eta_{\text{proj},\ell} \propto \rho^\ell$ (or equivalently N_ℓ at least $\propto \rho^{-\ell/\alpha}$ or larger), for $0 < \rho < 1$, a parameter that can be adjusted to the application at hand.

Remark 5. *A this stage, the reader may think of using the multi-step Nesterov algorithm to solve (33) instead of the one-step projected gradient descent. However, one must be aware that while the projected gradient is robust to errors in the computation of the projection operator as we discuss above, we do not have any proof of robustness of the multi-step Nesterov scheme to such errors. This is the main reason we did not use it here to solve TV-constrained inverse problems.*

VI. NUMERICAL EXAMPLES

A. Denoising

We have first tested our TV projection algorithms for denoising, i.e. $\Phi = \text{Id}$ in (31). In our experiment $y = f_0 + \varepsilon$ is an observed image of $N = 512^2$ pixels contaminated by a zero-mean additive white Gaussian noise (AWGN) ε of standard deviation $0.06\|f_0\|_\infty$ (PSNR=12.2 dB).

Figure 1 shows the projections f^* of the noisy image computed with our dual projection algorithms described either in Algorithm 1 or Algorithm 2 for a decreasing value of the constraint τ , so that only the strongest edges are present in the projected images.

Figure 2 compares the convergence speed of our one-step and multi-step projection algorithms summarized in Algorithm 1 and Algorithm 2, with the sub-gradient projection method proposed in [12], [13]. Since an iteration of the multi-step Nesterov projection is approximately twice the computational cost of one iteration of the two other algorithms, we displayed the errors generated by $f^{(k/2)}$ instead of $f^{(k)}$ for the curve of the multi-step algorithm. The one-step algorithm converges slightly faster compared to the sub-gradient projection. Moreover, and as predicted by our convergence analysis, the multi-step projection algorithm clearly outperforms the two other methods. Figure 3 shows the evolution of the total variation of the iterates.

B. Total Variation Texture Synthesis

Statistical approaches to texture synthesis draw an image at random from a set of potential textures defined by constraints that can be learned from an exemplar texture. A standard procedure considers the outputs of a filter-bank, such as wavelets, and constrains their marginal distributions to match those of the exemplar [21], [24]. Instead of considering the uniform distribution inside the set of constraints—assumed to have a non-empty intersection—and sample from this intersection, a simpler approach is to use alternating projections onto the constraints of an initial random white noise image [25]. Although the constrained sets are often non-convex, hopefully, this scheme converges to a point within the constraints that is expected to be visually similar to the exemplar image.



Fig. 1. Examples of total variation projections computed with our algorithm.




Fig. 2. Decay with k of the error $\log_{10}(\|f^{(k)} - f^*\|/\|f^*\|)$ for the one-step forward-backward dual projection in Algorithm 1 (solid line), for the multi-step dual projection in Algorithm 2 (dashed-dotted line), and the sub-gradient projection [12] (dashed line). Here $\tau = \|f_0\|_{\text{TV}}/4$.

We propose here to consider a total variation constraint to better preserve sharp edges during the synthesis

$$\mathcal{C}_{\text{TV}} = \{f \in \mathbb{R}^N \mid \|f\|_{\text{TV}} \leq \tau\},$$

where τ might be computed from an exemplar texture $\tau = \|f_0\|_{\text{TV}}$. This constraint is especially suitable for piecewise-smooth textures representing objects occluding each other.

To enforce other statistical constraints as well, we set up a simple texture model where only the histogram of grayscale intensity is imposed. In a discrete setting, this amounts to selecting a set $V = \{v_i\}_{i=0}^{N-1}$ of N values $v_i \in \mathbb{R}$, and consider the convex compact constraint set

$$\mathcal{C}_V = \{f \in \mathbb{R}^N \mid \{f[i, j]\}_{i, j=0}^{n-1} = V\}.$$

In other words, the constraint $f \in \mathcal{C}_V$ imposes that the contrast of the texture remains the same during the synthesis. We note




Fig. 3. Decay with k of the total variation error $\log_{10}(\|f^{(k)}\|_{\text{TV}}/\tau - 1)$ for the same algorithms as in Figure 2. Here $\tau = \|f_0\|_{\text{TV}}/4$.

that more complicated statistical constraints can be imposed as well, see for instance [25].

The synthesis algorithm, detailed in Algorithm 4, corresponds to the alternating projection onto the two (convex) constraint sets

$$f^{(\ell+1)} = \text{Proj}_{\mathcal{C}_V} \left(\text{Proj}_{\|\cdot\|_{\text{TV}} \leq \tau} (f^{(\ell)}) \right),$$

starting from an initial noise texture $f^{(0)}$.

Algorithm 4: Total variation texture synthesis algorithm.

Initialization: set $f^{(0)}$ = random white noise and set $\ell = 0$.

Main iteration:

while $\|f^{(\ell+1)} - f^{(\ell)}\| > \eta$ **do**

1) *Impose TV constraint:* compute

$$\tilde{f}^{(\ell)} = \text{Proj}_{\|\cdot\|_{\text{TV}} \leq \tau} f^{(\ell)}$$

using our dual projection Algorithm 1 or 2.

2) *Impose histogram constraint:* compute

$$f^{(\ell+1)} = \text{Proj}_{\mathcal{C}_V} \tilde{f}^{(\ell)}$$

using histogram equalization (36).

3) $\ell \leftarrow \ell + 1$.

The orthogonal projection $\text{Proj}_{\mathcal{C}_V}(f)$ of an image onto \mathcal{C}_V corresponds to an histogram equalization. We assume that the values of V are ordered by increasing values $v_i \leq v_{i+1}$, and we denote by $f[r(i)]$ the ordered values of f , where r is a permutation of the pixel indices. The orthogonal projection is then

$$\text{Proj}_{\mathcal{C}_V}(f) = \tilde{f} \quad \text{with} \quad \tilde{f}[r(i)] = v_i. \quad (36)$$

Figure 4 depicts examples of texture synthesis where the pixel values V are equally spaced in $[0, 1]$ so that the textures have a uniform distribution of gray values. When the total variation constraint τ decreases, short edges are removed. This allows to interpolate between a noisy texture and a cartoon image with sharp edges.




Fig. 4. Texture synthesis with TV projection and histogram equalization.

C. inpainting

Inpainting aims at restoring an image f_0 from which a set $\Omega \subset \{0, \dots, n-1\}^2$ of pixels is missing. It corresponds to the inversion of the ill-posed linear problem (31) where Φ is defined as

$$(\Phi f)[i, j] = \begin{cases} 0 & \text{if } (i, j) \in \Omega, \\ f[i, j] & \text{if } (i, j) \notin \Omega. \end{cases} \quad (37)$$

In this case, it is obvious that $\|\Phi\| = 1$, and we use $\nu = 1$ for the projected gradient descent Algorithm 3. The recursion (34) amounts to first imposing the known values outside Ω

$$\tilde{f}^{(\ell)}[i, j] = \begin{cases} f^{(\ell)}[i, j] & \text{if } (i, j) \in \Omega, \\ y[i, j] & \text{otherwise.} \end{cases}$$

and then projecting onto the TV ball

$$f^{(\ell+1)} = \text{Proj}_{\{\|\cdot\|_{\text{TV}} \leq \tau\}} \tilde{f}^{(\ell)}.$$

The top of Figure 5, exemplifies a damaged image y of $N = 512^2$ pixels, with $|\Omega|/N = 0.7$ of randomly removed pixels. The noise is AWGN with standard deviation $0.05\|f_0\|_\infty$ (PSNR=13 dB).

The inpainted image f^* is computed by solving (33) with the projected gradient descent, Algorithm 3, with a total variation constraint size $\tau = 0.6\|f_0\|_{\text{TV}}$. The number of iterations for the projection step 3) is controlled by setting $\eta_{\text{proj}} = 10^{-2}$. Roughly between 10 to 20 iterations of dual projections are required to maintain the total variation constraints during the gradient descent. Figure 6, top, depicts the decay in log scale of the iterates error, that exhibits a roughly linear convergence speed for large ℓ . This rate is likely to be a consequence of the special structure of the masking operator Φ . We think that this can be justified in the light of compressed sensing arguments (the mask is random here), and the convergence analysis in

[50]. We leave these aspects, which are beyond the scope of the paper, to a future work.

D. Deconvolution

An optical system produces a blur that is modeled by convolution with a low pass point spread function (PSF) φ . In such a case, the operator Φ in (31) represents the circular convolution with φ . The convolution by φ removes high frequency details and the total variation constraint (33) helps to recover sharp edges of the original image.

Figure 6, bottom, gives an example of blurred image y of $N = 512^2$ pixels. The PSF is a Gaussian kernel of standard deviation $s = 4$ pixels, normalized to a unit mass. Thus $\|\Phi\| = 1$. The noise is AWGN with standard deviation $0.02\|f_0\|_\infty$ (PSNR=17 dB). We use a gradient descent step-size $\nu = 1.9$ in Algorithm 3.

The deconvolved image f^* is computed by applying Algorithm 3 to y , with a total variation constraint size $\tau = 0.6\|f_0\|_{TV}$. Figure 6, bottom, shows the decay in log scale of the iterates error. This is consistent with the predicted convergence rate of Theorem 3.




Fig. 6. Decay of the error $\log_{10}(\|f^{(\ell)} - f^*\|/\|f_0\|)$ for inpainting (top) and deconvolution (bottom).

CONCLUSION

This paper proposes a new approach to compute the projection of an image on a total variation ball. This approach solves an unconstrained dual formulation of the primal problem, and

boils down to an iterative soft thresholding on the gradient field. Two algorithms for solving the dual minimization problem were proposed. We also studied their convergence behavior and established their convergence rates.

Even though we only focused on the total variation norm in the constraint, our dualization-based projection approach is quite general and extends to any positively 1-homogeneous functional for which the conjugate can be easily computed. This includes constraints involving functionals of the form $\|Gf\|_1$ for any linear operator G (of explicit adjoint) such as the analysis operator of a frame. The scheme also generalizes very easily to arbitrary dimension. In particular, our proof of convergence rate on the forward-backward can be extended to the infinite dimensional case.

We have illustrated the projection algorithm over several applications such as denoising when little is known about the noise statistics, or to enforce total variation constraint in texture synthesis.

A projected gradient descent was also proposed that uses this projection to solve linear inverse problems under a total variation constraint. Our projector can be advantageously used when additional constraints are involved, such as for example contrast bounds $a \leq f \leq b$ or affine constraints $Af = g$. Such a problem with compound constraints can be solved for instance using e.g. Douglas-Rachford splitting and its extensions.

Acknowledgment: The authors would like to thank Pierre Weiss for reading an earlier version of this paper and for fruitful discussions. This work is supported by ANR grant ANR-08-EMER-009.

REFERENCES

- [1] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Phys. D*, vol. 60, no. 1-4, pp. 259–268, 1992.
- [2] D. Goldfarb and W. Yin, "Second-order cone programming methods for total variation-based image restoration," *SIAM J. Scientific Computing*, vol. 27, no. 2, pp. 622–645, 2005.
- [3] T. F. Chan, G. H. Golub, and P. Mulet, "A nonlinear primal-dual method for total variation-based image restoration," *SIAM Journal on Scientific Computing*, vol. 20, no. 6, pp. 1964–1977, Nov. 1999.
- [4] M. Zhu and T. Chan, "An efficient primal-dual hybrid gradient algorithm for total variation image restoration," *UCLA CAM Report 08-34*, 2007.
- [5] C. R. Vogel and M. E. Oman, "Iterative methods for total variation denoising," *SIAM Journal on Scientific Computing*, vol. 17, no. 1, pp. 227–238, Jan. 1996.
- [6] A. Chambolle, "An algorithm for total variation minimization and applications," *Journal of Mathematical Imaging and Vision*, vol. 20, pp. 89–97, 2004.
- [7] P. Weiss, L. Blanc-Féraud, and G. Aubert, "Efficient schemes for total variation minimization under constraints in image processing," *SIAM Journal on Scientific Computing*, 2009, to appear.
- [8] J.-F. Aujol, "Some first-order algorithms for total variation based image restoration," *Journal of Mathematical Imaging and Vision*, 2009, to appear.
- [9] P. Rodríguez and B. Wohlberg, "Efficient minimization method for a generalized total variation functional," *IEEE Trans. Image Proc.*, vol. 18, no. 2, pp. 322–332, 2009.
- [10] Y. Wang, W. Yin, and Y. Zhang, "A fast algorithm for image deblurring with total variation regularization," *Rice University CAAM Technical Report TR07-10*, 2007.
- [11] J. Bect, L. Blanc-Féraud, G. Aubert, and A. Chambolle, "A ℓ_1 -unified variational framework for image restoration," in *Proc. of ECCV04*. Springer-Verlag, 2004, pp. Vol IV: 1–13.
- [12] P.-L. Combettes and J.-C. Pesquet, "Image restoration subject to a total variation constraint," *IEEE Trans. Image Proc.*, vol. 13, no. 9, pp. 1213–1222, 2004.

- [37] C. Chau, P. L. Combettes, J.-C. Pesquet, and V. R. Wajs, "A variational formulation for frame-based inverse problems," *Inv. Prob.*, vol. 23, pp. 1495–1518, 2007.
- [38] M. J. Fadili and J.-L. Starck, "Sparse representation-based image deconvolution by iterative thresholding," in *ADA IV*. France: Elsevier, 2006.
- [39] C. Lemaréchal and J.-B. Hiriart-Urruty, *Convex Analysis and Minimization Algorithms I and II*, 2nd ed. Springer, 1996.
- [40] J.-J. Moreau, "Fonctions convexes duales et points proximaux dans un espace hilbertien," *CRAS Sér. A Math.*, vol. 255, pp. 2897–2899, 1962.
- [41] I. Ekeland and R. Temam, *Analyse convexe et problèmes variationnels*. Dunod, 1974.
- [42] A. Chambolle, R. D. Vore, N. Lee, and B. Lucier, "Nonlinear wavelet image processing: Variational problems, compression, and noise removal through wavelet shrinkage," *IEEE Trans. Image Proc.*, vol. 7, no. 3, pp. 319–335, March 1998.
- [43] E. J. Candes and J. Romberg, "Practical signal recovery from random projections," in *Wavelet Applications in Signal and Image Processing XI*, vol. 5914. SPIE, 2004.
- [44] I. Daubechies, M. Fornasier, and I. Loris, "Accelerated projected gradient method for linear inverse problems with sparsity constraints," *J. Fourier Anal. Appl.*, vol. 14, no. 5-6, pp. 764–792, 2008.
- [45] J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra, "Efficient projection onto L1 ball for learning in high dimension," in *ICML '08: Proceedings of the 25th international conference on Machine learning*, July 2008, pp. 272–279.
- [46] E. VandenBerg, M. Schmidt, M. P. Friedlander, and K. Murphy, "Group sparsity via linear time projection," Department of Computer Science, University of British Columbia, Tech. Rep. TR-2008-09, June 2008.
- [47] J.-J. Moreau, "Proximité et dualité dans un espace hilbertien," *Bulletin de la S.M.F.*, vol. 93, pp. 273–299, 1965.
- [48] P. L. Combettes, "Solving monotone inclusions via compositions of nonexpansive averaged operators," *Optimization*, vol. 53, no. 5-6, pp. 475–504, 2004.
- [49] H. H. Bauschke, J. M. Borwein, and P. L. Combettes, "Bregman monotone optimization algorithms," *SIAM J. on Control and Optimization*, vol. 42, no. 2, pp. 596–636, 2003.
- [50] K. Bredies and D. Lorenz, "Linear convergence of iterative soft-thresholding," *J. Fourier Anal. Appl.*, vol. 14, pp. 813–837, Dec 2008.
- [51] Y. Nesterov, "Smooth minimization of non-smooth functions," *Math. Program.*, vol. 103, no. 1, Ser. A, pp. 127–152, 2005.
- [52] P. Weiss, "Algorithmes rapides d'optimisation convexe: application à la restauration d'images et à la détection de changements," Ph.D. dissertation, Université Sophia-Antipolis, Nice, France, 2008.
- [53] B. Polyak, *Introduction to Optimization*. New York: Optimization Software, Inc., Publications Division, 1987.