

MOCCA: Mirrored Convex/Concave Optimization for Nonconvex Composite Functions

Rina Foygel Barber

*Department of Statistics
University of Chicago
5747 South Ellis Avenue
Chicago, IL 60637, USA*

RINA@UCHICAGO.EDU

Emil Y. Sidky

*Department of Radiology
University of Chicago
5841 South Maryland Avenue
Chicago, IL 60637, USA*

SIDKY@UCHICAGO.EDU

Editor: Tong Zhang

Abstract

Many optimization problems arising in high-dimensional statistics decompose naturally into a sum of several terms, where the individual terms are relatively simple but the composite objective function can only be optimized with iterative algorithms. In this paper, we are interested in optimization problems of the form $F(Kx) + G(x)$, where K is a fixed linear transformation, while F and G are functions that may be nonconvex and/or nondifferentiable. In particular, if either of the terms are nonconvex, existing alternating minimization techniques may fail to converge; other types of existing approaches may instead be unable to handle nondifferentiability. We propose the MOCCA (mirrored convex/concave) algorithm, a primal/dual optimization approach that takes a local convex approximation to each term at every iteration. Inspired by optimization problems arising in computed tomography (CT) imaging, this algorithm can handle a range of nonconvex composite optimization problems, and offers theoretical guarantees for convergence when the overall problem is approximately convex (that is, any concavity in one term is balanced out by convexity in the other term). Empirical results show fast convergence for several structured signal recovery problems.

Keywords: MOCCA, ADMM, nonconvex, penalized likelihood, total variation, computed tomography

1. Introduction

We consider the problem of minimizing a composite objective function of the form

$$F(Kx) + G(x) \tag{1}$$

over $x \in \mathbb{R}^d$, where $K \in \mathbb{R}^{m \times d}$ is a fixed linear operator, and F and G are functions which are potentially nonconvex and/or nondifferentiable. Optimization problems of this form arise in many applications, and in particular, the algorithm developed here was motivated by an

image reconstruction problem for computed tomography (CT), an imaging technology used often in medicine and in other domains.

When F and G are both convex, many existing methods are well-equipped to handle this optimization problem, even in high dimensions. For example, the Alternating Direction Method of Multipliers (ADMM) (Boyd et al., 2011) and related primal/dual methods yield effective algorithms when the functions F and G both have inexpensive proximal maps, defined as

$$\text{Prox}_F(u) = \arg \min_w \left\{ \frac{1}{2} \|w - u\|^2 + F(w) \right\}$$

for any $u \in \mathbb{R}^m$, and same for Prox_G defined on \mathbb{R}^d . Methods such as ADMM are especially effective if F has an inexpensive proximal map but the linear transformation of the same function, i.e. the function $x \mapsto F(Kx)$, does not. If F does not offer an inexpensive proximal map but is instead smoothly differentiable (even if nonconvex), while G does have a fast proximal map, then methods such as proximal gradient descent (see e.g. Nesterov, 2013; Beck and Teboulle, 2009) can be applied instead of a primal/dual method.

In this paper, we consider a more challenging setting where F and G may both be nonconvex and nondifferentiable. For instance, we can consider working with functions that can be decomposed as $F = F_{\text{cvx}} + F_{\text{diff}}$ and $G = G_{\text{cvx}} + G_{\text{diff}}$, where we assume that $F_{\text{cvx}}, G_{\text{cvx}}$ are convex but do not need to be differentiable, while $F_{\text{diff}}, G_{\text{diff}}$ are potentially nonconvex but are differentiable. (If F_{diff} and G_{diff} are concave, then this type of optimization problem is often referred to as “convex/concave”.) As we will see, this formulation arises naturally in a range of applications, but in general, cannot be handled by existing methods that are designed with convexity in mind. In this work, we generalize to a more flexible framework where F and G can each be locally approximated at any point with a convex function.

A special case is the setting where G is convex while F is nonconvex, but $x \mapsto F(Kx)$ is convex (i.e. if F is twice differentiable, then this is equivalent to assuming that $K^\top(\nabla^2 F)K \succeq 0$ but $\nabla^2 F \not\preceq 0$). In this case, the overall optimization problem, i.e. minimizing $F(Kx) + G(x)$, is a convex problem, that is, any local minimum is guaranteed to be a global minimum, and thus we might expect that this problem would be simple to optimize. Surprisingly, this may not be the case—if F is nondifferentiable, then we cannot use gradient descent on F , while the nonconvexity of F means that existing primal/dual optimization techniques might not be applicable.

In the nonconvex, or more specifically, convex/concave setting, one of the most common techniques used in place of convex methods is the majorization/minimization approach (Ortega and Rheinboldt, 1970; Hunter and Lange, 2000), where at each iteration, we work with a convex upper bound on the nonconvex objective function. Specifically, for the setting considered here, at iteration t we would choose some convex functions $F^{(t)}, G^{(t)}$ satisfying $F^{(t)} \geq F$ and $G^{(t)} \geq G$, then solve the modified optimization problem $\min_x \{F^{(t)}(Kx) + G^{(t)}(x)\}$. However, the modified optimization problem may itself be very challenging to solve in this setting, so we often cannot apply the majorization/minimization technique in a straightforward way. Our work combines the ideas of majorization/minimization with primal/dual techniques to handle this composite optimization problem.

1.1 The MOCCA Algorithm

In this work, we propose the mirrored convex/concave (MOCCA) algorithm, which offers an approach that combines some of the techniques described above. We work with a primal/dual formulation of the problem, and incorporate a majorization/minimization type step at each iteration. To motivate our method, we first present an existing method for the case where F and G are both convex, the Chambolle-Pock (CP) algorithm (Chambolle and Pock, 2011) (this method is closely related to other existing algorithms, which we discuss later on).

The CP algorithm is derived by considering the problem in a different form. From this point on, for clarity, we will use variables x, y, z to denote points in \mathbb{R}^d and u, v, w to denote points in \mathbb{R}^m . Since we are considering the setting where F is convex, by duality we can write

$$\min_x \{F(Kx) + G(x)\} = \min_x \max_w \{\langle Kx, w \rangle - F^*(w) + G(x)\} .$$

where F^* is the conjugate to F (Rockafellar, 1997), also known as the Legendre-Fenchel transform of F , and is defined as

$$F^*(w) = \max_v \{\langle w, v \rangle - F(v)\} . \quad (2)$$

The primal variable x and dual variable w define a saddle-point problem. Given step sizes Σ, T which are positive diagonal matrices, the (preconditioned) CP algorithm (Chambolle and Pock, 2011; Pock and Chambolle, 2011) iterates the following steps:

$$\begin{cases} x_{t+1} = \arg \min_x \left\{ \langle Kx, w_t \rangle + G(x) + \frac{1}{2} \|x - x_t\|_{T^{-1}}^2 \right\} , \\ w_{t+1} = \arg \min_w \left\{ -\langle K\bar{x}_{t+1}, w \rangle + F^*(w) + \frac{1}{2} \|w - w_t\|_{\Sigma^{-1}}^2 \right\} , \end{cases} \quad (3)$$

where $\bar{x}_{t+1} = x_{t+1} + \theta(x_{t+1} - x_t)$ is an extrapolation term for some parameter $\theta \in [0, 1]$ (generally $\theta = 1$). Here the two norms are calculated via the definition $\|x\|_A := \sqrt{x^\top A x}$ (for any positive semidefinite matrix $A \succeq 0$). When $\|\Sigma^{1/2} K T^{1/2}\| < 1$, convergence properties for this algorithm have been proved, e.g. Chambolle and Pock (2011); Pock and Chambolle (2011); He and Yuan (2012).¹ Setting $\theta = 0$ reduces to an earlier approach, the Primal-Dual Hybrid Gradient algorithm (Esser et al., 2009); with $\theta = 1$, the CP algorithm is equivalent to a modification of ADMM (discussed later in Section 3.1).

We now ask how we could modify the discussed methods to handle nonconvexity of F and/or G . One approach would be to approximate the problem with a convex optimization problem, that is, to take some approximation to F and G that are chosen to be a convex function. Of course, in general, there may not be a convex function that will provide a globally accurate approximation to a nonconvex function, but local convex approximations may be possible.

Consider a family of approximations to the functions F and G , indexed by $(z, v) \in \mathbb{R}^d \times \mathbb{R}^m$, the (primal) points at which the local approximations are taken. We write

1. The original form of the Chambolle-Pock algorithm (Chambolle and Pock, 2011), without preconditioning, can be obtained from (3) by replacing Σ, T with $\sigma \mathbf{I}, \tau \mathbf{I}$ for scalar step size parameters $\sigma, \tau > 0$, with convergence results proved if $\sigma \tau \|K\|^2 < 1$; however, in general the preconditioned form gives better performance and we only consider the preconditioned version here.

$F_v : \mathbb{R}^m \rightarrow \mathbb{R}$ and $G_z : \mathbb{R}^d \rightarrow \mathbb{R}$ for these approximations, and from this point on, we implicitly assume that the approximations satisfy, for any points $(z, v) \in \text{dom}(G) \times \text{dom}(F)$,

$$\left\{ \begin{array}{l} \text{dom}(F_v) = \text{dom}(F) \text{ and } \text{dom}(G_z) = \text{dom}(G); \\ F_v \text{ and } G_z \text{ are convex and continuous functions;} \\ F_v \text{ and } G_z \text{ are accurate up to first order: } F - F_v \text{ and } G - G_z \text{ are differentiable,} \\ \text{with } (F - F_v)(v) = 0, \nabla(F - F_v)(v) = 0, (G - G_z)(z) = 0, \nabla(G - G_z)(z) = 0. \end{array} \right. \quad (4)$$

(Here $\text{dom}(\cdot)$ denotes the domain of a function, i.e. all points at which the function takes a finite value.) In particular this assumption implicitly requires that F and G both have convex domains.

As a special case, in some settings we can consider decomposing each function into the sum of a convex and a differentiable term, $F = F_{\text{cvx}} + F_{\text{diff}}$ and $G = G_{\text{cvx}} + G_{\text{diff}}$, as mentioned before; we can then take linear approximations to F_{diff} and to G_{diff} at the points v and z , respectively, to obtain

$$\left\{ \begin{array}{l} F_v(w) := F_{\text{cvx}}(w) + [F_{\text{diff}}(v) + \langle w - v, \nabla F_{\text{diff}}(v) \rangle], \\ G_z(x) := G_{\text{cvx}}(x) + [G_{\text{diff}}(z) + \langle x - z, \nabla G_{\text{diff}}(z) \rangle]. \end{array} \right. \quad (5)$$

In particular, if F_{diff} and G_{diff} are both concave (and thus F and G are each convex/concave), these two approximations are standard convex majorizations to F and G taken at points v and z (as might be used in a majorization/minimization algorithm in some settings).

Since F_v, G_z are convex, we can substitute them in place of F, G in the iterations of the CP algorithm (3):

$$\left\{ \begin{array}{l} x_{t+1} = \arg \min_x \left\{ \langle Kx, w_t \rangle + G_z(x) + \frac{1}{2} \|x - x_t\|_{\Sigma^{-1}}^2 \right\}, \\ w_{t+1} = \arg \min_w \left\{ -\langle K\bar{x}_{t+1}, w \rangle + F_v^*(w) + \frac{1}{2} \|w - w_t\|_{\Sigma^{-1}}^2 \right\}. \end{array} \right. \quad (6)$$

We will see later on (in Section 3.3) that this formulation is closely related to the ADMM algorithm, and in fact in the special case given in (5), if F_{diff} and G_{diff} are concave, then MOCCA can be viewed as a special case of Bregman ADMM (Wang and Banerjee, 2014; Wang et al., 2014a).

Of course, a key question remains: which primal points (z, v) should we use for constructing the convex approximations to F and to G , at iteration t of the algorithm? We find that, before solving for (x_{t+1}, w_{t+1}) , we should use expansion points

$$(z, v) = (x_t, \Sigma^{-1}(w_{t-1} - w_t) + K\bar{x}_t) .$$

We will return shortly to the question of how these values were chosen; with this choice in place, the MOCCA algorithm is defined in Algorithm 1.

For reasons of stability, it may sometimes be desirable to update the expansion points (z, v) only periodically, and we will incorporate this option into a more general version of MOCCA, given in Algorithm 2. Specifically, at the t th stage, we repeat the (x, w) updates L_t many times; we refer to these repeated updates as the “inner loop”. In fact, the t th “inner

Algorithm 1 MOCCA algorithm

Input: Functions F, G with local convex approximations F_v, G_z , linear operator K , positive diagonal step size matrices Σ, T , extrapolation parameter $\theta \in [0, 1]$.

Initialize: Primal point $x_0 \in \mathbb{R}^d$, dual point $w_0 \in \mathbb{R}^m$, expansion points $(z_0, v_0) \in \mathbb{R}^d \times \mathbb{R}^m$.

for $t = 0, 1, 2, \dots$ **do**

Update x and w variables: writing $\bar{x}_{t+1} = x_{t+1} + \theta(x_{t+1} - x_t)$, define

$$\begin{cases} x_{t+1} = \arg \min_x \left\{ \langle Kx, w_t \rangle + G_{z_t}(x) + \frac{1}{2} \|x - x_t\|_{T^{-1}}^2 \right\}, \\ w_{t+1} = \arg \min_w \left\{ -\langle K\bar{x}_{t+1}, w \rangle + F_{v_t}^*(w) + \frac{1}{2} \|w - w_t\|_{\Sigma^{-1}}^2 \right\}. \end{cases}$$

Update expansion points: define

$$\begin{cases} z_{t+1} = x_{t+1}, \\ v_{t+1} = \Sigma^{-1}(w_t - w_{t+1}) + K\bar{x}_{t+1} \in \partial F_{v_t}^*(w_{t+1}). \end{cases}$$

until some convergence criterion is reached.

Algorithm 2 MOCCA algorithm (stable version with “inner loop”)

Input / Initialize: Same as for Algorithm 1, along with inner loop lengths L_1, L_2, \dots

for $t = 0, 1, 2, \dots$ **do**

Define $(x_{t+1;0}, w_{t+1;0}) = (x_t, w_t)$.

for $\ell = 1, 2, \dots, L_{t+1}$ **do**

Update x and w variables: writing $\bar{x}_{t+1;\ell} = x_{t+1;\ell} + \theta(x_{t+1;\ell} - x_{t+1;\ell-1})$, define

$$\begin{cases} x_{t+1;\ell} = \arg \min_x \left\{ \langle Kx, w_{t+1;\ell-1} \rangle + G_{z_t}(x) + \frac{1}{2} \|x - x_{t+1;\ell-1}\|_{T^{-1}}^2 \right\}, \\ w_{t+1;\ell} = \arg \min_w \left\{ -\langle K\bar{x}_{t+1;\ell}, w \rangle + F_{v_t}^*(w) + \frac{1}{2} \|w - w_{t+1;\ell-1}\|_{\Sigma^{-1}}^2 \right\}. \end{cases}$$

end for

Define $(x_{t+1}, w_{t+1}) = \frac{1}{L_{t+1}} \sum_{\ell=1}^{L_{t+1}} (x_{t+1;\ell}, w_{t+1;\ell})$.

Update expansion points: define

$$\begin{cases} z_{t+1} = \frac{1}{L_{t+1}} \sum_{\ell=1}^{L_{t+1}} x_{t+1;\ell}, \\ v_{t+1} = \frac{1}{L_{t+1}} \sum_{\ell=1}^{L_{t+1}} (\Sigma^{-1}(w_{t+1;\ell-1} - w_{t+1;\ell}) + K\bar{x}_{t+1;\ell}). \end{cases}$$

until some convergence criterion is reached.

loop” is simply running the CP algorithm for the convex problem $\min_x \{F_{v_t}(Kx) + G_{z_t}(x)\}$. Then, we average over the inner loop and calculate a single update of the expansion points (z, v) . Observe that, if we set $L_t = 1$ for all t , we do only a single (x, w) update in each “inner loop” and thus have reduced to the basic form of MOCCA. In practice, the basic form (Algorithm 1) performs well, and the more stable version (Algorithm 2) is primarily

proposed here for theoretical purposes, as some of our convergence guarantees do not hold for the basic version. However, in some settings the added stability does help empirically.

We remark that F, G (and their approximations F_v, G_z) are allowed to take the value $+\infty$, for instance, to reflect a constraint. For example we might have $G(x) = \delta\{\|x\|_2 \leq 1\}$, the convex indicator function taking the value $+\infty$ if the constraint $\|x\|_2 \leq 1$ is violated and zero otherwise; this has the effect of imposing a constraint on the x update step of our algorithm. Furthermore, in settings where F_v may not be strongly convex, its conjugate F_v^* may not be finitely valued; we would have $F_v^*(w) = +\infty$ (for some, but not all, w). For instance if $F_v(w) = \|w\|_1$ then $F_v^* = \delta\{\|w\|_\infty \leq 1\}$, which has the effect of imposing a constraint on w in the w update step of our algorithm. Our theoretical results in this paper hold across all these settings, i.e. we do not assume that any of the functions F, G, F_v, G_z, F_v^* are everywhere finite, but instead work in the domains of these functions.

1.1.1 SIMPLE SPECIAL CASES

Before discussing the implementation and behavior of MOCCA for general nonconvex and/or nonsmooth problems, we pause to illustrate that MOCCA can be viewed as a generalization of many existing techniques.

- If F, G are both convex with easy proximal maps, then we can of course choose the trivial convex families $F_v = F$ and $G_z = G$; the MOCCA algorithm then reduces to the Chambolle-Pock (Chambolle and Pock, 2011) or Primal-Dual Hybrid Gradient (Esser et al., 2009) algorithm (depending on our choice of the extrapolation parameter θ). These methods can handle composite objective functions with convex terms; MOCCA extends these methods to a setting with nonconvex terms.
- In the setting where we want to minimize a function $G(x)$ which is a sum of a convex term and a differentiable term, $G(x) = g(x) + h(x)$, we can show that MOCCA reduces to proximal gradient descent (Nesterov, 2013; Beck and Teboulle, 2009) as a special case. Specifically, we define the approximations $G_z(x) = g(x) + h(z) + \langle \nabla h(z), x - z \rangle$. In this setting there is no F function, and hence no w variable; taking $T = \tau \cdot \mathbf{I}_d$, the steps of Algorithm 1 become

$$\begin{cases} x_{t+1} = \arg \min_x \left\{ G_{z_t}(x) + \frac{1}{2\tau} \|x - x_t\|_2^2 \right\} \\ \quad = \arg \min_x \left\{ g(x) + \langle \nabla h(z_t), x \rangle + \frac{1}{2\tau} \|x - x_t\|_2^2 \right\}, \\ z_{t+1} = x_{t+1}, \end{cases}$$

which simplifies to the update scheme

$$x_{t+1} = \text{Prox}_{\tau \cdot g}(x_t - \tau \cdot \nabla h(x_t)). \tag{7}$$

This is exactly the proximal gradient descent algorithm with step size τ .

Proximal gradient descent can handle a function which combines a differentiable term with a convex term as long as the convex term has an easy proximal map; MOCCA extends this method to a setting where the convex terms lack an easy proximal map due to linear transformations, leading to composite optimization problems.

We will discuss other existing methods in more detail later on in Section 3.

1.1.2 STEP SIZE PARAMETERS Σ AND T

We now turn to the question of choosing the diagonal step size matrices Σ and T . As we will see later on, good convergence properties are attained when Σ, T are chosen sufficiently small, to satisfy $\|\Sigma^{1/2}KT^{1/2}\| < 1$ —this condition on Σ and T is derived by Pock and Chambolle (2011) for the preconditioned CP algorithm, and appears in our theory as well. Here $\|\cdot\|$ is the matrix operator norm (i.e. the largest singular value). To choose matrices that satisfy this requirement, Pock and Chambolle (2011) propose a parametrized family of choices: after fixing some parameter $\lambda > 0$, define²

$$\Sigma_{ii} = \frac{\lambda}{\sum_j |K_{ij}|} \quad \text{and} \quad T_{jj} = \frac{\lambda^{-1}}{\sum_i |K_{ij}|}. \quad (8)$$

Empirically, we find that higher values of λ are more stable but lead to slower convergence; it seems that the best choice is the smallest possible λ such that the algorithm does not diverge. It may also be interesting to consider varying λ adaptively over the iterations of the algorithm, but we do not study this extension here.

1.1.3 UNDERSTANDING THE CHOICE OF EXPANSION POINTS

We now return to our choice of the expansion points $(z, v) \in \mathbb{R}^d \times \mathbb{R}^m$. We will give an intuition for the choices of these points in the MOCCA algorithm. To examine this question, first consider the goal for optimization: we would like to find

$$x^* \in \arg \min_x \{F(Kx) + G(x)\},$$

or if this problem is nonconvex then we may be satisfied to let x^* be a local minimizer or critical point of this objective function. We then need to find primal points $z \in \mathbb{R}^d$ and $v \in \mathbb{R}^m$, such that replacing F with F_v and G with G_z still yields the same solution, i.e. so that

$$x^* \in \arg \min_x \{F_v(Kx) + G_z(x)\}. \quad (9)$$

Examining the first-order optimality conditions for each of these problems, it follows that we should set $(z, v) = (x^*, Kx^*)$ to ensure that (9) holds.

Of course, x^* is unknown and so we cannot set $(z, v) = (x^*, Kx^*)$. Instead, a logical approach would be to set $(z_t, v_t) = (x_t, Kx_t)$, before solving for (x_{t+1}, w_{t+1}) . Then, hopefully, as x_t converges to x^* we will also have (z_t, v_t) converging to (x^*, Kx^*) . However, in practice, we find that this approach does not always perform as well as expected. Specifically, the problem lies with the choice $v_t = Kx_t$, relative to the primal/dual structure of the algorithm.

To understand why, imagine that F and G are actually convex, but we nonetheless are taking local approximations F_v and G_z (which are also convex), perhaps for computational reasons. Then we would like our x and w update steps (6) to coincide with the updates (3) of the original CP algorithm. Examining the optimality conditions, this will occur when

$$\partial G(x_{t+1}) = \partial G_{z_t}(x_{t+1}) \quad \text{and} \quad \partial F^*(w_{t+1}) = \partial F_{v_t}^*(w_{t+1}). \quad (10)$$

2. In fact these choices for Σ, T satisfy the matrix norm constraint more weakly, with $\|\Sigma^{1/2}KT^{1/2}\| \leq 1$ rather than a strict inequality, but this is sufficient in practice.

(Here we are ignoring issues of multivalued subdifferentials since our aim is only to give intuition.) Using the definitions of \mathbf{G} and \mathbf{G}_{z_t} , for the x step our requirement in (10) is equivalent to

$$\nabla(\mathbf{G} - \mathbf{G}_{z_t})(x_{t+1}) = 0,$$

which will certainly hold if $z_t = x_{t+1}$ by our assumption (4) on the function \mathbf{G}_{z_t} . Since we have not yet solved for x_{t+1} , we instead choose the expansion point $z_t = x_t$ for the function \mathbf{G} , as previously proposed.

For the w step, our outcome will be different. Subgradients satisfy a duality property, namely, $w \in \partial\mathbf{F}(u)$ if and only if $u \in \partial\mathbf{F}^*(w)$ for any convex function \mathbf{F} and its conjugate \mathbf{F}^* . The requirement $\partial\mathbf{F}^*(w_{t+1}) = \partial\mathbf{F}_{v_t}^*(w_{t+1})$ in (10) therefore yields $w_{t+1} \in \partial\mathbf{F}(\partial\mathbf{F}_{v_t}^*(w_{t+1}))$ by this duality property, and so we have

$$\begin{aligned} w_{t+1} &\in \partial\mathbf{F}(\partial\mathbf{F}_{v_t}^*(w_{t+1})) \\ &= \partial\mathbf{F}_{v_t}(\partial\mathbf{F}_{v_t}^*(w_{t+1})) + \nabla(\mathbf{F} - \mathbf{F}_{v_t})(\partial\mathbf{F}_{v_t}^*(w_{t+1})) \\ &= w_{t+1} + \nabla(\mathbf{F} - \mathbf{F}_{v_t})(\partial\mathbf{F}_{v_t}^*(w_{t+1})) \end{aligned}$$

where the last step again holds from the duality property of subgradients. So, we see that we would like

$$\nabla(\mathbf{F} - \mathbf{F}_{v_t})(\partial\mathbf{F}_{v_t}^*(w_{t+1})) = 0$$

which, according to our assumption (4) on the expansions \mathbf{F}_{v_t} , will hold if

$$v_t \in \partial\mathbf{F}_{v_t}^*(w_{t+1}).$$

In other words, we would like v_t to be the primal point that corresponds to the dual point w_{t+1} —that is, the primal point that *mirrors* the dual point w_{t+1} . Of course, this is not possible since we have not yet computed w_{t+1} , and furthermore v_t appears on both sides of this equation. Instead, we take $v_t \in \partial\mathbf{F}_{v_{t-1}}^*(w_t)$. Looking at the first-order optimality conditions for the update step for w_t , we see that we can satisfy this expression by choosing

$$v_t = \Sigma^{-1}(w_{t-1} - w_t) + K\bar{x}_t \in \partial\mathbf{F}_{v_{t-1}}^*(w_t).$$

In fact, we will see in Section 3.3 that this choice for v_t is very logical in light of the connection between the CP algorithm and the Alternating Direction Method of Multipliers (ADMM) (Boyd et al., 2011).

For the stable form of MOCCA, Algorithm 2, our choice for expansion points (z, v) takes an average over each inner loop, which we will see gives sufficient stability for our convergence results to hold.

1.1.4 CHECKING CONVERGENCE

Here we give a simple way to check whether the basic MOCCA algorithm, Algorithm 1, is near a critical point (e.g. a local minimum). (We treat this question more formally, for the more general Algorithm 2, in our theoretical results later on.) Due to the first-order accuracy of the convex approximations to \mathbf{F} and \mathbf{G} as specified in (4), a critical point $x \in \mathbb{R}^d$ for the objective function $\mathbf{F}(Kx) + \mathbf{G}(x)$ is characterized by the first-order condition

$$0 \in K^\top \partial\mathbf{F}_{Kx}(Kx) + \partial\mathbf{G}_x(x). \quad (11)$$

Equivalently, we can search for a dual variable $w \in \mathbb{R}^m$ such that

$$-K^\top w \in \partial \mathbf{G}_x(x) \quad \text{and} \quad w \in \partial \mathbf{F}_{Kx}(Kx).$$

We can expand this condition to include additional variables $(z, v) \in \mathbb{R}^d \times \mathbb{R}^m$:

$$\begin{cases} -K^\top w \in \partial \mathbf{G}_z(x), \\ w \in \partial \mathbf{F}_v(Kx) \Leftrightarrow Kx \in \partial \mathbf{F}_v^*(w), \\ z = x, \\ v = Kx. \end{cases}$$

To check whether these conditions hold approximately, we can take the following ‘‘optimality gap’’:

$$\text{OptimalityGap}(x, w, z, v) = \left\| -K^\top w - \partial \mathbf{G}_z(x) \right\|_2^2 + \|Kx - \partial \mathbf{F}_v^*(w)\|_2^2 + \|z - x\|_2^2 + \|v - Kx\|_2^2.$$

Here, if any of the subdifferentials are multivalued, we can interpret these norms as choosing some element of the corresponding subdifferentials. Now we consider the value of this gap at an iteration of the MOCCA algorithm (in its original form, Algorithm 1). By the definitions of $x_{t+1}, w_{t+1}, z_t, v_t$, we can show that

$$\begin{cases} 0 \in K^\top w_t + \partial \mathbf{G}_{z_t}(x_{t+1}) + \mathbf{T}^{-1}(x_{t+1} - x_t), \\ 0 \in -K\bar{x}_{t+1} + \partial \mathbf{F}_{v_t}^*(w_{t+1}) + \Sigma^{-1}(w_{t+1} - w_t), \\ z_t = x_t, \\ v_t = \Sigma^{-1}(w_t - w_{t-1}) + K\bar{x}_t. \end{cases}$$

Therefore, plugging these calculations in to the definition of the optimality gap, we see that

$$\begin{aligned} & \text{OptimalityGap}(x_{t+1}, w_{t+1}, z_t, v_t) \\ &= \left\| -K(w_t - w_{t+1}) + \mathbf{T}^{-1}(x_t - x_{t+1}) \right\|_2^2 + \left\| K(x_t - x_{t+1}) + \Sigma^{-1}(w_t - w_{t+1}) \right\|_2^2 \\ & \quad + \|x_t - x_{t+1}\|_2^2 + \left\| K(x_{t-1} - 2x_t + x_{t+1}) + \Sigma^{-1}(w_{t-1} - w_t) \right\|_2^2 \\ &= \mathcal{O} \left(\left\| \begin{pmatrix} x_{t-1} - x_t \\ w_{t-1} - w_t \end{pmatrix} \right\|_2^2 + \left\| \begin{pmatrix} x_t - x_{t+1} \\ w_t - w_{t+1} \end{pmatrix} \right\|_2^2 \right). \end{aligned}$$

In other words, if the change in the variables (x_t, w_t) converges to zero as $t \rightarrow \infty$, then the optimality gap is also converging to zero.

1.1.5 PREVIEW OF THEORETICAL RESULTS

We present two theoretical results in this work. The first is fairly standard in the related literature: in Theorem 1 we show that if the algorithm does converge to a point, then we have reached a critical point of the original optimization problem. (Since the simple form, Algorithm 1, is a special case of the stable form, Algorithm 2, we prove this result for the stable algorithm only.)

The novelty of our theory lies in our convergence guarantee, given in Theorem 2, where we prove that the stable form of MOCCA, given in Algorithm 2, is guaranteed to converge to a nearly-globally-optimal solution, under some assumptions on convexity and curvature. Specifically, we consider a scenario where convexity and nonconvexity in F and G counter-balance each other, so that the overall function

$$x \mapsto F(Kx) + G(x) \tag{12}$$

is itself either strongly convex or satisfies restricted strong convexity assumptions (which we will discuss in detail in Section 4.2.1). It is important to note that even the globally convex setting is by no means trivial—even if (12) is strongly convex, if F itself is nonconvex it may be the case that ADMM and other primal/dual or alternating minimization algorithms diverge or converge to the wrong solution, as we discuss later in Section 3.2. Crucially, our results allow F and G to be nondifferentiable as well as nonconvex, a setting that is necessary in practice but is not covered by existing theory.

1.1.6 OUTLINE OF PAPER

The remainder of the paper is organized as follows. In Section 2, we present several important applications where the minimization problem considered here, with a nonconvex composite objective function as in (1), arises naturally: regression problems with errors in covariates, isotropic total variation penalties, nonconvex total variation penalties, and image reconstruction problems in computed tomography (CT) imaging. In Section 3 we give background on several types of existing algorithms for convex and nonconvex composite objective functions, and compare a range of existing results to the work presented here. Theoretical results on the convergence properties of our algorithm are given in Section 4. We study the empirical performance of MOCCA in Section 5. Proofs are given in Section 6, with technical details deferred to the Appendix. In Section 7 we discuss our findings and outline directions for future research.

2. Applications

We now highlight several applications of the MOCCA algorithm, in high-dimensional statistics and in imaging.

2.1 Regression with Errors in Variables

Recent work by Loh and Wainwright (2013) considers optimization for nonconvex statistical problems, proving that under some special conditions, nonconvexity may not pose a challenge to recovering optimal parameters. In particular, they consider the following example (Loh and Wainwright, 2011, 2013): suppose that we observe a response $y \in \mathbb{R}^n$ which is generated with a Gaussian linear model,

$$b = Ax_{\text{true}} + \epsilon \text{ with } \epsilon \sim N(0, \sigma^2 \mathbf{I}_n) ,$$

where $A \in \mathbb{R}^{n \times d}$ is a design matrix and $x_{\text{true}} \in \mathbb{R}^d$ is the unknown vector of coefficients. In this case, we might seek to recover x_{true} with the least squares estimator, perhaps with

some penalty added to promote some desired structure in x ,

$$\begin{aligned} \hat{x} &= \arg \min_x \left\{ \frac{1}{2n} \|b - Ax\|_2^2 + \text{Penalty}(x) \right\} \\ &= \arg \min_x \left\{ \frac{1}{2} x^\top \left(\frac{A^\top A}{n} \right) x - x^\top \left(\frac{A^\top b}{n} \right) + \text{Penalty}(x) \right\}. \end{aligned} \quad (13)$$

In some settings, however, the design matrix A itself may not be known with perfect accuracy. Instead, suppose we observe

$$Z = A + W$$

where $W \perp A$ has independent mean-zero entries, with $\mathbb{E}[W_{ij}^2] = \sigma_A^2$ for all i, j . In this case, a naive approach might be to substitute Z for A in (13), before finding the minimizer. However, unless σ_A^2 is negligible, this may not produce a good approximation to \hat{x} since, when substituting $Z^\top Z$ for $A^\top A$ in the quadratic term in (13), we have

$$\mathbb{E}[Z^\top Z \mid A] = \mathbb{E}[(A + W)^\top (A + W) \mid A] = A^\top A + \mathbb{E}[W^\top W] = A^\top A + n\sigma_A^2 \mathbf{I}_d \neq A^\top A.$$

In contrast, for the linear term in (13), we have $\mathbb{E}[Z^\top b \mid A] = A^\top b$, as desired. To correct for the bias in $Z^\top Z$, we should take

$$\hat{x}_{\text{noisy}} = \arg \min_x \{ \mathcal{L}(x) + \text{Penalty}(x) \},$$

where

$$\mathcal{L}(x) := \frac{1}{2} x^\top \left(\frac{Z^\top Z}{n} - \sigma_A^2 \mathbf{I}_d \right) x - x^\top \left(\frac{Z^\top b}{n} \right).$$

Of course, this optimization problem is no longer convex due to the negative quadratic term, and in particular, for a Lipschitz penalty and a high-dimensional setting ($n < d$), the value tends to $-\infty$ as x grows large in any direction in the null space of Z . Remarkably, Loh and Wainwright (2013) show that, if x_{true} is sparse and $\text{Penalty}(x)$ is similar to the ℓ_1 norm, then as long as $(Z^\top Z)$ satisfies a restricted strong convexity assumption (as is standard in the sparse regression literature), then x_{true} can be accurately recovered from *any* local minimum or critical point of the constrained optimization problem

$$\hat{x}_{\text{noisy}} = \arg \min_{\|x\|_1 \leq R} \{ \mathcal{L}(x) + \text{Penalty}(x) \}. \quad (14)$$

The approach taken by Loh and Wainwright (2013) is to perform proximal gradient descent, with steps taking the form

$$x_{t+1} = \arg \min_{\|x\|_1 \leq R} \left\{ \frac{1}{2} \left\| x - \left(x_t - \frac{1}{\eta} \nabla \mathcal{L}(x_t) \right) \right\|_2^2 + \frac{1}{\eta} \text{Penalty}(x) \right\}$$

where $\frac{1}{\eta}$ is a step size parameter. When $\text{Penalty}(x)$ is (some multiple of) the ℓ_1 norm, or some other function with a simple proximal operator (that is, it is simple to compute $\arg \min_x \{ \frac{1}{2} \|x - z\|_2^2 + \text{Penalty}(x) \}$ for any z), this algorithm is very efficient.

2.1.1 TOTAL VARIATION AND GENERALIZED CONVEX ℓ_1 PENALTIES

Consider a setting where the penalty term in (14) is given by a generalized ℓ_1 norm,

$$\text{Penalty}(x) = \nu \cdot \|Kx\|_1$$

for some matrix $K \in \mathbb{R}^{m \times d}$. In particular, if K is the one-dimensional differences matrix $\nabla_{1d} \in \mathbb{R}^{(d-1) \times d}$, which has entries $(\nabla_{1d})_{ii} = 1$ and $(\nabla_{1d})_{i,i+1} = -1$ for each i , then this defines the (one-dimensional) total variation norm on x , $\|x\|_{\text{TV}} = \|\nabla_{1d}x\|_1$; this method is also known as the fused Lasso (Tibshirani et al., 2005). We can also consider a two- or three-dimensional total variation norm, $K = \nabla_{2d}$ or $K = \nabla_{3d}$, defined analogously as the differences matrix for a two- or three-dimensional grid. Total variation type penalties are commonly used in imaging applications and many other fields to obtain solutions that are locally constant or locally smooth.

In this setting, proximal gradient descent is not practical except in some special cases, such as when K is diagonal, because the proximal operator $\arg \min_x \{ \frac{1}{2} \|z - x\|_2^2 + \nu \cdot \|Kx\|_1 \}$ does not have a closed form solution for general K and would itself require an iterative algorithm to be run to convergence. For a total variation penalty on a one-dimensional grid of points, e.g. $K = \nabla_{1d}$, some fast algorithms do exist for the proximal map (Johnson, 2013). Additional methods for convex problems with two-dimensional total variation and related penalties such as total variation over a graph can be found in Chambolle and Darbon (2009); Wang et al. (2014b, 2015b). We are not aware of a non-iterative algorithm for general K . Here we apply MOCCA to allow for arbitrary K and for a nonconvex loss term $\mathcal{L}(x)$.

2.1.2 APPLYING THE MOCCA ALGORITHM

We consider applying the MOCCA algorithm to this nonconvex optimization problem with $\text{Penalty}(x) = \nu \|Kx\|_1$. We define the convex function

$$F(w) = \nu \|w\|_1$$

with the trivial convex approximations $F_v(w) = F(w)$ at any expansion point $v \in \mathbb{R}^m$. We also let

$$G(x) = \begin{cases} \mathcal{L}(x), & \text{if } \|x\|_1 \leq R, \\ +\infty, & \text{if } \|x\|_1 > R, \end{cases}$$

with convex approximations given by taking the linear approximation to the loss,

$$G_z(x) = \begin{cases} \mathcal{L}(z) + \langle x - z, \nabla \mathcal{L}(z) \rangle, & \text{if } \|x\|_1 \leq R, \\ +\infty, & \text{if } \|x\|_1 > R, \end{cases}$$

for any expansion point $z \in \mathbb{R}^d$. Then the optimization problem (14) can be expressed as

$$\hat{x}_{\text{noisy}} = \arg \min_x \{ F(Kx) + G(x) \} .$$

Applying Algorithm 1, the update steps take the form

$$\begin{cases} x_{t+1} = \arg \min_{\|x\|_1 \leq R} \left\{ \|x - [x_t - \text{T}(\nabla \mathcal{L}(x_t) + K^\top w_t)]\|_{\text{T}^{-1}}^2 \right\}, \\ w_{t+1} = \text{Truncate}_\nu(w_t + \Sigma K \bar{x}_{t+1}), \\ z_{t+1} = x_{t+1}, \end{cases} \quad (15)$$

where $\text{Truncate}_\nu(w)$ truncates the entries of the vector w to the range $[-\nu, \nu]$. Note that the x update step is a simple shrinkage step and therefore easy to solve (and $\nabla\mathcal{L}(x_t)$ is simple to compute), while the w and z updates are computationally trivial.

As a second option, we can incorporate more convexity into our approximations G_z by taking

$$G_z(x) = \begin{cases} \mathcal{L}(x) + \frac{\sigma_z^2}{2} \|x - z\|_2^2, & \text{if } \|x\|_1 \leq R, \\ +\infty, & \text{if } \|x\|_1 > R, \end{cases}$$

which is convex since $\mathcal{L}(x)$ has negative curvature bounded by $\frac{\sigma_z^2}{2}$. In this case, after simplifying, our update steps become

$$\begin{cases} x_{t+1} = \arg \min_{\|x\|_1 \leq R} \left\{ \left\| x - \left[x_t - \left(T^{-1} + \frac{Z^\top Z}{n} \right)^{-1} (\nabla\mathcal{L}(x_t) + K^\top w_t) \right] \right\|_{\left(T^{-1} + \frac{Z^\top Z}{n} \right)}^2 \right\}, \\ w_{t+1} = \text{Truncate}_\nu(w_t + \Sigma K \bar{x}_{t+1}), \\ z_{t+1} = x_{t+1}. \end{cases} \quad (16)$$

While the x update step may appear difficult due to the combination of the non-diagonal matrix $\left(T^{-1} + \frac{Z^\top Z}{n} \right)$ which scales the norm, combined with the ℓ_1 constraint, in practice the constraint R is chosen to be large so that it is inactive in all or most steps; the x update step is then solved by $x_{t+1} = x_t - \left(T^{-1} + \frac{Z^\top Z}{n} \right)^{-1} (\nabla\mathcal{L}(x_t) + K^\top w_t)$.

An important point is that MOCCA can be applied to this problem for arbitrary K , including a difference operator such as ∇_{2d} or ∇_{3d} ; in contrast, proximal gradient descent can only be performed approximately except for certain special cases, as mentioned above. We explore this setting's theoretical properties in Section 4.2.3, and give empirical results for this problem in Section 5.

2.2 Isotropic Total Variation and Generalized ℓ_1/ℓ_2 Penalties

For locally constant images or signals in two dimensions, the form of the total variation penalty given above is known as ‘‘anisotropic’’, meaning that it imposes a sparsity pattern which is specific to the alignment of the image onto a horizontal and vertical axis. In contrast, the isotropic total variation penalty (Rudin et al., 1992), on an image x parametrized with values $x_{i,j}$ at grid location (i, j) , is given by

$$\|x\|_{\text{isoTV}} = \sum_{(i,j)} \sqrt{(x_{i,j} - x_{i,j+1})^2 + (x_{i,j} - x_{i+1,j})^2}.$$

Optimization methods for the denoising problem with an isotropic total variation penalty, i.e. problems of the form $\min_x \left\{ \frac{1}{2} \|b - x\|_2^2 + \nu \cdot \|x\|_{\text{isoTV}} \right\}$, were studied in Chambolle and Pock (2015). In practice the isotropic penalty is often preferred as it leads to smoother contours, avoiding the artificial horizontal or vertical edges that may result from anisotropic total variation regularization.

The isotropic total variation penalty can be generalized to penalties of the form

$$\text{Penalty}(x) = \nu \cdot \sum_{\ell=1}^L \|K_\ell x\|_2,$$

where each $K_\ell \in \mathbb{R}^{m_\ell \times d}$ is some fixed matrix. To see how this determines an isotropic total variation penalty in two dimensions, we let ℓ index all locations (i, j) ; then the corresponding matrix K_ℓ has two rows, which when multiplying the image x , extracts the differences $x_{i,j} - x_{i,j+1}$ and $x_{i,j} - x_{i+1,j}$. To see how this specializes to the usual generalized ℓ_1 penalty, $\|Kx\|_1$ for some fixed matrix $K \in \mathbb{R}^{m \times d}$, simply take K_ℓ to be the ℓ th row of the matrix K for each $\ell = 1, \dots, m$.

2.2.1 APPLYING THE MOCCA ALGORITHM

We now show the steps of the MOCCA algorithm to the problem of minimizing

$$\mathcal{L}(x) + \nu \cdot \sum_{\ell=1}^L \|K_\ell x\|_2, \tag{17}$$

where $\mathcal{L}(x)$ is a differentiable likelihood term (such the nonconvex likelihood for regression with errors in variables as above). We define the matrix K by vertically stacking the K_ℓ 's, and define the convex function

$$F(w) = \nu \cdot \sum_{\ell=1}^L \|w_{B_\ell}\|_2,$$

where w_{B_ℓ} is understood to be the ℓ th block of the vector w , i.e.

$$w_{B_\ell} = (w_{m_1+\dots+m_{\ell-1}+1}, \dots, w_{m_1+\dots+m_\ell}).$$

We take the trivial convex approximations $F_v(w) = F(w)$ at any expansion point $v \in \mathbb{R}^m$, and define $G(x)$ and $G_z(x)$ exactly as before (as for the previous application, we allow the option of restricting to $\|x\|_1 \leq R$ if desired). Then the objective function (17) is equivalent to minimizing $F(Kx) + G(x)$.

Applying Algorithm 1, the update steps take the form

$$\begin{cases} x_{t+1} = \arg \min_{\|x\|_1 \leq R} \left\{ \|x - [x_t - T(\nabla \mathcal{L}(x_t) + K^\top w_t)]\|_{T^{-1}}^2 \right\}, \\ w_{t+1} = \text{Truncate}_{\nu \cdot (\mathbb{B}_{m_1} \times \dots \times \mathbb{B}_{m_L})}(w_t + \Sigma K \bar{x}_{t+1}), \\ z_{t+1} = x_{t+1}, \end{cases}$$

where $\text{Truncate}_{\nu \cdot (\mathbb{B}_{m_1} \times \dots \times \mathbb{B}_{m_L})}(w)$ projects each block $w_{B_\ell} \in \mathbb{R}^{m_\ell}$ of the vector w to the ball of radius ν , $\nu \cdot \mathbb{B}_{m_\ell} \subseteq \mathbb{R}^{m_\ell}$ (here \mathbb{B}_{m_ℓ} is the unit ball of dimension m_ℓ in the ℓ_2 norm). Specifically, the w update step can be computed for each block as

$$(w_{t+1})_{B_\ell} = (w_t + \Sigma K \bar{x}_{t+1})_{B_\ell} \cdot \min \left\{ 1, \frac{\nu}{\|(w_t + \Sigma K \bar{x}_{t+1})_{B_\ell}\|_2} \right\}$$

for each $\ell = 1, \dots, L$, and is therefore trivial to compute.

2.3 Nonconvex Total Variation Penalties

As discussed in Section 2.1, total variation penalties are common in many applications where the underlying signal exhibits smooth or locally constant spatial structure (in one, two, or more dimensions). In convex optimization, we are faced with a well-known tradeoff between sparsity and bias—using $\nu \cdot \|x\|_{\text{TV}}$ as our penalty function for some parameter $\nu > 0$, we want to be sure to choose ν large enough that the resulting solution is total-variation-sparse, to avoid overfitting when the sample size is small relative to the dimension of the image; however, larger ν leads to increased shrinkage of the signal, leading to an estimate that is biased towards zero. One way to avoid this tradeoff is to use a nonconvex penalty, which should behave like the total variation norm in terms of promoting sparsity, but reduce the amount of shrinkage for larger signal strength. In this section, we will use ∇_{TV} to denote the differences matrix in the appropriate space (e.g. $\nabla_{\text{TV}} = \nabla_{2\text{d}}$ in two dimensions), so that $\|x\|_{\text{TV}} = \|\nabla_{\text{TV}}x\|_1$.

For sparse regression problems (i.e. where the signal x is itself sparse, rather than sparsity in $\nabla_{\text{TV}}x$), many nonconvex alternatives to the ℓ_1 norm penalty $\|x\|_1$ have been studied, demonstrating more accurate signal recovery empirically as well as theoretical properties of reduced bias, such as the Smoothly Clipped Absolute Deviation (SCAD) penalty (Fan and Li, 2001), the ℓ_q penalty which penalizing $\sum_i |x_i|^q$ for some $q \in (0, 1)$ (Knight and Fu, 2000; Chartrand, 2007), and the Minimax Concave Penalty (MCP) which seeks to minimize concavity while avoiding bias (Zhang, 2010). Another option is to use a reweighted ℓ_1 norm (Candès et al., 2008), where signals estimated to be large at the first pass are then penalized less in the next pass to reduce bias in their estimates; in fact, Candès et al. (2008) show that this procedure is related to a nonconvex log-sum penalty, given by penalizing each component of x_i as $\log(|x_i| + \epsilon)$ for some fixed $\epsilon > 0$. For the problem of total variation sparsity, a variety of nonconvex approaches have also been studied, including applying SCAD (Chopra and Lian, 2010), an ℓ_q norm penalty for $0 < q < 1$ (Sidky et al., 2014; Lu and Huang, 2014), or a log-sum total variation penalty (Selesnick et al., 2015; Parekh and Selesnick, 2015) to the total variation sparsity setting.

We now consider the problem applying a log-sum penalty to the problem of total variation sparsity. Here we consider the form of this penalty given by

$$\log\text{TV}_\beta(x) = \text{logL1}_\beta(\nabla_{\text{TV}}x) \text{ where } \text{logL1}_\beta(w) = \sum_i \beta \log(1 + |w_i|/\beta) ,$$

where $\beta > 0$ is a nonconvexity parameter. (We can also consider applying this nonconvex penalty to the isotropic version of total variation, as discussed in Section 2.2, but for simplicity we do not give that version explicitly here.)

To understand this function, observe that for any t , the function

$$t \mapsto \beta \log(1 + |t|/\beta)$$

is approximately equal to $|t|$ when $t \approx 0$ (that is, near zero it behaves like the ℓ_1 norm), but is nonconvex and penalizes large values of t much less heavily than an absolute value penalty of $|t|$. The parameter β controls the amount of nonconvexity: small β gives a highly nonconvex penalty, while for large β the penalty is nearly convex, with $\log\text{TV}_\beta(x) \approx \|x\|_{\text{TV}}$; see Figure 1 for an illustration.




Figure 1: Illustration of the nonconvex sparsity-promoting penalty discussed in Section 2.3. The figure plots the function $t \mapsto \beta \cdot \log(1 + |t|/\beta)$ across a range of values of t , for $\beta \in \{1, 2, 10, \infty\}$; for $\beta = \infty$, we should interpret this as the absolute value function, $t \mapsto |t|$. We see that all the functions appear similar for $t \approx 0$, with a nondifferentiable point at $t = 0$ which ensures sparsity when this function is used as a penalty. For larger values of t , smaller values of β correspond to greater nonconvexity.

Consider the problem of minimizing an objective function

$$\mathcal{L}(x) + \nu \cdot \log \text{TV}_\beta(x), \tag{18}$$

where $\mathcal{L}(x)$ is some likelihood or loss term. In the image denoising setting, where $\mathcal{L}(x) = \frac{1}{2} \|y - x\|_2^2$ (i.e. when y is a noisy observation of the signal x), Parekh and Selesnick (2015) approach this optimization problem with a majorization/minimization algorithm, iterating the steps: (1) find a majorization of $\log \text{TV}_\beta(x)$ at the current estimate x_t , which takes the form of a reweighted TV norm, (2) compute x_{t+1} as the minimizer of $\mathcal{L}(x) + \nu \cdot$ (the majorized penalty). In other settings, however, for a general loss $\mathcal{L}(x)$, step (2) may not be possible.

We now show how the MOCCA algorithm can be used to optimize objective functions of the form (18). For simplicity we show the steps for the case that $\mathcal{L}(x)$ is convex and has a simple proximal operator, but this can be generalized as needed.

First, we define a new function

$$h_\beta(w) = \log \text{L1}_\beta(w) - \|w\|_1 = \sum_i \beta \log(1 + |w_i|/\beta) - \|w\|_1.$$

We note an important property of this function: $h_\beta(w)$ is differentiable with $(\nabla h_\beta(w))_i = -\frac{w_i}{\beta + |w_i|}$.

For a first approach, we will take

$$F(w) = \nu \cdot \log \text{L1}_\beta(w) \text{ and } G(x) = \mathcal{L}(x).$$

Now we define a local convex approximation to F by writing $F(w) = \nu (\|w\|_1 + h_\beta(w))$ and taking the linear approximation to h_β , namely,

$$F_v(w) = \nu \cdot \|w\|_1 + \nu [h_\beta(v) + \langle w - v, \nabla h_\beta(v) \rangle].$$

And, since by assumption $\mathcal{L}(x)$ is convex and has a simple proximal operator, we can simply define

$$\mathbf{G}_z(x) = \mathbf{G}(x) = \mathcal{L}(x).$$

Then the objective function (18) is equal to $\mathbf{F}(\nabla_{\text{TV}}x) + \mathbf{G}(x)$, and can be minimized with MOCCA. Applying Algorithm 1, the update steps take the form

$$\begin{cases} x_{t+1} = \arg \min_x \left\{ \mathcal{L}(x) + \frac{1}{2} \|x - [x_t - \mathbf{T}\nabla_{\text{TV}}^\top w_t]\|_{\mathbf{T}^{-1}}^2 \right\}, \\ w_{t+1} = \text{Truncate}_\nu(w_t + \Sigma \nabla_{\text{TV}} \bar{x}_{t+1} - \nu \nabla h_\beta(v_t)) + \nu \nabla h_\beta(v_t), \\ v_{t+1} = \Sigma^{-1}(w_t - w_{t+1}) + K \bar{x}_{t+1}. \end{cases}$$

Of course, we have the flexibility to arrange the functions differently if we wish—for example, we could instead define

$$\mathbf{F}(w) = \nu \cdot \|w\|_1 \quad \text{and} \quad \mathbf{G}(x) = \mathcal{L}(x) + \nu \cdot h_\beta(\nabla_{\text{TV}}x).$$

In this case, we will define the local approximations as

$$\mathbf{F}_v(w) = \mathbf{F}(w) = \nu \cdot \|w\|_1$$

and

$$\mathbf{G}_z(w) = \mathcal{L}(x) + \nu [h_\beta(\nabla_{\text{TV}}z) + \langle \nabla_{\text{TV}}(x - z), \nabla h_\beta(\nabla_{\text{TV}}z) \rangle].$$

In this case the objective function (18) is equal to $\mathbf{F}(\nabla_{\text{TV}}x) + \mathbf{G}(x)$ as before. In this case, the update steps are

$$\begin{cases} x_{t+1} = \arg \min_x \left\{ \mathcal{L}(x) + \frac{1}{2} \|x - [x_t - \mathbf{T}\nabla_{\text{TV}}^\top (\nu \nabla h_\beta(\nabla_{\text{TV}}z_t) + w_t)]\|_{\mathbf{T}^{-1}}^2 \right\}, \\ w_{t+1} = \text{Truncate}_\nu(w_t + \Sigma \nabla_{\text{TV}} \bar{x}_{t+1}), \\ z_{t+1} = x_{t+1}. \end{cases}$$

However, in a sense this decomposition is less natural as it splits the penalty $\log \text{TV}_\beta(x)$ across \mathbf{F} and \mathbf{G} , and in fact in our experiments in Section 5, we will see that this second formulation gives poorer convergence results when the MOCCA algorithm is applied.

2.4 Application to CT Image Reconstruction

The initial motivation for developing the algorithm presented here, is a problem arising in computed tomography (CT) imaging. Here we briefly summarize the problem and our approach via the MOCCA algorithm; we describe this setting fully, and give detailed results, in Barber et al. (2016).

In CT imaging, an X-ray beam is sent along many rays through the object of interest. Typically, the measurement is the total energy that has passed through the object, along each ray; comparing the energy retrieved against the energy of the entering X-ray beam, gives information about the materials inside the object, since different materials have different beam attenuation properties. A recent technological development is the photon counting

detector, which measures the raw number of photons that successfully pass through the object rather than the total integrated energy. As a first pass, this transmission model can be written as follows, where the location $\vec{r}_\ell(t)$ inside the object parametrizes the ray ℓ :

$$\text{Count}_\ell \sim \text{Poisson} \left(\int_{\text{energy } E} \left(\begin{array}{c} \text{beam} \\ \text{intensity at} \\ \text{energy } E \end{array} \right) \cdot \left(\begin{array}{c} \text{detector} \\ \text{sensitivity} \\ \text{to energy } E \end{array} \right) \cdot \exp \left\{ - \int_t \mu(E, \vec{r}_\ell(t)) dt \right\} dE \right), \quad (19)$$

where the only unknowns are the coefficients $\mu = (\mu(E, \vec{r}))$, indexed over energy level E and location \vec{r} inside the object. Here $\mu(E, \vec{r})$ is the attenuation for photons at energy E at the location \vec{r} —higher values of $\mu(E, \vec{r})$ indicate that a photon (at energy E) is more likely to be absorbed (at location \vec{r}). The expression $\exp\{-\int_t \mu(E, \vec{r}_\ell(t)) dt\}$ determines, for a photon at energy E that enters the object along the trajectory defined by ray ℓ , the probability that the photon will not be absorbed by the object (i.e. will pass through the object). These coefficients $\mu(E, \vec{r})$ can be further decomposed as

$$\mu(E, \vec{r}) = \sum_{\text{materials } m} \mu_m(E) \cdot x_m(\vec{r}) \quad (20)$$

where $\mu_m(E)$ is a known quantity determining the absorption properties of material m at energy E , while $x_m(\vec{r})$ is unknown, representing the amount of material m that is present at location \vec{r} . In practice, the object space is discretized into pixels, so that x is finite-dimensional.

In this application, the optimization problem is then given by

$$\hat{x} = \arg \min_x \left\{ \mathcal{L}(x) + \sum_{\text{materials } m} (\text{total variation constraint on material map } x_m) \right\},$$

perhaps with other constraints added as well (e.g. nonnegativity of the material maps). Here $\mathcal{L}(x)$ is the negative log-likelihood of x given the Poisson model for the observed photon counts as a function of x , given by (19) and (20). $\mathcal{L}(x)$ is a nonconvex function due to the integration across energy levels. Both $\mathcal{L}(x)$ and the total variation constraints are better represented as functions of linear transformations of x ; the presence of these multiple terms, including nonconvexity (from \mathcal{L}) and nondifferentiability (from total variation), mean that existing methods cannot be applied to solve this optimization problems.

The MOCCA algorithm, applied to this problem, gives strong performance in terms of fast convergence and accurate image reconstruction on simulated and real imaging data. We do not give the details of the algorithm implementation or any empirical results in this paper, but instead refer the reader to our work in Barber et al. (2016) for more details on the method and for empirical results on simulated CT image data (results on real data are forthcoming).

3. Background: Optimizing Convex and Nonconvex Composite Functions

In this sections, we give background on several related algorithms for solving the minimization problem in the simpler setting where F and G are convex, and in the more challenging nonconvex setting.

3.1 Optimization When F is Convex

When F is convex, a variety of existing methods are available to recover the (possibly non-unique) minimizer

$$x^* = \arg \min \{F(Kx) + G(x)\} .$$

In many settings, the functions F and G may each have easily computable proximal operators, but the linear transformation inside $x \mapsto F(Kx)$ might make the function difficult to optimize directly—for example, we might have $F(w) = \|w\|_1$ with $K = \nabla_{2d}$ chosen so that $F(Kx) = \|x\|_{TV}$, the two-dimensional total variation norm of x . For this type of setting, the Alternating Direction Method of Multipliers (ADMM) reframes the problem as

$$(x^*, u^*) = \arg \min_{x, u} \{F(u) + G(x) : Kx = u\}$$

and solves for (x^*, u^*) by working with the augmented Lagrangian

$$\min_{x, u} \max_{\Delta} L(x, u, \Delta) \text{ where } L(x, u, \Delta) = \left\{ F(u) + G(x) + \langle \Delta, Kx - u \rangle + \frac{\rho}{2} \|Kx - u\|_2^2 \right\} .$$

Here $\Delta \in \mathbb{R}^m$ is the dual variable whose role is to enforce the constraint $Kx = u$. The steps of the algorithm are, for each $t \geq 1$,

$$\begin{cases} x_t = \arg \min_x \{L(x, u_{t-1}, \Delta_{t-1})\}, \\ u_t = \arg \min_u \{L(x_t, u, \Delta_{t-1})\}, \\ \Delta_t = \Delta_{t-1} + \rho(Kx_t - u_t). \end{cases}$$

Examining the update step for u , we see that this step entails a single use of the proximal map for F . However, the x update step is more complicated due to the linear operator K ; this step cannot be solved with one use of the proximal map for G , except in the special case that $K^\top K$ is a multiple of the identity matrix. To resolve this, we can add additional curvature to the x update step:

$$x_t = \arg \min_x \left\{ L(x, u_{t-1}, \Delta_{t-1}) + \frac{1}{2} (x - x_{t-1})^\top (\lambda \mathbf{I} - \rho K^\top K) (x - x_{t-1}) \right\},$$

where $\lambda \geq \rho \|K\|^2$. In this setting, the x update step now becomes solveable with the proximal map for G . In fact, this preconditioned form of the ADMM algorithm is equivalent to the CP algorithm (3); for details of the equivalence, see Chambolle and Pock (2011).

In addition to the ADMM and CP algorithms, and their variants, we mention one other option here. If F is differentiable, then proximal gradient descent offers a simple procedure, alternating between taking a gradient descent step on the term $F(Kx)$ and a proximal operator step on G . As we discussed in Section 1.1.1, the proximal gradient descent algorithm can be viewed as a simple special case of MOCCA (in that section, the terms were arranged slightly differently, with both differentiable and convex terms all included in G , but the scenarios are equivalent). Of course, this type of method cannot handle scenarios with a nondifferentiable F , which can arise through total variation penalties and in other settings.

3.2 The Issue of Nonconvexity

Next, we turn to the setting where F and/or G may be nonconvex.

To begin, we consider the following interesting scenario, introduced in Section 1: suppose that G is itself convex, with trivial approximations $G_z = G$, and that $x \mapsto F(Kx)$ is strictly convex even though F is nonconvex. Since $x \mapsto F(Kx) + G(x)$ is therefore strictly convex, the optimization problem has a unique global minimizer $x^* \in \mathbb{R}^d$. However, since F itself is nonconvex, then the strategies for optimization described in Section 3.1 may not be directly applicable for the task of finding x^* .

Here we outline the difficulties faced by the main existing approaches outlined in Section 3.1 for settings where F and/or G are nonconvex (including the scenario outlined above), and summarize the most relevant results in the literature.

To see this, first consider the ADMM algorithm. Suppose that F is nonconvex, and in particular, for some vector $w \in \mathbb{R}^m$, the function $s \mapsto F(s \cdot w)$ is strongly concave—specifically,

$$F(s \cdot w) \leq C - c \cdot s^2 \tag{21}$$

for some $C < \infty, c > 0$ and all $s > 0$. Then we see that the update step

$$u_t = \arg \min_u \{L(x_t, u, \Delta_{t-1})\}$$

is not well-defined; the function $u \mapsto L(x_t, u, \Delta_{t-1})$ diverges to $-\infty$ when we set $u = s \cdot w$ and let $s \rightarrow \infty$. Therefore, for some types of nonconvex functions, the ADMM algorithm will not be implementable due to this divergence.

In the literature, theoretical guarantees for the performance of ADMM on nonconvex objective functions have been considered under several different settings. Broadly speaking, we can summarize existing work as falling into one of two categories. First, there are settings where the original form of the ADMM updates perform well. For instance, Magnússon et al. (2015) proves convergence results (for a slightly different algorithm) when optimization is over a bounded set, thus avoiding the issue of divergence arising from directions of strong concavity in F as mentioned above; this paper also proves results guaranteeing optimality of any limit point, if one exists, in the unbounded optimization setting, but does not guarantee that a limit point is reached. In Li and Pong (2015), convergence results are proved assuming that one of the two terms (i.e. F or G) is smooth; in contrast, for the applications considered here, including total variation type penalties incorporated into F or hard constraints in G , it is critical to allow for nondifferentiable F and G . The special case of applying ADMM to nonconvex consensus problems is considered by Hong et al. (2016), with convergence guarantees again in a bounded setting, in this case assuming that any nonconvex functions must obey a lower bound.

Second, when the original ADMM updates cannot be expected to perform well—if for instance F has directions of strong concavity—then the ADMM can be modified by adding curvature to each update via a Bregman divergence term, as studied by Wang et al. (2014a) (with extensions to multi-block ADMM, with more than two terms in the objective function (Wang et al., 2015a)). This work proves convergence guarantees for the algorithm, but requires that the function F (after converting to our notation) is differentiable and smooth, in contrast to our work where allowing for nondifferentiable F is critical.

Next, consider the CP algorithm. When F is nonconvex, it is no longer the case in general that $F(Kx) = \max_w \{\langle Kx, w \rangle - F^*(w)\}$. In fact, by definition of conjugate functions, this maximum defines the “conjugate of the conjugate”, i.e.

$$F^{**}(Kx) = \max_w \{\langle Kx, w \rangle - F^*(w)\} .$$

It is known that any conjugate function must be convex, i.e. F^{**} is convex. This implies that $F^{**} \neq F$. Therefore, the saddle-point problem does not correspond to the original optimization problem: we have

$$\min_x \max_w \{\langle Kx, w \rangle - F^*(w) + G(x)\} = \min_x \{F^{**}(Kx) + G(x)\} ,$$

which is different from the original optimization problem since $F^{**} \neq F$. If $F^*(w)$ and $F^{**}(w)$ take finite values on some domain, then the CP algorithm can be expected to converge, but it will converge to the solution of an optimization problem that is different from the one intended due to the issue that $F^{**} \neq F$. Problems also arise in a setting where F exhibits negative curvature as in (21), in which case $F^*(w) = \infty$ for all w , so we do not have a well-defined saddle point problem to begin with.

To our knowledge, no general results exist for the CP algorithm with nonconvex and nondifferentiable F and/or G . Valkonen (2014) considers an interesting related problem, namely, a variant of the CP algorithm, where the objective function is now $F(\mathcal{K}(x)) + G(x)$, for convex F, G but with a nonlinear map $\mathcal{K}(x)$ in place of the previous linear map Kx , as the argument to F . In this case, convergence to a stationary point is proved, even when the nonlinearity of $\mathcal{K}(x)$ may make the overall problem nonconvex. Relatedly, Ochs et al. (2015) study the setting where $F(\mathcal{K}(x))$ is a nonconvex elementwise penalty on the convex transform $\mathcal{K}(x)$ while G is convex; their approach uses the CP algorithm as a subroutine for solving convex approximations of the objective function, at each step.

Next, we consider the option of proximal gradient descent, in the case that G has a simple proximal operator. Loh and Wainwright (2013) study penalized likelihood problems,

$$\min_x \{\mathcal{L}(x) + \text{Penalty}(x)\} ,$$

where the likelihood term $\mathcal{L}(x)$ and/or the penalty term $\text{Penalty}(x)$ may exhibit nonconvexity. In many settings that arise in high-dimensional statistics, for instance, the likelihood term $\mathcal{L}(x)$ may be strongly concave in some directions, but will be strongly convex in all “plausible” directions, that is, all directions x that are not prohibited by the penalty term $\text{Penalty}(x)$. For instance, if $\text{Penalty}(x)$ is a sparsity-promoting penalty, with low values only at solutions x with many (near-)zero values, then $\mathcal{L}(x)$ might be strongly convex in all sparse directions. This relates to the notion of *restricted strong convexity*, introduced by Negahban et al. (2009), which we discuss in greater detail in Section 4.2.1. Under restricted convexity and smoothness assumptions on the likelihood term, and with bounds on the amount of nonconvexity allowed in the penalty term, Loh and Wainwright (2013) prove convergence to a point that is near the global optimum, for a proximal gradient descent method, with some additional details restricting steps to a bounded set to avoid diverging towards directions of strong concavity. Ochs et al. (2014)’s iPiano method gives an inertial (i.e. accelerated) proximal gradient descent method for this same setting where the loss is

differentiable with the penalty has an easy proximal map. The (accelerated) proximal gradient descent method for nonconvex problems is studied also by Ghadimi and Lan (2016); Li and Lin (2015). Note that these algorithms are applicable only when the terms in the objective function are all either differentiable (the likelihood) or have an easy-to-compute proximal operator (the penalty), and therefore, cannot be applied to many of the problems that that we have considered.

Finally, Bolte et al. (2014) propose an algorithm, Proximal Alternating Linearized Minimization (PALM), to solve a related problem of the form

$$\min_{x,w} \{F(w) + G(x) + H(w, x)\}, \quad (22)$$

where H is differentiable while F, G each have easy to compute proximal maps (and may be nonconvex); this formulation is related to the problem we study, but PALM cannot be used to solve general problems of the form $F(Kx) + G(x)$ where F is nondifferentiable (since, if we add a variable $w = Kx$, the constraint $w = Kx$ cannot be enforced with any differentiable function $H(w, x)$ unless we allow modifications such as a relaxation to a penalty on $\|w - Kx\|_2^2$), and therefore again cannot be applied to some of the problems considered here.

3.3 Connection Between MOCCA and ADMM

In the convex setting, the CP method with parameter $\theta = 1$ is known to be equivalent to a preconditioned ADMM algorithm (Chambolle and Pock, 2011). Specifically, reformulating the original optimization problem in the ADMM form

$$\min_{x,u} \{F(u) + G(x) : u = Kx\} = \min_{x,u} \max_{\Delta} \left\{ F(u) + G(x) + \langle \Delta, Kx - u \rangle + \frac{1}{2} \|Kx - u\|_{\Sigma}^2 \right\},$$

the CP iterations given in (3) are equivalent to the following preconditioned ADMM iterations, where we choose a preconditioning matrix $T^{-1} - K^{\top} \Sigma K \succeq 0$:

$$\begin{cases} x_{t+1} = \arg \min_x \left\{ G(x) + \langle K^{\top} \Delta_t, x \rangle + \frac{1}{2} \|Kx - u_t\|_{\Sigma}^2 + \frac{1}{2} \|x - x_t\|_{T^{-1} - K^{\top} \Sigma K}^2 \right\}, \\ \Delta_{t+1} = \Delta_t + \Sigma(Kx_{t+1} - u_t), \\ u_{t+1} = \arg \min_u \left\{ F(u) - \langle \Delta_{t+1}, u \rangle + \frac{1}{2} \|Kx_{t+1} - u\|_{\Sigma}^2 \right\}, \end{cases}$$

(Here we use a slightly nonstandard indexing, writing the variable updates in the order x, Δ, u for later convenience.) We do not derive the equivalence here (see Chambolle and Pock, 2011, for details), but remark that the variables (x_t, w_t) at iteration t of the CP algorithm (3) can be recovered by taking $(x_t, \Sigma(Kx_t - u_t) + \Delta_t)$ from the ADMM iterations.

Similarly, in the more general nonconvex setting considered here, we can equivalently formulate the MOCCA method as a combination of the preconditioned ADMM iterations and taking convex expansions to F and G . The steps are given in Algorithm 3. This algorithm is exactly equivalent to the basic version of MOCCA, Algorithm 1, with extrapolation parameter $\theta = 1$, but is expressed as an ADMM type algorithm with preconditioning and with convex approximations to F and G . (The stable “inner loop” version of MOCCA, Algorithm 2, can also be interpreted as an extension of ADMM, but we do not give details here). The

equivalence between Algorithms 1 and 3 is simply an extension of the connection between the Chambolle-Pock algorithm and a preconditioned ADMM, as shown in Chambolle and Pock (2011).

Algorithm 3 MOCCA algorithm: ADMM version

Input: Convex functions $F_{\text{cvx}}, G_{\text{cvx}}$, differentiable functions $F_{\text{diff}}, G_{\text{diff}}$, linear operator K , positive diagonal step size matrices Σ, T .

Initialize: Primal variables $x_0 \in \mathbb{R}^d$, $u_0 \in \mathbb{R}^m$, dual variable $\Delta_0 \in \mathbb{R}^m$.

for $t = 0, 1, 2, \dots$ **do**

Update all variables:

$$\begin{cases} x_{t+1} = \arg \min_x \left\{ G_{x_t}(x) + \langle x, K^\top \Delta_t \rangle + \frac{1}{2} \|Kx - u_t\|_\Sigma^2 + \frac{1}{2} \|x - x_t\|_{T^{-1} - K^\top \Sigma K}^2 \right\}, \\ \Delta_{t+1} = \Delta_t + \Sigma(Kx_{t+1} - u_t), \\ u_{t+1} = \arg \min_u \left\{ F_{u_t}(u) - \langle \Delta_{t+1}, u \rangle + \frac{1}{2} \|Kx_{t+1} - u\|_\Sigma^2 \right\}, \end{cases}$$

until some convergence criterion is reached.

This ADMM formulation of MOCCA gives us a clearer understanding of the choice of the expansion points (z, v) in the MOCCA algorithm. Recalling the simpler form of the MOCCA algorithm given in Algorithm 1 where the expansion points are updated at each iteration (i.e. there is no “inner loop”), the expansion points were defined as

$$z_{t+1} = x_{t+1} \text{ and } v_{t+1} = \Sigma^{-1}(w_t - w_{t+1}) + K\bar{x}_{t+1}.$$

Using our conversion between the CP variables (x, w) and the ADMM variables (x, Δ, u) , we see that

$$\begin{aligned} v_{t+1} &= \Sigma^{-1}(w_t - w_{t+1}) + K(2x_{t+1} - x_t) \quad \text{since } \theta = 1 \text{ so } \bar{x}_{t+1} = 2x_{t+1} - x_t \\ &= \Sigma^{-1}(\Sigma(Kx_t - u_t) + \Delta_t - \Sigma(Kx_{t+1} - u_{t+1}) - \Delta_{t+1}) + K(2x_{t+1} - x_t) \\ &= \Sigma^{-1}\left(\Sigma(Kx_t - u_t) + \Delta_t - \Sigma(Kx_{t+1} - u_{t+1})\right. \\ &\quad \left. - (\Delta_t + \Sigma(Kx_{t+1} - u_t))\right) + K(2x_{t+1} - x_t) \\ &= u_{t+1}, \end{aligned}$$

where the next-to-last step uses the definition of the update step for Δ_{t+1} . In other words, after the t th step, our estimated minimizers for $\{F(u) + G(x) : Kx = u\}$ are given by u_{t+1} and x_{t+1} , and our convex approximations to F and to G for the next step are consequently taken at the values $v = u_{t+1}$ and $z = x_{t+1}$.

4. Theoretical Results

In this section we present our two main theoretical results.

4.1 Convergence to a Critical Point

First, we show that if the algorithm converges, then its limit point is a solution to our original problem.

Theorem 1 *Assume that the families of approximations F_v and G_z satisfy (4), and furthermore that*

$$\begin{cases} (v, w) \mapsto \nabla(F - F_v)(w) \text{ is continuous jointly in } (v, w), \text{ and} \\ (z, x) \mapsto \nabla(G - G_z)(x) \text{ is continuous jointly in } (z, x). \end{cases} \quad (23)$$

Suppose that Algorithm 2 converges to a point, with

$$\begin{aligned} x_{t;\ell} &\rightarrow \hat{x} \quad \text{where the sequence } (x_{t;\ell}) \text{ is interpreted as } (x_{1;0}, x_{1;1}, \dots, x_{1;L_1}, x_{2;0}, \dots), \\ w_{t;\ell} &\rightarrow \hat{w} \quad \text{where } (w_{t;\ell}) \text{ is interpreted analogously,} \\ z_t &\rightarrow \hat{z}, \\ v_t &\rightarrow \hat{v}. \end{aligned}$$

Then \hat{x} is a critical point of the original optimization problem, in the sense that

$$0 \in K^\top \partial F_{K\hat{x}}(K\hat{x}) + \partial G_{\hat{x}}(\hat{x}).$$

4.2 Guarantees of Convergence

We now turn to theoretical results proving that the algorithm converges (and proving rates of convergence) under specific assumptions on F and G . In this section, we only consider the “inner loop” form of MOCCA, given in Algorithm 2. We show that if our inner loop length (i.e. L_t) tends to infinity, then we can bound the error of the algorithm.

We begin with an assumption on the step size parameters:

Assumption 1 *The extrapolation parameter is set at $\theta = 1$, and the diagonal matrices Σ and T are chosen such that*

$$M = \begin{pmatrix} T^{-1} & -K^\top \\ -K & \Sigma^{-1} \end{pmatrix} \succ 0.$$

Pock and Chambolle (2011) introduce this assumption for the (convex) preconditioned Chambolle-Pock algorithm, and give a simple construction for one choice of Σ, T to satisfy this without calculating any matrix norms or other high-cost operations, specified in (8) above.

Next, we turn to the convexity and smoothness assumptions required for our convergence guarantee.

4.2.1 RESTRICTED CONVEXITY AND SMOOTHNESS

In practice, F and/or G may each consist of multiple terms, combining characteristics of the problem such as a likelihood calculation or a penalty or constraint on the underlying signal. To accommodate a range of potential applications, in particular those arising in the regression and imaging applications described in Section 2, we consider a broad setting where our main assumptions involve the interplay between convexity and negative curvature in the functions F, G .

The notion of restricted strong convexity (RSC), introduced by Negahban et al. (2009), has often been used in high-dimensional statistics to express the idea that likelihood functions and optimization problems, which may not have desirable strong convexity properties

globally, nonetheless exhibit strong convexity in “directions of interest”. For example, in a least-squares regression problem with design matrix $A \in \mathbb{R}^{n \times d}$, with $n \ll d$, the least squares loss function $\mathcal{L}(x) = \frac{1}{2} \|y - Ax\|_2^2$ is not strongly convex since $A^\top A$ is rank deficient, but can yield good statistical properties if $A^\top A$ is strongly convex in all sparse directions, that is, $x^\top A^\top A x \geq c \cdot \|x\|_2^2$ for all sparse (or approximately sparse) vectors x . In this case, the loss function is globally convex, but it is the RSC property that ensures high accuracy for sparse regression problems. More recently, Loh and Wainwright (2013) proved that the RSC property, along with an analogous restricted smoothness property, can in fact be leveraged even in nonconvex optimization problems, such as the regression-with-errors-in-variables scenario described in Section 2.1. Their work relies on optimizing the variable x within some bounded set, to ensure that the RSC property will push x towards a good (local) minimum rather than allowing x to diverge. For instance, if the loss function has some directions of strong concavity—as is the case for regression-with-errors-in-variables in (14)—then staying within a bounded set is critical. In theory, their work focuses on problems that take the form of minimizing a penalized loss function over a bounded set $\{x : \|x\|_1 \leq R\}$, where we think of R as a large bound, requiring only a loose bound on the ℓ_1 norm of the true signal. In practice, if an optimization algorithm is initialized at zero, then it is often the case that the iterations will never leave a bounded region, without imposing any explicit constraint.

In general, results using the RSC condition take the following form: first, the loss function or objective function $\mathcal{L}(x)$ is shown to satisfy a RSC property of the form

$$\langle x - x', \nabla \mathcal{L}(x) - \nabla \mathcal{L}(x') \rangle \geq c \|x - x'\|_2^2 - \tau^2 \|x - x'\|_{\text{restrict}}^2 ,$$

for some structured norm $\|\cdot\|_{\text{restrict}}$ (for example, the ℓ_1 norm). Here $c > 0$ is a constant while τ is vanishingly small, for instance $c \sim 1$ and $\tau \sim \sqrt{\frac{\log(d)}{n}}$ in many high dimensional regression applications with sample size n . The solution \hat{x} is then shown to converge to the true signal x^* up to an error of size τR , where R is some bound on the signal complexity, for instance $R \sim \sqrt{k}$ where k is the true sparsity level of a sparse regression problem. In these settings, it is assumed that $\tau R = o(1)$, and that errors of this magnitude are negligible. We will follow this general framework in our convergence guarantee as well. However, since we consider settings where the signals may not have natural sparsity but would instead have a different type of structure (such as total variation sparsity), we replace the ℓ_1 norm with a general measure of signal complexity, $\|\cdot\|_{\text{restrict}}$, chosen with respect to the problem at hand (for instance, a total variation norm).

We now specify our assumptions on convexity and smoothness for the functions involved in the optimization, using the restricted strong convexity / restricted smoothness framework from the literature. Roughly speaking, the following assumption requires that the errors of the convex approximations $F - F_v$, $G - G_z$ are counterbalanced by strong convexity in the composite approximations $F_v(Kx) + G_z(x)$. For the term G , we allow for some flexibility by considering restricted strong convexity and restricted smoothness, relative to the structured norm $\|x\|_{\text{restrict}}$.

Assumption 2 *The approximations F_v and G_z satisfy the conditions (4), with additional assumptions as follows. For the function F and its family of local approximations F_v , we*

assume that F_v is strongly convex, while $F - F_v$ is smooth: for all u, v, w ,³

$$\begin{cases} \text{Strong convexity of } F_v: \langle u, \partial F_v(w+u) - \partial F_v(w) \rangle \geq \|u\|_{\Lambda_F}^2, \\ \text{Smoothness of } F - F_v: |\langle u, \nabla(F - F_v)(v+w) \rangle| \leq \frac{1}{2} \left(\|u\|_{\Theta_F}^2 + \|w\|_{\Theta_F}^2 \right), \end{cases}$$

for some $\Lambda_F, \Theta_F \succeq 0$.⁴ We also assume a gradient condition on F_v ,

$$\left\{ F_v \text{ satisfies a gradient condition: } \|\partial F_v(w) - \partial F_v(w')\|_2 \leq C_{\text{Lip}} + C_{\text{grad}} \|w - w'\|_2, \right.$$

for some $C_{\text{Lip}}, C_{\text{grad}} < \infty$. (For example, this is satisfied if F_v can be written as the sum of a Lipschitz function and a smooth function.)

For the function G and its family of local approximations G_z , we assume that G_z satisfies restricted strong convexity, while $G - G_z$ satisfies a restricted smoothness assumption: for all x, y, z ,

$$\begin{cases} \text{Restricted strong convexity of } G_z: \langle y, \partial G_z(x+y) - \partial G_z(x) \rangle \geq \|y\|_{\Lambda_G}^2 - \tau^2 \|y\|_{\text{restrict}}^2, \\ \text{Restricted smoothness of } G - G_z: \\ \quad |\langle y, \nabla(G - G_z)(z+x) \rangle| \leq \frac{1}{2} \left(\|x\|_{\Theta_G}^2 + \|y\|_{\Theta_G}^2 \right) + \frac{\tau^2}{2} \left(\|x\|_{\text{restrict}}^2 + \|y\|_{\text{restrict}}^2 \right), \end{cases}$$

for some $\Lambda_G, \Theta_G \succeq 0$ and $\tau < \infty$.

Finally, the total convexity in the local approximations F_v and G_z must (approximately) outweigh the total curvature of the differences $F - F_v$ and $G - G_z$. Specifically, for all $x \in \mathbb{R}^d$, we require

$$x^\top (K^\top \Lambda_F K + \Lambda_G) x \geq x^\top (K^\top \Theta_F K + \Theta_G) x + C_{\text{cvx}} \|x\|_2^2 - \tau^2 \|x\|_{\text{restrict}}^2,$$

for some $C_{\text{cvx}} > 0$ and $\tau < \infty$.

In general, greater convexity (i.e. Λ_F, Λ_G as strongly positive definite as possible) and tighter bounds on smoothness (i.e. Θ_F, Θ_G as small as possible) allow for a better (i.e. larger) constant C_{cvx} and, therefore, faster convergence of the algorithm. The value of τ is typically of a very small order in many problems arising in high-dimensional statistics, as discussed above for the sparse regression setting.

It is critical to note that this assumption does *not* require either F_v or G_z to be strictly convex—if the matrices Λ_F or Λ_G are not full rank, then strict convexity has not been assumed. Instead, F_v is strongly convex in any direction of \mathbb{R}^m contained in the column span of Λ_F , and similarly for G_z and Λ_G in \mathbb{R}^d . Our assumption essentially requires that the combination of these directions leads to overall (approximate) convexity, after accounting for concavity that might be introduced by the errors $F - F_v$ and $G - G_z$.

For simplicity in the statements and proofs of our results, we group the norms of all matrices from Assumptions 1 and 2 into a single constant:

$$C_{\text{matrix}} = \max \{ \|\Lambda_F\|, \|\Theta_F\|, \|\Lambda_G\|, \|\Theta_G\|, \|M\|, \|M^{-1}\| \}.$$

Throughout, we will treat $C_{\text{matrix}}, C_{\text{cvx}}, C_{\text{Lip}}$, and C_{grad} as fixed finite positive constants, and dependence on these values will not be given explicitly except in the proofs. On the other hand, the role of the restricted convexity/smoothness parameter τ will be shown explicitly.

3. We implicitly restrict all variables to the domain of the appropriate function, throughout.
 4. For a function F with a multivalued subdifferential, this notation is taken to mean that the statement must hold for any elements of the subdifferentials, throughout the paper.

4.2.2 CONVERGENCE GUARANTEE

Choose any point $x^* \in \mathbb{R}^d$ with $\|x^*\|_{\text{restrict}} \leq R$, which is a critical point for the optimization problem

$$\min_{\|x\|_{\text{restrict}} \leq R} \{F(Kx) + G(x)\} .$$

For convenience, we will now absorb the constraint $\|x\|_{\text{restrict}} \leq R$ into the functions themselves, by replacing G with the function

$$x \mapsto \begin{cases} G(x), & \text{if } \|x\|_{\text{restrict}} \leq R, \\ +\infty, & \text{if } \|x\|_{\text{restrict}} > R, \end{cases}$$

and replacing G_z (for each z) with the function

$$x \mapsto \begin{cases} G_z(x), & \text{if } \|x\|_{\text{restrict}} \leq R, \\ +\infty, & \text{if } \|x\|_{\text{restrict}} > R, \end{cases}$$

In practice, as mentioned before, we typically do not need to explicitly incorporate this constraint into the optimization algorithm, as we will generally only see updates that all lie within a bounded region. However, in our statements and proofs of theoretical results from this point onward, we will assume that G, G_z restrict the domain of the variable x , that is,

$$G(x) = G_z(x) = +\infty \text{ whenever } \|x\|_{\text{restrict}} > R . \quad (24)$$

We will also assume that τR is bounded by a constant without further comment; since our results give convergence guarantees up to the accuracy level τR , the results are meaningful only if τR is small.

We now state our convergence guarantee for the stable form of the MOCCA algorithm, given in Algorithm 2:

Theorem 2 *Assume that Assumptions 1 and 2, and that G, G_z satisfy (24). Then there exists constants $C_{\text{converge}}, L_{\text{min}} < \infty$ and $\delta > 0$, such that if $\min_{t \geq 1} L_t \geq L_{\text{min}}$, then for all $t \geq 1$, the iterations of Algorithm 2 satisfy*

$$\|x_t - x^*\|_2 \leq C_{\text{converge}} \left(\frac{1}{\sqrt{L'_t}} + \tau R \right)$$

where

$$L'_t = \min \{L_t, (1 + \delta)L_{t-1}, \dots, (1 + \delta)^{t-1}L_1\} .$$

As an example, we can set $L_t \sim (1 + \delta)^t$. Then after the t th inner loop, $\|x_t - x^*\|_2 \sim (1 + \delta)^{-t/2} + \tau R$, and the total number of iterations taken is $L_1 + \dots + L_t \sim (1 + \delta)^t$. In other words, the error $\|x - x^*\|_2$ scales as $\frac{1}{\sqrt{T}} + \tau R$ where T is the total number of update steps, i.e. error is inversely proportional to the square root of computational cost, up to the accuracy level τR .

We also remark that, if we were to assume additionally that F is differentiable and smooth, the result would improve dramatically: we would obtain error decaying exponentially in the number of update steps (up to the accuracy level τR). The resulting convergence

guarantee would then be comparable to the results obtained in Loh and Wainwright (2013, Theorem 3), which show a result with error at time t scaling as $c^t + \tau R$ (for a constant $c < 1$). However, convergence in this setting is of limited interest for the applications we have in mind, since total variation penalties, and many other natural penalties or losses falling into the F term of the composite objective functions, are not differentiable.

4.2.3 CONVEXITY AND SMOOTHNESS ASSUMPTIONS: AN EXAMPLE

To illustrate the many different matrices and constants appearing in Assumption 2 with a concrete example, we return to the problem studied in Section 2.1, where a least squares regression with errors in variables is combined with a total variation (or generalized ℓ_1) penalty. Recalling this setting, we seek to minimize $\mathcal{L}(x) + \nu \cdot \|Kx\|_1$, and we set $F_v(w) = F(w) = \nu \cdot \|w\|_1$, and

$$\mathbf{G}(x) = \begin{cases} \mathcal{L}(x), & \text{if } \|x\|_1 \leq R, \\ +\infty, & \text{if } \|x\|_{\text{restrict}} > R, \end{cases}$$

and

$$\mathbf{G}_z(x) = \begin{cases} \mathcal{L}(x) + \frac{\sigma_A^2}{2} \|x - z\|_2^2, & \text{if } \|x\|_{\text{restrict}} \leq R, \\ +\infty, & \text{if } \|x\|_1 > R, \end{cases}$$

where

$$\mathcal{L}(x) = \frac{1}{2} x^\top \left(\frac{Z^\top Z}{n} - \sigma_A^2 \mathbf{I}_d \right) x - x^\top \left(\frac{Z^\top b}{n} \right).$$

Under the Gaussian noise model, the noisy design matrix given by entries $Z_{ij} = A_{ij} + \text{Normal}(0, \sigma_A^2)$ and the response is given by $b = A \cdot x_{\text{true}} + \text{Normal}(0, \sigma^2 \mathbf{I})$. The MOCCA update steps for this problem are given in (16).

In this setting, Assumption 2 is satisfied with the following parameters. First, since F is convex but not strongly convex, we set $\Lambda_F = \mathbf{0}$; we can also set $\Theta_F = \mathbf{0}$ as $F_v = F$ for any expansion point v . F is ν -Lipschitz so we can take $C_{\text{Lip}} = 2\nu, C_{\text{grad}} = 0$. Next, for \mathbf{G} , we see that

$$\mathbf{G}_z(x) = \frac{1}{2} x^\top \left(\frac{Z^\top Z}{n} \right) x - x^\top \left(\frac{Z^\top b}{n} + \sigma_A^2 z \right)$$

(on the domain $\|x\|_{\text{restrict}} \leq R$), and so we can set $\Lambda_G = \frac{Z^\top Z}{n}$. To check the smoothness condition, we have

$$\langle y, \nabla(\mathbf{G} - \mathbf{G}_z)(z + x) \rangle = \langle y, \sigma_A^2 x \rangle \leq \sigma_A^2 \cdot \frac{1}{2} (\|x\|_2^2 + \|y\|_2^2),$$

and so we can take $\Theta_G = \sigma_A^2 \mathbf{I}_d$.

Finally, for the ‘‘total convexity’’ condition of Assumption 2, we need to check that $K^\top (\Lambda_F - \Theta_F) K + (\Lambda_G - \Theta_G) = \frac{Z^\top Z}{n} - \sigma_A^2 \mathbf{I}_d$ satisfies restricted strong convexity. In Loh and Wainwright (2011, Corollary 1), it is shown that if the rows of the (original) design matrix A are drawn i.i.d. from a subgaussian distribution with covariance Σ_A then, assuming that the sample size n satisfies $n \gg \log(d)$, the matrix $\frac{Z^\top Z}{n} - \sigma_A^2 \mathbf{I}_d$ (which is an unbiased estimate of the desired term $\frac{A^\top A}{n}$ using the unknown original design matrix A) satisfies

$$x^\top \left(\frac{Z^\top Z}{n} - \sigma_A^2 \mathbf{I}_d \right) x \geq \frac{1}{2} \lambda_{\min}(\Sigma_A) \cdot \|x\|_2^2 - (\text{constant}) \cdot \frac{\log(d)}{n} \cdot \|x\|_{\text{restrict}}^2 \quad \text{for all } x \in \mathbb{R}^d \quad (25)$$

with high probability, when we choose $\|x\|_{\text{restrict}} = \|x\|_1$; similar results will hold for other structured choices of $\|\cdot\|_{\text{restrict}}$ such as total variation norm or a generalized ℓ_1 norm. Therefore we can set $C_{\text{cvx}} = \frac{1}{2}\lambda_{\min}(\Sigma_A)$ and $\tau \sim \sqrt{\frac{\log(d)}{n}}$ to obtain the desired condition in Assumption 2. Note that the guarantees of Theorem 2 give a meaningful convergence result even if we choose a fairly large radius R .

5. Experiments

We now implement the MOCCA algorithm to examine its performance in practice. Throughout this section, we work with the simpler formulation of MOCCA, given in Algorithm 1, with no “inner loop”. All computations were performed in MATLAB (MATLAB, 2015).⁵

For all simulations, we choose not to place a bound on $\|x\|_{\text{restrict}}$, although technically this is required by our convergence guarantees and those of the related results in Loh and Wainwright (2013) (which we compare to, in Simulation 2). Empirically we observe good convergence without imposing such a bound, but can easily add such a bound if desired.

We consider two examples: Simulation 1 studies nonconvex total variation regularization with a least squares loss (as described in Section 2.3), and Simulation 2 considers convex total variation regularization with a nonconvex loss arising from regression with errors in variables (as described in Section 2.1). While other algorithms which are developed specifically for these problems are available—for example, denoising with total variation penalties is studied by e.g. Chambolle and Darbon (2009); Wang et al. (2014b, 2015b), and could be combined with a proximal gradient method for Simulation 2—here our purpose is simply to illustrate applications of MOCCA to several concrete examples in order to demonstrate its flexibility for a broad range of problems. Specific problems will often have specialized algorithms which would far outperform our general-purpose method; however, slight modifications to the optimization problem (for example, replacing total variation regularization with a more general penalty $\|Kx\|_1$ for a generic dense matrix K , or with isotropic total variation) will often mean that specialized algorithms can no longer be applied, while MOCCA can adapt easily to accommodate these changes.

5.1 Simulation 1: Nonconvex Total Variation Penalty

In the first simulation, we study the nonconvex total variation penalty considered in Section 2.3, using a two-dimensional spatial structure. We generate data as follows: first, we define the true signal $x_{\text{true}} \in \mathbb{R}^d$ with dimension $d = 625$, obtained by vectorizing the two-dimensional locally constant array

$$\begin{pmatrix} \mathbf{1}_{5 \times 5} & \mathbf{0}_{5 \times 15} & \mathbf{0}_{5 \times 5} \\ \mathbf{0}_{15 \times 5} & \mathbf{1}_{15 \times 15} & \mathbf{0}_{15 \times 5} \\ \mathbf{0}_{5 \times 5} & \mathbf{0}_{5 \times 15} & \mathbf{1}_{5 \times 5} \end{pmatrix} \in \mathbb{R}^{25 \times 25}.$$

The two-dimensional total variation of the true signal is very low, because $\nabla_{2d} x_{\text{true}}$ is sparse. We then take a linear regression model with $n = 200$ observations, with design matrix

5. Code for fully reproducing these simulations is available at <http://www.stat.uchicago.edu/~rina/mocca.html>.

$A \in \mathbb{R}^{n \times d}$ with $A_{ij} \stackrel{\text{iid}}{\sim} \text{Normal}(0, 1)$ and $b \in \mathbb{R}^n$ with entries

$$b_i = (A \cdot x_{\text{true}})_i + \text{Normal}(0, 1) .$$

We would then like to solve a penalized least-squares problem using the nonconvex total variation penalty introduced in Section 2.3, namely,

$$\hat{x} = \arg \min_x \{ \text{Obj}(x) \} \text{ for } \text{Obj}(x) = \frac{1}{2} \|b - Ax\|_2^2 + \nu \cdot \log \text{TV}_\beta(x) , \quad (26)$$

where we choose penalty parameter $\nu = 20$ and nonconvexity parameter $\beta = 3$ (recall that a low value of β corresponds to greater nonconvexity), and where the $\log \text{TV}_\beta(\cdot)$ penalty is defined with respect to two-dimensional total variation—recall

$$\log \text{TV}_\beta(x) = \log \text{L1}_\beta(\nabla_{2d}x) = \sum_i \beta \log(1 + |(\nabla_{2d}x)_i|/\beta) .$$

Here $\nabla_{2d} \in \mathbb{R}^{m \times d}$ is the two-dimensional first differences matrix for the vectorized $d_1 \times d_2$ grid, where $d = d_1 \cdot d_2$ is the total dimension of the signal while $m = d_1(d_2 - 1) + d_2(d_1 - 1)$ is the number of first-order differences measured; in our case, we have $d_1 = d_2 = 25$, $d = 625$, and $m = 1200$.

Next, we implement the MOCCA algorithm with the two variants described in Section 2.3: setting $K = \nabla_{2d}$, we consider the more natural form where the penalty term is contained in F , given by

$$\begin{cases} F(w) = \nu \cdot \log \text{L1}_\beta(w), \text{ with } F_v(w) = \nu \cdot \|w\|_1 + \nu [h_\beta(v) + \langle w - v, \nabla h_\beta(v) \rangle], \\ G(x) = G_z(x) = \frac{1}{2} \|b - Ax\|_2^2, \end{cases} \quad (27)$$

where $h_\beta(w) = \log \text{L1}_\beta(w) - \|w\|_1$ is a differentiable concave function as discussed in Section 2.3. We also consider the less natural form where the penalty term is split across F and G , given by

$$\begin{cases} F(w) = F_v(w) = \nu \cdot \|w\|_1, \\ G(x) = \frac{1}{2} \|b - Ax\|_2^2 + \nu \cdot h_\beta(\nabla_{2d}x), \\ \text{with } G_z(x) = \frac{1}{2} \|b - Ax\|_2^2 + \nu [h_\beta(\nabla_{2d}z) + \langle \nabla_{2d}(x - z), \nabla h_\beta(\nabla_{2d}z) \rangle], \end{cases} \quad (28)$$

We will refer to these two versions as MOCCA(natural) and MOCCA(split), respectively. Finally, we choose step size parameters

$$\Sigma = \lambda \cdot \frac{1}{2} \mathbf{I}_m \quad \text{and} \quad T = \lambda^{-1} \cdot \frac{1}{4} \mathbf{I}_d ,$$

which ensures that the positive semidefinite assumption, Assumption 1, will hold (although perhaps not strictly) as in Pock and Chambolle (2011). We test the algorithm across a range of λ values, $\lambda \in \{4, 8, 16, 32, 64\}$.

The results are shown in Figure 2, which plots the log value of the objective function $\text{Obj}(x)$ at each iteration, and also plots the log of the change in each iteration,

$$\text{Change}_t = \left\| \begin{pmatrix} x_{t-1} - x_t \\ w_{t-1} - w_t \end{pmatrix} \right\|_2 . \quad (29)$$

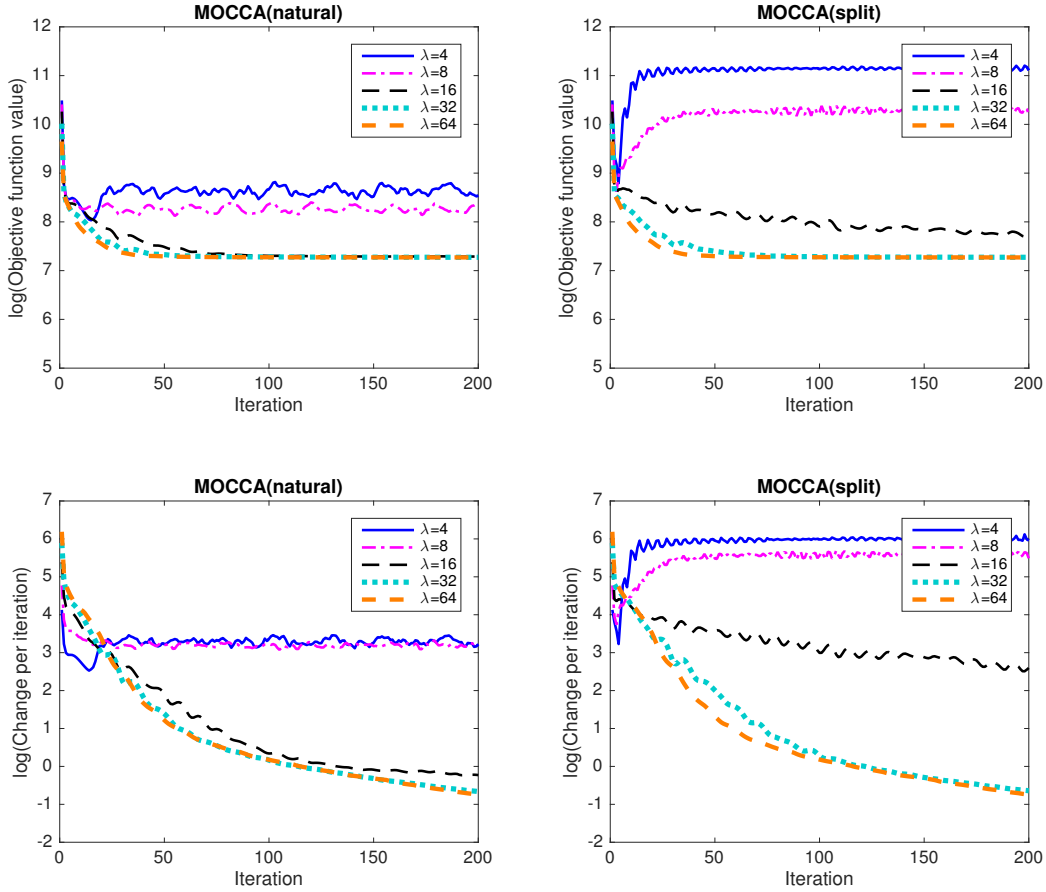


Figure 2: Results from Simulation 1. The top row plots the log of the objective function value defined in (26) against iteration number, while the bottom row plots the log of the change in the (x, w) variables (29) at each iteration, for several step size parameters λ . (Section 1.1.4 gives a direct correspondence between convergence of the (x, w) variables, and the optimality of the x variable.) The plots show results from the “natural” (left) and “split” (right) versions of the MOCCA algorithm, defined in (27) and (28), respectively.

(Recall from Section 1.1.4 that, if $\text{Change}_t \rightarrow 0$, then the optimality gap of the solution tends to zero; that is, the x variable is close to being a critical point for the optimization problem.) Looking first at the results for MOCCA(natural), we see that smaller λ values tend to lead to faster convergence at the very early stages, but poorer performance or instability at later stages. (In fact, this suggests the possibility of varying λ as we run more iterations, which we leave to future work.)

Turning to MOCCA(split), we see that the performance is worse at all λ values as compared with MOCCA(natural); the difference is minor for the largest λ values, but the lower λ values give far poorer results and far more instability for MOCCA(split) as compared to MOCCA(natural). This highlights the importance of the “mirroring” step in our algorithm,

which gives us the flexibility of placing the nonconvex terms into F , i.e. the function which will be optimized via the dual variable. In other scenarios, of course, a different arrangement of the terms may be preferable.

5.2 Simulation 2: Total Variation Penalty, with Errors in Variables

In the second simulation, we treat the errors-in-variables setting discussed in Section 2.1. We generate the signal x_{true} , the design matrix A , and the response vector b as in Simulation 1. Next, suppose our measurement of A is itself noisy: define $Z \in \mathbb{R}^{n \times d}$ with $Z_{ij} = A_{ij} + \text{Normal}(0, \sigma_A^2)$, where $\sigma_A = 0.2$. Finally, we would like to minimize the objective function

$$\frac{1}{2} x^\top \left(Z^\top Z - n \cdot \sigma_A^2 \mathbf{I}_d \right) x - x^\top Z^\top b + \nu \|x\|_{\text{TV}}, \quad (30)$$

with penalty parameter $\nu = 20$, where again we use two-dimensional total variation, $\|x\|_{\text{TV}} = \|\nabla_{2d} x\|_1$. (Here we use a different scaling of the likelihood term relative to Section 2.1 for simpler implementation and tuning.) Of course, due to the negative quadratic term, this objective function is strongly concave in some directions and so its global minimum is “at infinity”; within a bounded set, however, the penalty will ensure that the objective function is approximately convex. In practice, initializing our algorithm at $x = 0$ does not lead to any problems, and we converge to a bounded solution that can be viewed as a local minimum within a bounded set, e.g. $\{x : \|x\|_{\text{TV}} \leq R\}$ for some appropriate choice of R , as discussed in the context of restricted strong convexity in Section 4.2.1.

A proximal gradient descent method, as proposed by Loh and Wainwright (2013) for this type of nonconvex penalized likelihood, would in theory iterate the steps

$$\begin{cases} \tilde{x}_{t+1} = x_t - \frac{1}{\eta} \left((Z^\top Z - n\sigma_A^2 \mathbf{I}_d) x_t - (Z^\top b) \right), \\ x_{t+1} = \arg \min_x \left\{ \frac{1}{2} \|x - \tilde{x}_{t+1}\|_2^2 + \frac{\nu}{\eta} \|x\|_{\text{TV}} \right\}, \end{cases}$$

where $\frac{1}{\eta}$ is a step size parameter. However, the second step is a proximal operator for the total variation norm $\|x\|_{\text{TV}} = \|\nabla_{2d} x\|_1$, which cannot be calculated in closed form. Instead, we could apply the CP algorithm to the convex (sub)problem of this proximity operator with parameters $\Sigma = \lambda \cdot \frac{1}{2} \mathbf{I}_m$ and $\mathbf{T} = \lambda^{-1} \cdot \frac{1}{4} \mathbf{I}_d$, and could terminate this inner “prox loop” after some convergence criterion is reached, e.g. after some fixed number n_{step} of steps, or once the relative change in x is below ϵ_{thresh} . We do not show details of the derivation, but the complete procedure iterates these steps (taking extrapolation parameter $\theta = 1$):

$$\left\{ \begin{array}{l} \text{Gradient step: } \tilde{x}_{t+1} = x_t - \frac{1}{\eta} \left((Z^\top Z - n\sigma_A^2 \mathbf{I}_d) x_t - (Z^\top b) \right), \\ \text{Initialize prox loop: } x'_{t+1;0} = x_t, u'_{t+1;0} = u_t. \\ \text{Run prox loop: for } \ell = 1, 2, \dots, \text{ writing } \bar{x}'_{t+1;\ell} = 2x'_{t+1;\ell} - x'_{t+1;\ell-1}, \\ \left\{ \begin{array}{l} x'_{t+1;\ell} = \left(1 + \frac{1}{4\lambda}\right)^{-1} \left(x'_{t+1;\ell-1} + \frac{1}{4\lambda} \tilde{x}_{t+1} - \frac{1}{4\lambda} \nabla_{2d}^\top u'_{t+1;\ell-1} \right), \\ u'_{t+1;\ell} = \text{Truncate}_{[-\nu/\eta, \nu/\eta]} \left(u'_{t+1;\ell-1} + \frac{\lambda}{2} \nabla_{2d} \bar{x}'_{t+1;\ell} \right), \end{array} \right. \\ \text{until a convergence criterion is reached} \\ \left(\text{i.e. } \ell = n_{\text{step}} \text{ or } \left\| x'_{t+1;\ell} - x'_{t+1;\ell-1} \right\|_2 / \left\| x'_{t+1;\ell-1} \right\|_2 \leq \epsilon_{\text{thresh}} \right). \\ \text{Gather results from prox loop: } x_{t+1} = x'_{t+1;\ell}, u_{t+1} = u'_{t+1;\ell}. \end{array} \right. \quad (31)$$

We will refer to this method as the Approximate Proximal Gradient Descent (APGD) algorithm, where “approximate” describes the fact that the proximal operator step is only ever solved approximately via a finite number of steps in the inner loop.

In fact, if we examine this algorithm carefully, we can find that by taking a single step of the inner “prox loop” (that is, setting $n_{\text{step}} = 1$), we arrive back at the steps of the MOCCA algorithm. Specifically, as in the implementation (15) in Section 2.1, we choose $K = \nabla_{2d}$, $F(w) = F_v(w) = \nu \|w\|_1$, and

$$\mathbf{G}(x) = \frac{1}{2} x^\top \left(Z^\top Z - n\sigma_A^2 \mathbf{I}_d \right) x - x^\top (Z^\top b)$$

with local approximations given by linear expansions,

$$\begin{aligned} \mathbf{G}_z(x) &= \mathbf{G}(z) + \langle x - z, \nabla \mathbf{G}(x) \rangle \\ &= \langle x, (Z^\top Z - n\sigma_A^2 \mathbf{I}_d) z - Z^\top b \rangle + (\text{terms constant with respect to } x). \end{aligned}$$

The update steps of the MOCCA algorithm are then given by

$$\begin{cases} x_{t+1} = \arg \min_x \left\{ \langle \nabla_{2d} x, w_t \rangle + \mathbf{G}_{z_t}(x) + \frac{1}{2} \|x - x_t\|_{\mathbf{T}^{-1}}^2 \right\}, \\ w_{t+1} = \arg \min_y \left\{ -\langle \nabla_{2d} \bar{x}_{t+1}, w \rangle + \mathbf{F}^*(w) + \frac{1}{2} \|w - w_t\|_{\Sigma^{-1}}^2 \right\}, \\ z_{t+1} = x_{t+1}, \end{cases}$$

which we can simplify to

$$\begin{cases} x_{t+1} = x_t - \mathbf{T} (\nabla_{2d}^\top w_t + (Z^\top Z - n\sigma_A^2 \mathbf{I}_d) x_t - (Z^\top b)), \\ w_{t+1} = \text{Truncate}_{[-\nu, \nu]} (w_t + \Sigma \nabla_{2d} \bar{x}_{t+1}). \end{cases}$$

If we choose

$$\Sigma = \frac{\lambda \eta}{2} \cdot \mathbf{I}_{d-1} \quad \text{and} \quad \mathbf{T} = \frac{1}{(4\lambda + 1)\eta} \cdot \mathbf{I}_d,$$

it can be shown that this is equivalent to the proximal gradient algorithm (31) with a single inner loop step, i.e. with $n_{\text{step}} = 1$ (specifically, the iterates x_t stay the same, while the other variables are related as $w_t = \eta \cdot u_t$).

Now we compare the performance of the approximate proximal gradient descent (APGD) algorithm, with various stopping criteria for the inner “prox loop”, against the performance of the MOCCA algorithm, which we can view as the APGD algorithm taking exactly one step in each inner “prox loop”. For simplicity, we consider only a few values for the step size parameters, setting $\eta = \lambda = 100$ or $\eta = \lambda = 200$. As for Simulation 1, we will see that higher values for these parameters gives more stability at the cost of slower convergence.

We consider stopping rules for the inner loop as follows: either we run the inner loop for a fixed number of steps, $n_{\text{step}} \in \{1, 5\}$ (with $n_{\text{step}} = 1$ yielding MOCCA), or we use a convergence criterion $\epsilon_{\text{thresh}} \in \{0.1, 0.05, 0.01\}$. Figure 3 shows the results; for the figure on the left, we see that running the inner loop longer does help to make our solutions more accurate (i.e. the objective function is lower) over the range of iterations. However, each iteration has greater computational cost when we increase the time spent running the inner loop. Since we would like to see the performance as a function of computational cost, the

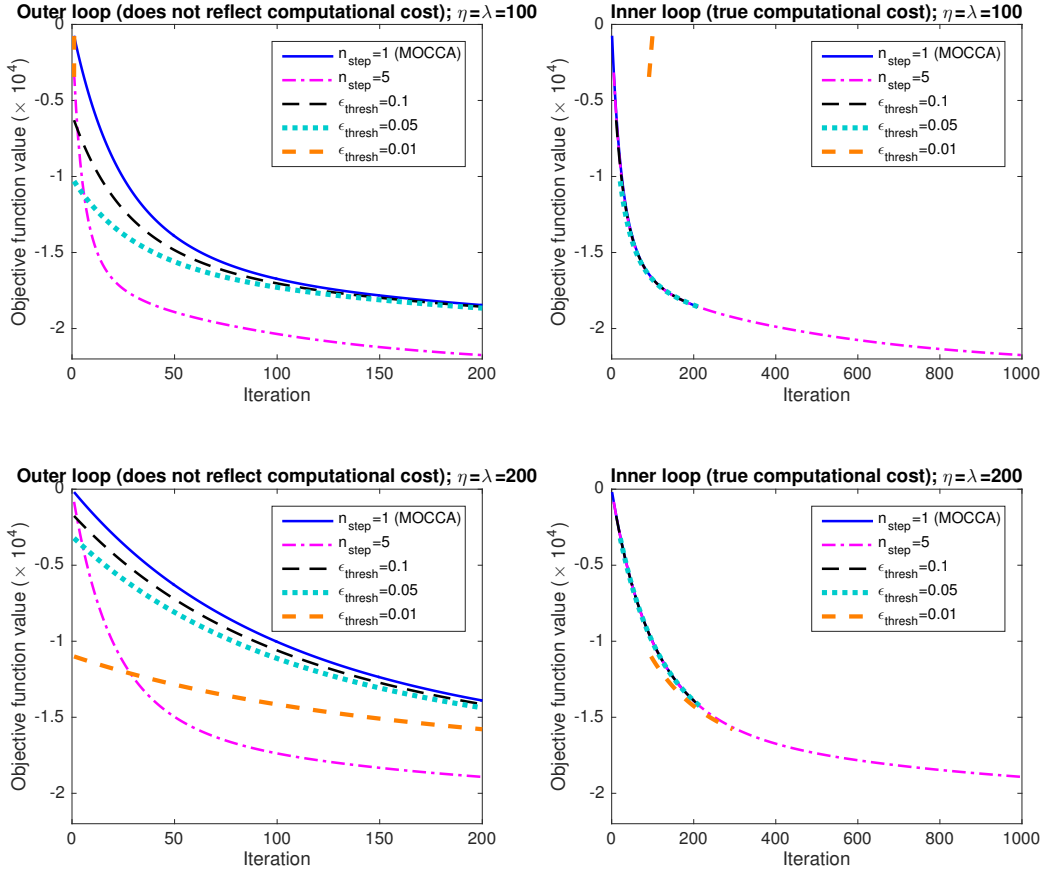


Figure 3: Results from Simulation 2, plotting the value of the objective function (30) against iteration number, counting either iterations of the “outer loop” (left) or of the “inner loop” (right), for various stopping rules for the inner loop in (31). Parameters are set as $\eta = \lambda = 100$ (top) or $\eta = \lambda = 200$ (bottom). Counting the number of passes through the inner loop is an accurate reflection of the true computational cost of (31), and so the right-hand plots give the correct interpretation of the results, where the various versions of the algorithms perform nearly identically in both settings (the lines are indistinguishable) except for the setting $\epsilon_{\text{thresh}} = 0.01$ which diverges for the setting $\eta = \lambda = 100$. Setting $n_{\text{step}} = 1$ yields the MOCCA algorithm as discussed in the text.

plot on the right-hand side of Figure 3 shows the same results plotted against the true number of iterations, i.e. where we count each pass through the inner loop of (31) rather than counting only passes through the outer loop of (31). In this setting, we see that in fact the various versions of the algorithms perform nearly identically—in other words, a one-step approximation to the proximal map performs just as well as a more conservative inner loop that is run for longer—with the exception of setting $\epsilon_{\text{thresh}} = 0.01$ and $\eta = \lambda = 100$, in which case the algorithm diverges immediately. It is interesting to note that this small

choice for ϵ_{thresh} is closest in spirit to proximal gradient descent (that is, the proximal step is the most accurate), although of course there may be some effect of tuning parameters. The choice ϵ_{thresh} does achieve convergence with the more conservative choice $\eta = \lambda = 200$, but convergence is noticeably slower in this case. Thus, we can conclude that when the proximal step does not have a closed form solution, it may be better to use a coarse approximation (which, implicitly, is the strategy taken by MOCCA for this problem) rather than aiming for near-convergence in the proximal step for each iteration.

6. Proofs

6.1 Critical Points (Theorem 1)

For this proof we will use two facts: for a continuous convex function h ,

$$\text{If } a_t \in \partial h(b_t) \text{ and } a_t \rightarrow a, b_t \rightarrow b \text{ then } a \in \partial h(b), \quad (32)$$

and

$$a \in \partial h(b) \text{ if and only if } b \in \partial h^*(a). \quad (33)$$

First, by definition of the $w_{t+1;1}$ update step,

$$\partial \mathbf{F}_{v_t}^*(w_{t+1;1}) \ni \Sigma^{-1}(w_{t+1;0} - w_{t+1;1}) + K\bar{x}_{t+1;1},$$

and therefore by (33),

$$w_{t+1;1} \in \partial \mathbf{F}_{v_t}(\Sigma^{-1}(w_{t+1;0} - w_{t+1;1}) + K\bar{x}_{t+1;1}).$$

We can rewrite this as

$$w_{t+1;1} \in \underbrace{\partial \mathbf{F}_{\hat{v}}(\Sigma^{-1}(w_{t+1;0} - w_{t+1;1}) + K\bar{x}_{t+1;1})}_{(\text{Term 1})} + \underbrace{\nabla(\mathbf{F}_{v_t} - \mathbf{F}_{\hat{v}})(\Sigma^{-1}(w_{t+1;0} - w_{t+1;1}) + K\bar{x}_{t+1;1})}_{(\text{Term 2})}. \quad (34)$$

Next, since the solution converges, we see that

$$w_{t+1;1} \rightarrow \hat{w}, v_t \rightarrow \hat{v}, \Sigma^{-1}(w_{t+1;0} - w_{t+1;1}) + K\bar{x}_{t+1;1} \rightarrow K\hat{x}.$$

Our assumption (23) implies that

$$(v, w) \mapsto \nabla(\mathbf{F}_v - \mathbf{F}_{\hat{v}})(w) = -\nabla(\mathbf{F} - \mathbf{F}_v)(w) + \nabla(\mathbf{F} - \mathbf{F}_{\hat{v}})(w)$$

is jointly continuous in (v, w) , and so

$$(\text{Term 2}) = \nabla(\mathbf{F}_{v_t} - \mathbf{F}_{\hat{v}})(\Sigma^{-1}(w_{t+1;0} - w_{t+1;1}) + K\bar{x}_{t+1;1}) \rightarrow \nabla(\mathbf{F}_{\hat{v}} - \mathbf{F}_{\hat{v}})(K\hat{x}) = 0.$$

Therefore, applying the property (32) to the expression in (34), we see that

$$\hat{w} \in \partial \mathbf{F}_{\hat{v}}(K\hat{x}). \quad (35)$$

Next, for each t , by definition,

$$\begin{aligned} v_{t+1} &= \frac{1}{L_{t+1}} \sum_{\ell=1}^{L_{t+1}} (\Sigma^{-1}(w_{t+1;\ell-1} - w_{t+1;\ell}) + K\bar{x}_{t+1;\ell}) \\ &= \frac{1}{L_{t+1}} (\Sigma^{-1}(w_{t+1;0} - w_{t+1;L_{t+1}}) + K(x_{t+1;L_{t+1}} - x_{t+1;0})) + Kx_{t+1} . \end{aligned}$$

Taking limits on each side, we see that $\hat{v} = K\hat{x}$. Returning to (35) above this proves that

$$\hat{w} \in \partial\mathbf{F}_{K\hat{x}}(K\hat{x}) .$$

Next, by definition of the $x_{t+1;1}$ update step,

$$\partial\mathbf{G}_{z_t}(x_{t+1;1}) = \mathbf{T}^{-1}(x_{t+1;0} - x_{t+1;1}) - K^\top w_{t+1;0} .$$

Taking limits on each side as $t \rightarrow \infty$, and applying (32) as before,

$$\partial\mathbf{G}_{\hat{z}}(\hat{x}) \ni -K^\top \hat{w} .$$

And, we know that $z_{t+1} = x_{t+1}$ for each t , therefore $\hat{z} = \hat{x}$, and so

$$\partial\mathbf{G}_{\hat{x}}(\hat{x}) \ni -K^\top \hat{w} .$$

Combining the work above, then,

$$0 = K^\top \hat{w} - K^\top \hat{w} \in K^\top \partial\mathbf{F}_{K\hat{x}}(K\hat{x}) + \partial\mathbf{G}_{\hat{x}}(\hat{x}) ,$$

as desired.

6.2 Convergence Guarantee (Theorem 2)

We first introduce some notation and supporting lemmas before turning to the main proof.

6.2.1 NOTATION

Fixing any expansion points $(z, v) \in \mathbb{R}^d \times \mathbb{R}^m$, we define a primal-dual update step:

$$(x', w') = \text{Step}_{z,v}(x, w) ,$$

given by

$$\begin{cases} x' = \arg \min_{x''} \left\{ \langle Kx'', w \rangle + \mathbf{G}_z(x'') + \frac{1}{2} \|x'' - x\|_{\mathbf{T}^{-1}}^2 \right\} , \\ w' = \arg \min_{w''} \left\{ -\langle K(2x' - x), w'' \rangle + \mathbf{F}_v^*(w'') + \frac{1}{2} \|w'' - w\|_{\Sigma^{-1}}^2 \right\} . \end{cases}$$

This is one step of the CP algorithm applied to the convex objective function

$$\min_x \{ \mathbf{F}_v(Kx) + \mathbf{G}_z(x) \}$$

with extrapolation parameter $\theta = 1$.

Next, defining x^* to be any critical point of the original problem as before, let $w^* \in \partial F_{Kx^*}(Kx^*)$ be an element of the subdifferential such that

$$0 \in K^\top w^* + \partial G_{x^*}(x^*) .$$

Finally, for any expansion points (z, v) , define the (not necessarily unique) solution for the convex optimization problem when we use approximations F_v, G_z :

$$x_{z,v}^* = \arg \min_x \{F_v(Kx) + G_z(x)\} ,$$

and let $w_{z,v}^* \in \partial F_v(Kx_{z,v}^*)$ be an element of the subdifferential such that

$$0 \in K^\top w^* + \partial G_z(x_{z,v}^*) .$$

6.2.2 LEMMAS

The proof of Theorem 2 can be split into several key results. First we state these lemmas and explain their role, then we will formally prove the theorem. The lemmas are proved in Appendix A.

The first lemma shows that, if we use expansion points (z, v) close to the true solution, i.e. $(z, v) \approx (x^*, Kx^*)$, then the minimizer $x_{z,v}^*$ for the convex approximation will be close to x^* .

Lemma 3 *Suppose that Assumptions 1 and 2 hold. Define*

$$\Theta = \begin{pmatrix} \Theta_G & 0 \\ 0 & \Theta_F \end{pmatrix} + \frac{C_{\text{cvx}}}{(C_{\text{matrix}})^2} \mathbf{I} \succ 0 .$$

Then there exist constants $C_{\text{contr}} > 0, C_{\text{excess}} < \infty$, which depend only on $C_{\text{matrix}}, C_{\text{cvx}}, C_{\text{Lip}}, C_{\text{grad}}$, such that for any $(z, v) \in \text{dom}(G) \times \text{dom}(F)$,

$$\left\| \begin{pmatrix} x_{z,v}^* - x^* \\ Kx_{z,v}^* - Kx^* \end{pmatrix} \right\|_{\Theta} \leq (1 - C_{\text{contr}}) \left\| \begin{pmatrix} z - x^* \\ v - Kx^* \end{pmatrix} \right\|_{\Theta} + C_{\text{excess}} \cdot \tau R$$

and

$$\left\| \begin{pmatrix} x_{z,v}^* - x^* \\ w_{z,v}^* - w^* \end{pmatrix} \right\|_2 \leq C_{\text{Lip}} + C_{\text{excess}} \left(\left\| \begin{pmatrix} z - x^* \\ v - Kx^* \end{pmatrix} \right\|_{\Theta} + \tau R \right) .$$

The second lemma shows that, after running an ‘‘inner loop’’, the (x, w) variables are nearly optimal for the current convex approximation, and the next expansion points are also near this optimum.

Lemma 4 *Suppose that Assumptions 1 and 2 hold. For any $L \geq 1$, and any points $(x^{(0)}, w^{(0)}), (z, v) \in \text{dom}(G) \times \text{dom}(F)$, suppose we iterate $\text{Step}_{z,v}(\cdot)$ for L times,*

$$\begin{aligned} (x^{(1)}, w^{(1)}) &= \text{Step}_{z,v}(x^{(0)}, w^{(0)}), \quad (x^{(2)}, w^{(2)}) = \text{Step}_{z,v}(x^{(1)}, w^{(1)}), \dots, \\ (x^{(L)}, w^{(L)}) &= \text{Step}_{z,v}(x^{(L-1)}, w^{(L-1)}) , \end{aligned}$$

and then define the averages $(\tilde{x}, \tilde{w}) = \frac{1}{L} \sum_{\ell=1}^L (x^{(\ell)}, w^{(\ell)})$ and averaged expansion points (\tilde{z}, \tilde{v}) as

$$\tilde{z} = \frac{1}{L} \sum_{\ell=1}^L x^{(\ell)} \quad \text{and} \quad \tilde{v} = \frac{1}{L} \sum_{\ell=1}^L \left(\Sigma^{-1}(w^{(\ell-1)} - w^{(\ell)}) + K(2x^{(\ell)} - x^{(\ell-1)}) \right).$$

Then there exists a constant $C_{\text{iter}} < \infty$, depending only on $C_{\text{matrix}}, C_{\text{cvx}}, C_{\text{Lip}}, C_{\text{grad}}$ (and in particular, not dependent on L), such that

$$\left\| \begin{pmatrix} \tilde{z} - x_{z,v}^* \\ \tilde{v} - Kx_{z,v}^* \end{pmatrix} \right\|_{\Theta} \leq C_{\text{iter}} \left(\frac{1}{\sqrt{L}} \left\| \begin{pmatrix} x^{(0)} - x_{z,v}^* \\ w^{(0)} - w_{z,v}^* \end{pmatrix} \right\|_2 + \tau R \right).$$

and

$$\left\| \begin{pmatrix} \tilde{x} - x_{z,v}^* \\ \tilde{w} - w_{z,v}^* \end{pmatrix} \right\|_2 \leq C_{\text{Lip}} + C_{\text{iter}} \left(\frac{1}{\sqrt{L}} \left\| \begin{pmatrix} x^{(0)} - x_{z,v}^* \\ w^{(0)} - w_{z,v}^* \end{pmatrix} \right\|_2 + \tau R \right)$$

6.2.3 PROOF OF THEOREM 2

We will assume for simplicity that the expansion point is initialized with some z_0 satisfying $\|z_0\|_{\text{restrict}} \leq R$; if this is not the case at step $t = 0$, then our results can be easily adjusted since we will have $z_1 = x_1 \in \text{dom}(\mathbf{G}_{z_0})$ and therefore $\|z_1\|_{\text{restrict}} \leq R$, so we can simply shift our calculations by one time point.

First, choose any δ such that $0 < \delta < (1 - C_{\text{contr}})^{-2} - 1$. Define constants

$$\begin{aligned} C_1 &= \max \left\{ \left\| \begin{pmatrix} x_1 - x_{z_1, v_1}^* \\ w_1 - w_{z_1, v_1}^* \end{pmatrix} \right\|_2, \right. \\ &\quad \left. 6C_{\text{Lip}} + 4C_{\text{excess}} + 2 \left(C_{\text{iter}} + 2C_{\text{excess}} \left(\frac{C_{\text{iter}} + C_{\text{excess}}}{C_{\text{contr}}} + 1 \right) \right) \tau R \right\}, \\ C_2 &= \max \left\{ \sqrt{L'_1} \cdot \left\| \begin{pmatrix} z_1 - x^* \\ v_1 - Kx^* \end{pmatrix} \right\|_{\Theta}, \frac{C_{\text{iter}} C_1}{1 - (1 - C_{\text{contr}}) \sqrt{1 + \delta}} \right\}, \\ C_3 &= \frac{C_{\text{iter}} + C_{\text{excess}}}{C_{\text{contr}}}, \end{aligned}$$

and define

$$L_{\min} = \max\{4(C_{\text{iter}})^2, (C_2)^2\}.$$

To prove the desired result, we will prove that

$$\left\| \begin{pmatrix} x_{t+1} - x_{z_{t+1}, v_{t+1}}^* \\ w_{t+1} - w_{z_{t+1}, v_{t+1}}^* \end{pmatrix} \right\|_2 \leq C_1 \tag{36}$$

and

$$\left\| \begin{pmatrix} z_{t+1} - x^* \\ v_{t+1} - Kx^* \end{pmatrix} \right\|_{\Theta} \leq \frac{C_2}{\sqrt{L'_{t+1}}} + C_3 \tau R \tag{37}$$

for all $t \geq 0$.

Assuming that these bounds hold, we then have

$$\|x_{t+1} - x^*\|_2 = \|z_{t+1} - x^*\|_2 \leq \|\Theta^{-1}\| \cdot \left(\frac{C_2}{\sqrt{L'_{t+1}}} + C_3\tau R \right),$$

where the first step holds by definition of z_{t+1} ; this proves the desired theorem with $C_{\text{converge}} := \|\Theta^{-1}\| \cdot \max\{C_2, C_3\}$.

Now we prove (36) and (37) by induction. For $t = 0$, both statements are true trivially by our definitions of C_1 and C_2 . Now we will assume that the statements are true for all $t = 0, \dots, m-1$ and will prove that they hold for $t = m$. First, for (37), we have

$$\begin{aligned} & \left\| \begin{pmatrix} z_{m+1} - x^* \\ v_{m+1} - Kx^* \end{pmatrix} \right\|_{\Theta} \\ & \leq \left\| \begin{pmatrix} z_{m+1} - x_{z_m, v_m}^* \\ v_{m+1} - Kx_{z_m, v_m}^* \end{pmatrix} \right\|_{\Theta} + \left\| \begin{pmatrix} x_{z_m, v_m}^* - x^* \\ Kx_{z_m, v_m}^* - Kx^* \end{pmatrix} \right\|_{\Theta} \quad \text{by the triangle inequality} \\ & \leq \left\| \begin{pmatrix} z_{m+1} - x_{z_m, v_m}^* \\ v_{m+1} - Kx_{z_m, v_m}^* \end{pmatrix} \right\|_{\Theta} \\ & \quad + (1 - C_{\text{contr}}) \cdot \left\| \begin{pmatrix} z_m - x^* \\ v_m - Kx^* \end{pmatrix} \right\|_{\Theta} + C_{\text{excess}}\tau R \quad \text{by Lemma 3} \\ & \leq C_{\text{iter}} \left(\frac{1}{\sqrt{L_{m+1}}} \left\| \begin{pmatrix} x_m - x_{z_m, v_m}^* \\ w_m - w_{z_m, v_m}^* \end{pmatrix} \right\|_2 + \tau R \right) \\ & \quad + (1 - C_{\text{contr}}) \cdot \left\| \begin{pmatrix} z_m - x^* \\ v_m - Kx^* \end{pmatrix} \right\|_{\Theta} + C_{\text{excess}}\tau R \quad \text{by Lemma 4} \\ & \leq C_{\text{iter}} \left(\frac{1}{\sqrt{L_{m+1}}} C_1 + \tau R \right) + (1 - C_{\text{contr}}) \cdot \left(\frac{C_2}{\sqrt{L'_m}} + C_3\tau R \right) \\ & \quad + C_{\text{excess}}\tau R \quad \text{by (36) and (37) applied with } t = m-1 \\ & \leq C_{\text{iter}} \left(\frac{1}{\sqrt{L'_{m+1}}} C_1 + \tau R \right) + (1 - C_{\text{contr}}) \cdot \left(\frac{C_2\sqrt{1+\delta}}{\sqrt{L'_{m+1}}} + C_3\tau R \right) \\ & \quad + C_{\text{excess}}\tau R \quad \text{since } L'_{m+1} \leq L_{m+1}, (1+\delta)L'_m \\ & \leq \frac{C_2}{\sqrt{L'_{m+1}}} + C_3\tau R, \end{aligned}$$

where the last step holds by our definition of the constants C_2, C_3 . This concludes the proof of (37) for $t = m$.

Next, we turn to the proof of (36). For both $t = m - 1$ and $t = m$,

$$\begin{aligned}
 \left\| \begin{pmatrix} x_{z_{t+1}, v_{t+1}}^* - x^* \\ w_{z_{t+1}, v_{t+1}}^* - w^* \end{pmatrix} \right\|_2 &\leq C_{\text{Lip}} + C_{\text{excess}} \left(\left\| \begin{pmatrix} z_{t+1} - x^* \\ v_{t+1} - Kx^* \end{pmatrix} \right\|_{\Theta} + \tau R \right) \quad \text{by Lemma 3} \\
 &= C_{\text{Lip}} + C_{\text{excess}} \left(\frac{C_2}{\sqrt{L'_{t+1}}} + (C_3 + 1)\tau R \right) \quad \text{by (37) at step } t \\
 &\leq C_{\text{Lip}} + C_{\text{excess}} \left(\frac{C_2}{\sqrt{L_{\min}}} + (C_3 + 1)\tau R \right) \quad \text{since } L'_{t+1} \geq L_{\min},
 \end{aligned}$$

and also, we have

$$\begin{aligned}
 &\left\| \begin{pmatrix} x_{m+1} - x_{z_m, v_m}^* \\ w_{m+1} - w_{z_m, v_m}^* \end{pmatrix} \right\|_2 \\
 &\leq C_{\text{Lip}} + C_{\text{iter}} \left(\frac{1}{\sqrt{L_{m+1}}} \left\| \begin{pmatrix} x_m - x_{z_m, v_m}^* \\ w_m - w_{z_m, v_m}^* \end{pmatrix} \right\|_2 + \tau R \right) \quad \text{by Lemma 4} \\
 &\leq C_{\text{Lip}} + C_{\text{iter}} \left(\frac{1}{\sqrt{L_{m+1}}} C_1 + \tau R \right) \quad \text{by (36) applied with } t = m - 1 \\
 &\leq C_{\text{Lip}} + C_{\text{iter}} \left(\frac{1}{\sqrt{L_{\min}}} C_1 + \tau R \right) \quad \text{since } L_{m+1} \geq L_{\min}.
 \end{aligned}$$

Therefore, combining these calculations,

$$\begin{aligned}
 &\left\| \begin{pmatrix} x_{m+1} - x_{z_{m+1}, v_{m+1}}^* \\ w_{m+1} - w_{z_{m+1}, v_{m+1}}^* \end{pmatrix} \right\|_2 \\
 &\leq \left\| \begin{pmatrix} x_{m+1} - x_{z_m, v_m}^* \\ w_{m+1} - w_{z_m, v_m}^* \end{pmatrix} \right\|_2 + \left\| \begin{pmatrix} x_{z_m, v_m}^* - x^* \\ w_{z_m, v_m}^* - w^* \end{pmatrix} \right\|_2 + \left\| \begin{pmatrix} x_{z_{m+1}, v_{m+1}}^* - x^* \\ w_{z_{m+1}, v_{m+1}}^* - w^* \end{pmatrix} \right\|_2 \\
 &\leq 3C_{\text{Lip}} + C_{\text{iter}} \left(\frac{1}{\sqrt{L_{\min}}} C_1 + \tau R \right) + 2C_{\text{excess}} \left(\frac{C_2}{\sqrt{L_{\min}}} + (C_3 + 1)\tau R \right) \\
 &\leq 3C_{\text{Lip}} + C_{\text{iter}} \left(\frac{1}{\sqrt{4(C_{\text{iter}})^2}} C_1 + \tau R \right) \\
 &\quad + 2C_{\text{excess}} \left(\frac{C_2}{\sqrt{(C_2)^2}} + \left(\frac{C_{\text{iter}} + C_{\text{excess}}}{C_{\text{contr}}} + 1 \right) \tau R \right) \quad \text{by definition of } C_3 \text{ and of } L_{\min} \\
 &\leq C_1,
 \end{aligned}$$

by definition of C_1 . This proves the desired bound (36) for $t = m$, and thus we have proved the theorem.

7. Discussion

We have developed a primal/dual algorithm for minimizing composite objective functions of the form $F(Kx) + G(x)$, which is able to handle nondifferentiability and nonconvexity (even strong concavity) in each individual term, beyond what is possible with many existing

approaches based on alternating minimization or proximal gradient methods. The key step of the MOCCA algorithm is the careful choice of local convex approximations to F and G at each step, which respects the mirroring between the primal and dual variables of the algorithm. Our method allows for accurate and efficient optimization for a range of problems arising in high-dimensional statistics, such as nonconvex total variation penalties (which reduce the bias caused by shrinkage, when compared to using a convex total variation norm), as well as inverse problems in computed tomography (CT) imaging.

Our present theoretical results give a convergence guarantee, in the case that the overall objective function is approximately convex, for a more stable form of the MOCCA algorithm. In future work, we hope to better understand the relative performance of the various forms of the algorithm, and to find a tighter characterization of the convergence behavior of the algorithm. It would also be interesting to consider a more general form of objective function, $F(w) + G(x)$ where F, G are nonconvex and nondifferentiable, and where instead of the linear constraint $w = Kx$, the variables w and x are linked via a nonlinear map; such an extension would greatly increase the range of applications of the method.

Acknowledgments

This work was partially supported by NIH grants CA158446, CA182264, and EB018102. The contents of this article are solely the responsibility of the authors and do not necessarily represent the official views of the National Institutes of Health. The authors are grateful to Wooseok Ha for discussions on the connections with existing algorithms, and to collaborators Taly Gilat Schmidt and Xiaochuan Pan for discussions on the application to CT imaging.

Appendix A. Proofs of Lemmas

A.1 Proof of Lemma 3

By definition of $x_{z,v}^*$,

$$0 \in K^\top \partial F_v(Kx_{z,v}^*) + \partial G_z(x_{z,v}^*) \quad (38)$$

and since x^* is a critical point of the original objective function,

$$0 \in K^\top \partial F_{Kx^*}(Kx^*) + \partial G_{x^*}(x^*) . \quad (39)$$

Since $(F_v - F_{Kx^*}) = (F - F_{Kx^*}) - (F - F_v)$ and $(G_z - G_{x^*}) = (G - G_{x^*}) - (G - G_z)$ are differentiable, we can rewrite (39) as

$$\begin{aligned} 0 \in K^\top \partial F_v(Kx^*) + K^\top \nabla(F - F_v)(Kx^*) - K^\top \nabla(F - F_{Kx^*})(Kx^*) \\ + \partial G_z(x^*) + \nabla(G - G_z)(x^*) - \nabla(G - G_{x^*})(x^*). \end{aligned}$$

By the first-order conditions (4), we know that $\nabla(F - F_{Kx^*})(Kx^*) = 0$ and $\nabla(G - G_{x^*})(x^*) = 0$, so this reduces to

$$0 \in K^\top \partial F_v(Kx^*) + K^\top \nabla(F - F_v)(Kx^*) + \partial G_z(x^*) + \nabla(G - G_z)(x^*). \quad (40)$$

We also see that $\|x^*\|_{\text{restrict}}, \|x_{z,v}^*\|_{\text{restrict}} \leq R$, since $x^*, x_{z,v}^*$ must lie in $\text{dom}(G) = \text{dom}(G_z)$.

Then we have

$$\begin{aligned}
 & \|x_{z,v}^* - x^*\|_{\Theta_G}^2 + \|Kx_{z,v}^* - Kx^*\|_{\Theta_F}^2 + C_{\text{cvx}} \|x_{z,v}^* - x^*\|_2^2 - 2\tau^2 \cdot \|x_{z,v}^* - x^*\|_{\text{restrict}}^2 \\
 & \leq \|Kx_{z,v}^* - Kx^*\|_{\Lambda_F}^2 + \|x_{z,v}^* - x^*\|_{\Lambda_G}^2 - \tau^2 \cdot \|x_{z,v}^* - x^*\|_{\text{restrict}}^2 \quad \text{by Assumption 2} \\
 & \leq \langle Kx_{z,v}^* - Kx^*, \partial F_v(Kx_{z,v}^*) - \partial F_v(Kx^*) \rangle \\
 & \quad + \langle x_{z,v}^* - x^*, \partial G_z(x_{z,v}^*) - \partial G_z(x^*) \rangle \quad \text{by Assumption 2} \\
 & = \langle Kx_{z,v}^* - Kx^*, \nabla(F - F_v)(Kx^*) \rangle + \langle x_{z,v}^* - x^*, \nabla(G - G_z)(x^*) \rangle \quad \text{by (38) and (40)} \\
 & \leq \frac{1}{2} \|Kx_{z,v}^* - Kx^*\|_{\Theta_F}^2 + \frac{1}{2} \|v - Kx^*\|_{\Theta_F}^2 + \frac{1}{2} \|x_{z,v}^* - x^*\|_{\Theta_G}^2 + \frac{1}{2} \|z - x^*\|_{\Theta_G}^2 \\
 & \quad + \frac{\tau^2}{2} \left(\|x_{z,v}^* - x^*\|_{\text{restrict}}^2 + \|z - x^*\|_{\text{restrict}}^2 \right) \quad \text{by Assumption 2.}
 \end{aligned}$$

Next, recall that $\|x_{z,v}^*\|_{\text{restrict}}, \|x^*\|_{\text{restrict}} \leq R$ from before and $\|z\|_{\text{restrict}} \leq R$ by assumption. After rearranging terms and multiplying by 2, this gives

$$\begin{aligned}
 & \|x_{z,v}^* - x^*\|_{\Theta_G}^2 + \|Kx_{z,v}^* - Kx^*\|_{\Theta_F}^2 + 2C_{\text{cvx}} \|x_{z,v}^* - x^*\|_2^2 \\
 & \leq \|z - x^*\|_{\Theta_G}^2 + \|v - Kx^*\|_{\Theta_F}^2 + 24\tau^2 R^2. \quad (41)
 \end{aligned}$$

Now, using the definition of Θ , we see that

$$\begin{pmatrix} \Theta_G & 0 \\ 0 & \Theta_F \end{pmatrix} \preceq (1 - C_{\text{contr}})^2 \cdot \Theta$$

where

$$C_{\text{contr}} := \frac{\frac{C_{\text{cvx}}}{(C_{\text{matrix}})^3}}{2 \left(1 + \frac{C_{\text{cvx}}}{(C_{\text{matrix}})^3} \right)} > 0.$$

We then get

$$\begin{aligned}
 & (1 - C_{\text{contr}})^2 \left\| \begin{pmatrix} z - x^* \\ v - Kx^* \end{pmatrix} \right\|_{\Theta}^2 + 24\tau^2 R^2 \\
 & \geq \|z - x^*\|_{\Theta_G}^2 + \|v - Kx^*\|_{\Theta_F}^2 + 24\tau^2 R^2 \\
 & \geq \|x_{z,v}^* - x^*\|_{\Theta_G}^2 + \|Kx_{z,v}^* - Kx^*\|_{\Theta_F}^2 + 2C_{\text{cvx}} \|x_{z,v}^* - x^*\|_2^2 \quad \text{from (41)} \\
 & \geq \|x_{z,v}^* - x^*\|_{\Theta_G}^2 + \|Kx_{z,v}^* - Kx^*\|_{\Theta_F}^2 \\
 & \quad + C_{\text{cvx}} \|x_{z,v}^* - x^*\|_2^2 + \frac{C_{\text{cvx}}}{(C_{\text{matrix}})^2} \|Kx_{z,v}^* - Kx^*\|_2^2 \quad \text{since } \|K\| \leq C_{\text{matrix}} \\
 & \geq \|x_{z,v}^* - x^*\|_{\Theta_G + C_{\text{cvx}}/(C_{\text{matrix}})^2 \cdot \mathbf{I}}^2 + \|Kx_{z,v}^* - Kx^*\|_{\Theta_F + C_{\text{cvx}}/(C_{\text{matrix}})^2 \cdot \mathbf{I}}^2 \quad \text{since } C_{\text{matrix}} \geq 1 \\
 & = \left\| \begin{pmatrix} x_{z,v}^* - x^* \\ Kx_{z,v}^* - Kx^* \end{pmatrix} \right\|_{\Theta}^2.
 \end{aligned}$$

Using the fact that $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$, we obtain

$$\left\| \begin{pmatrix} x_{z,v}^* - x^* \\ Kx_{z,v}^* - Kx^* \end{pmatrix} \right\|_{\Theta} \leq (1 - C_{\text{contr}}) \left\| \begin{pmatrix} z - x^* \\ v - Kx^* \end{pmatrix} \right\|_{\Theta} + \sqrt{24}\tau R. \quad (42)$$

Next, by definition, we also have

$$w_{z,v}^* \in \partial F_v(Kx_{z,v}^*) \text{ and } w^* \in \partial F_{Kx^*}(Kx^*) .$$

This second expression can be rewritten as

$$w^* \in \partial F_v(Kx^*) + \nabla(F - F_v)(Kx^*) - \nabla(F - F_{Kx^*})(Kx^*) = \partial F_v(Kx^*) + \nabla(F - F_v)(Kx^*) ,$$

where the last step uses the first-order condition (4) to see that $\nabla(F - F_{Kx^*})(Kx^*) = 0$. Therefore,

$$\begin{aligned} \|w_{z,v}^* - w^*\|_2 &\leq \|\partial F_v(Kx^*) - \partial F_v(Kx_{z,v}^*)\|_2 + \|\nabla(F - F_v)(Kx^*)\|_2 \\ &\leq C_{\text{Lip}} + C_{\text{grad}} \|Kx_{z,v}^* - Kx^*\|_2 + \sqrt{\|\Theta_{\text{F}}\|} \|v - Kx^*\|_{\Theta_{\text{F}}} \quad \text{by Assumption 2} \\ &\leq C_{\text{Lip}} + C_{\text{grad}} C_{\text{matrix}} \|x_{z,v}^* - x^*\|_2 + \sqrt{C_{\text{matrix}}} \|v - Kx^*\|_{\Theta_{\text{F}}} , \end{aligned}$$

and so

$$\begin{aligned} \left\| \begin{pmatrix} x_{z,v}^* - x^* \\ w_{z,v}^* - w^* \end{pmatrix} \right\|_2 &\leq C_{\text{Lip}} + (1 + C_{\text{grad}} C_{\text{matrix}}) \|x_{z,v}^* - x^*\|_2 + \sqrt{C_{\text{matrix}}} \|v - Kx^*\|_{\Theta_{\text{F}}} \\ &\leq C_{\text{Lip}} + \frac{1 + C_{\text{grad}} C_{\text{matrix}}}{\sqrt{C_{\text{cvx}}/C_{\text{matrix}}}} \left\| \begin{pmatrix} x_{z,v}^* - x^* \\ Kx_{z,v}^* - Kx^* \end{pmatrix} \right\|_{\Theta} \\ &\quad + \sqrt{C_{\text{matrix}}} \|v - Kx^*\|_{\Theta_{\text{F}}} \quad \text{since } \Theta \succeq \frac{C_{\text{cvx}}}{(C_{\text{matrix}})^2} \mathbf{I} \\ &\leq C_{\text{Lip}} + \frac{1 + C_{\text{grad}} C_{\text{matrix}}}{\sqrt{C_{\text{cvx}}/C_{\text{matrix}}}} \left((1 - C_{\text{contr}}) \left\| \begin{pmatrix} z - x^* \\ v - Kx^* \end{pmatrix} \right\|_{\Theta} + \sqrt{24}\tau R \right) \\ &\quad + \sqrt{C_{\text{matrix}}} \left\| \begin{pmatrix} z - x^* \\ v - Kx^* \end{pmatrix} \right\|_{\Theta} , \end{aligned}$$

where the last step applies (42) from above. Setting

$$C_{\text{excess}} = \max \left\{ \sqrt{24}, \sqrt{C_{\text{matrix}}} + (1 - C_{\text{contr}}) \cdot \frac{1 + C_{\text{grad}} C_{\text{matrix}}}{\sqrt{C_{\text{cvx}}/C_{\text{matrix}}}}, \sqrt{24} \cdot \frac{1 + C_{\text{grad}} C_{\text{matrix}}}{\sqrt{C_{\text{cvx}}/C_{\text{matrix}}}} \right\} ,$$

this is sufficient to prove the lemma.

A.2 Proof of Lemma 4

First we state and prove a supporting lemma which considers only a single step of the ‘‘inner loop’’.

Lemma 5 *There exists a constant $C_{\text{monotone}} > 0$, which depends only on C_{matrix} , C_{cvx} , C_{Lip} , C_{grad} , such that, for any x, y, z with $\|x\|_{\text{restrict}} \leq R$, if $(x', w') = \text{Step}_{z,v}(x, w)$, then*

$$\begin{aligned} C_{\text{monotone}} \left(\|x' - x_{z,v}^*\|_2^2 + \left\| \begin{pmatrix} x - x' \\ w - w' \end{pmatrix} \right\|_2^2 \right) \\ \leq \left\| \begin{pmatrix} x - x_{z,v}^* \\ w - w_{z,v}^* \end{pmatrix} \right\|_M^2 - \left\| \begin{pmatrix} x' - x_{z,v}^* \\ w' - w_{z,v}^* \end{pmatrix} \right\|_M^2 + \tau^2 R^2 . \end{aligned}$$

Proof By definition of the (x, w) update step, we have

$$\mathbb{T}^{-1}(x - x') - K^\top w \in \partial \mathbf{G}_z(x') \quad (43)$$

and

$$\Sigma^{-1}(w - w') + K(2x' - x) \in \partial \mathbf{F}_v^*(w') ,$$

and from (33) we know that this last expression implies

$$w' \in \partial \mathbf{F}_v(\Sigma^{-1}(w - w') + K(2x' - x)) . \quad (44)$$

Similarly, by definition of $(x_{z,v}^*, w_{z,v}^*)$, we also have

$$-K^\top w_{z,v}^* \in \partial \mathbf{G}_z(x_{z,v}^*) \quad (45)$$

and

$$w_{z,v}^* \in \partial \mathbf{F}_v(Kx_{z,v}^*) . \quad (46)$$

Therefore, combining these expressions, we have

$$\begin{aligned} & \left\langle \left(\begin{array}{c} x' - x_{z,v}^* \\ \Sigma^{-1}(w - w') + K(2x' - x) - Kx_{z,v}^* \end{array} \right), \left(\begin{array}{c} \partial \mathbf{G}_z(x') - \partial \mathbf{G}_z(x_{z,v}^*) \\ \partial \mathbf{F}_v(K(2x' - x) + \Sigma^{-1}(w - w')) - \partial \mathbf{F}_v(Kx_{z,v}^*) \end{array} \right) \right\rangle \\ & \ni \left\langle \left(\begin{array}{c} x' - x_{z,v}^* \\ \Sigma^{-1}(w - w') + K(2x' - x) - Kx_{z,v}^* \end{array} \right), \left(\begin{array}{c} \mathbb{T}^{-1}(x - x') - K^\top(w - w_{z,v}^*) \\ w' - w_{z,v}^* \end{array} \right) \right\rangle \\ & = \left\langle \left(\begin{array}{c} x' - x_{z,v}^* \\ w' - w_{z,v}^* \end{array} \right), \left(\begin{array}{c} \mathbb{T}^{-1}(x - x') - K^\top(w - w') \\ \Sigma^{-1}(w - w') - K(x - x') \end{array} \right) \right\rangle \quad \text{by reorganizing terms} \\ & = \left(\begin{array}{c} x' - x_{z,v}^* \\ w' - w_{z,v}^* \end{array} \right)^\top M \left(\begin{array}{c} x - x' \\ w - w' \end{array} \right) . \end{aligned} \quad (47)$$

As in Chambolle and Pock (2011) we can calculate

$$\begin{aligned} & \left| \left(\begin{array}{c} x' - x_{z,v}^* \\ w' - w_{z,v}^* \end{array} \right)^\top M \left(\begin{array}{c} x - x' \\ w - w' \end{array} \right) \right| \\ & = \frac{1}{2} \left\| \left(\begin{array}{c} x - x_{z,v}^* \\ w - w_{z,v}^* \end{array} \right) \right\|_M^2 - \frac{1}{2} \left\| \left(\begin{array}{c} x' - x_{z,v}^* \\ w' - w_{z,v}^* \end{array} \right) \right\|_M^2 - \frac{1}{2} \left\| \left(\begin{array}{c} x - x' \\ w - w' \end{array} \right) \right\|_M^2 . \end{aligned}$$

On the other hand, by the convexity of F_v and G_z as stated in Assumption 2, we have

$$\begin{aligned}
 & \left\langle \left(\Sigma^{-1}(w - w') + K(2x' - x) - Kx_{z,v}^* \right), \left(\partial G_z(x') - \partial G_z(x_{z,v}^*) \right) \right\rangle \\
 & \geq \|x' - x_{z,v}^*\|_{\Lambda_G}^2 + \|\Sigma^{-1}(w - w') + K(2x' - x) - Kx_{z,v}^*\|_{\Lambda_F}^2 - \tau^2 \|x' - x_{z,v}^*\|_{\text{restrict}}^2 \\
 & \geq \|x' - x_{z,v}^*\|_{\Lambda_G}^2 + \frac{1}{2} \|Kx' - Kx_{z,v}^*\|_{\Lambda_F}^2 - \|\Sigma^{-1}(w - w') - K(x - x')\|_{\Lambda_F}^2 \\
 & \quad - \tau^2 \|x' - x_{z,v}^*\|_{\text{restrict}}^2 \quad \text{using the fact that } (a + b)^2 \geq \frac{1}{2}a^2 - b^2 \text{ for all } a, b \\
 & \geq \frac{C_{\text{cvx}}}{2} \|x' - x_{z,v}^*\|_2^2 - \frac{3\tau^2}{2} \|x' - x_{z,v}^*\|_{\text{restrict}}^2 \\
 & \quad - \|\Sigma^{-1}(w - w') - K(x - x')\|_{\Lambda_F}^2 \quad \text{by Assumption 2} \\
 & \geq \frac{C_{\text{cvx}}}{2} \|x' - x_{z,v}^*\|_2^2 - \frac{3\tau^2}{2} \|x' - x_{z,v}^*\|_{\text{restrict}}^2 - \left\| \begin{pmatrix} x - x' \\ w - w' \end{pmatrix} \right\|_M^2 \cdot \|M^{-1}\|_{\Lambda_F} \left(\|\Sigma^{-1}\| + \|K\| \right) \\
 & \geq \frac{C_{\text{cvx}}}{2} \|x' - x_{z,v}^*\|_2^2 - 6\tau^2 R^2 - \left\| \begin{pmatrix} x - x' \\ w - w' \end{pmatrix} \right\|_M^2 \cdot 2(C_{\text{matrix}})^3,
 \end{aligned}$$

where the last step holds because $\|x'\|_{\text{restrict}}, \|x_{z,v}^*\|_{\text{restrict}} \leq R$, since x' and $x_{z,v}^*$ must both lie in $\text{dom}(G) = \text{dom}(G_z)$ by their definitions. Now, examining these calculations, we see that the left-hand side must be nonnegative, so we can also write

$$\begin{aligned}
 & \left\langle \left(\Sigma^{-1}(w - w') + K(2x' - x) - Kx_{z,v}^* \right), \left(\partial G_z(x') - \partial G_z(x_{z,v}^*) \right) \right\rangle \\
 & \geq c \left(\frac{C_{\text{cvx}}}{2} \|x' - x_{z,v}^*\|_2^2 - 6\tau^2 R^2 - \left\| \begin{pmatrix} x - x' \\ w - w' \end{pmatrix} \right\|_M^2 \cdot 2(C_{\text{matrix}})^3 \right)
 \end{aligned}$$

for any $c \in [0, 1]$. Choosing $c = \frac{1}{\max\{12, 8C_{\text{matrix}}\}}$, we obtain

$$\begin{aligned}
 & \left\langle \left(\Sigma^{-1}(w - w') + K(2x' - x) - Kx_{z,v}^* \right), \left(\partial G_z(x') - \partial G_z(x_{z,v}^*) \right) \right\rangle \\
 & \geq \frac{C_{\text{cvx}}}{2 \max\{12, 8C_{\text{matrix}}\}} \|x' - x_{z,v}^*\|_2^2 - \frac{1}{2}\tau^2 R^2 - \frac{1}{4} \left\| \begin{pmatrix} x - x' \\ w - w' \end{pmatrix} \right\|_M^2
 \end{aligned}$$

Combining all our work, then,

$$\begin{aligned}
 & \frac{C_{\text{cvx}}}{4 \max\{12, 8C_{\text{matrix}}\}} \|x' - x_{z,v}^*\|_2^2 \\
 & \leq \left\| \begin{pmatrix} x - x_{z,v}^* \\ w - w_{z,v}^* \end{pmatrix} \right\|_M^2 - \left\| \begin{pmatrix} x' - x_{z,v}^* \\ w' - w_{z,v}^* \end{pmatrix} \right\|_M^2 - \frac{1}{2} \left\| \begin{pmatrix} x - x' \\ w - w' \end{pmatrix} \right\|_M^2 + \tau^2 R^2.
 \end{aligned}$$

Setting $C_{\text{monotone}} = \min \left\{ \frac{1}{2C_{\text{matrix}}}, \frac{C_{\text{cvx}}}{4 \max\{12, 8C_{\text{matrix}}\}} \right\}$, we have proved the lemma. \blacksquare

Now we turn to the proof of Lemma 4. By Lemma 5, for each $\ell = 1, \dots, L$, we have

$$\begin{aligned} C_{\text{monotone}} \left(\left\| x^{(\ell)} - x_{z,v}^* \right\|_2^2 + \left\| \begin{pmatrix} x^{(\ell-1)} - x^{(\ell)} \\ w^{(\ell-1)} - w^{(\ell)} \end{pmatrix} \right\|_2^2 \right) \\ \leq \left\| \begin{pmatrix} x^{(\ell-1)} - x_{z,v}^* \\ w^{(\ell-1)} - w_{z,v}^* \end{pmatrix} \right\|_M^2 - \left\| \begin{pmatrix} x^{(\ell)} - x_{z,v}^* \\ w^{(\ell)} - w_{z,v}^* \end{pmatrix} \right\|_M^2 + \tau^2 R^2. \end{aligned}$$

Summing this inequality over $\ell = 1, \dots, L$, taking a telescoping sum on the right-hand side, and dividing by L , we have

$$\frac{C_{\text{monotone}}}{L} \sum_{\ell=1}^L \left(\left\| x^{(\ell)} - x_{z,v}^* \right\|_2^2 + \left\| \begin{pmatrix} x^{(\ell-1)} - x^{(\ell)} \\ w^{(\ell-1)} - w^{(\ell)} \end{pmatrix} \right\|_2^2 \right) \leq \frac{1}{L} \left\| \begin{pmatrix} x^{(0)} - x_{z,v}^* \\ w^{(0)} - w_{z,v}^* \end{pmatrix} \right\|_M^2 + \tau^2 R^2.$$

Next, by convexity of $w \mapsto \|w\|_2^2$, we have

$$\left\| \tilde{x} - x_{z,v}^* \right\|_2^2 \leq \frac{1}{L} \sum_{\ell=1}^L \left\| x^{(\ell)} - x_{z,v}^* \right\|_2^2$$

and

$$\left\| \frac{1}{L} \begin{pmatrix} x^{(0)} - x^{(L)} \\ w^{(0)} - w^{(L)} \end{pmatrix} \right\|_2^2 \leq \frac{1}{L} \sum_{\ell=1}^L \left\| \begin{pmatrix} x^{(\ell-1)} - x^{(\ell)} \\ w^{(\ell-1)} - w^{(\ell)} \end{pmatrix} \right\|_2^2.$$

So,

$$C_{\text{monotone}} \left(\left\| \tilde{x} - x_{z,v}^* \right\|_2^2 + \left\| \frac{1}{L} \begin{pmatrix} x^{(0)} - x^{(L)} \\ w^{(0)} - w^{(L)} \end{pmatrix} \right\|_2^2 \right) \leq \frac{1}{L} \left\| \begin{pmatrix} x^{(0)} - x_{z,v}^* \\ w^{(0)} - w_{z,v}^* \end{pmatrix} \right\|_M^2 + \tau^2 R^2.$$

Next, by definition of \tilde{z}, \tilde{v} , we can write

$$\begin{aligned} \left\| \begin{pmatrix} \tilde{z} - x_{z,v}^* \\ \tilde{v} - Kx_{z,v}^* \end{pmatrix} \right\|_{\Theta} &\leq \left\| \begin{pmatrix} \tilde{x} - x_{z,v}^* \\ K\tilde{x} - Kx_{z,v}^* \end{pmatrix} \right\|_{\Theta} \\ &\quad + \left\| \frac{1}{L} \begin{pmatrix} 0 \\ \Sigma^{-1}(w^{(0)} - w^{(L)}) + K(x^{(L)} - x^{(0)}) \end{pmatrix} \right\|_{\Theta} \end{aligned}$$

and so,

$$\begin{aligned} \left\| \begin{pmatrix} \tilde{z} - x_{z,v}^* \\ \tilde{v} - Kx_{z,v}^* \end{pmatrix} \right\|_{\Theta}^2 &\leq \|\Theta\| (1 + 2\|K\| + \|\Sigma^{-1}\|)^2 \left(\left\| \tilde{x} - x_{z,v}^* \right\|_2^2 + \left\| \frac{1}{L} \begin{pmatrix} x^{(0)} - x^{(L)} \\ w^{(0)} - w^{(L)} \end{pmatrix} \right\|_2^2 \right) \\ &\leq \frac{\|\Theta\| (1 + 2\|K\| + \|\Sigma^{-1}\|)^2}{C_{\text{monotone}}} \left(\frac{1}{L} \left\| \begin{pmatrix} x^{(0)} - x_{z,v}^* \\ w^{(0)} - w_{z,v}^* \end{pmatrix} \right\|_M^2 + \tau^2 R^2 \right) \\ &\leq \frac{C_{\text{matrix}}(1 + 3C_{\text{matrix}})^2}{C_{\text{monotone}}} \left(\frac{1}{L} \left\| \begin{pmatrix} x^{(0)} - x_{z,v}^* \\ w^{(0)} - w_{z,v}^* \end{pmatrix} \right\|_M^2 + \tau^2 R^2 \right). \end{aligned}$$

Finally, for each ℓ , by definition of the step,

$$w^{(\ell)} \in \partial F_v \left(\Sigma^{-1}(w^{(\ell-1)} - w^{(\ell)}) + K\bar{x}^{(\ell)} \right)$$

while

$$w_{z,v}^* \in \partial F_v(Kx_{z,v}^*).$$

Therefore, by Assumption 2,

$$\left\| w^{(\ell)} - w_{z,v}^* \right\|_2 \leq C_{\text{Lip}} + C_{\text{grad}} \left\| \Sigma^{-1}(w^{(\ell-1)} - w^{(\ell)}) + K\bar{x}^{(\ell)} - Kx_{z,v}^* \right\|_2.$$

By convexity, then,

$$\begin{aligned} & \left\| \tilde{w} - w_{z,v}^* \right\|_2 \\ & \leq \frac{1}{L} \sum_{\ell=1}^L \left(C_{\text{Lip}} + C_{\text{grad}} \left\| \Sigma^{-1}(w^{(\ell-1)} - w^{(\ell)}) + K\bar{x}^{(\ell)} - Kx_{z,v}^* \right\|_2 \right) \\ & \leq C_{\text{Lip}} + C_{\text{grad}} (\|K\| + \|\Sigma^{-1}\|) \sqrt{\frac{1}{L} \sum_{\ell=1}^L \|x^{(\ell)} - x_{z,v}^*\|_2^2 + \left\| \begin{pmatrix} x^{(\ell-1)} - x^{(\ell)} \\ w^{(\ell-1)} - w^{(\ell)} \end{pmatrix} \right\|_2^2} \\ & \leq C_{\text{Lip}} + C_{\text{grad}} (\|K\| + \|\Sigma^{-1}\|) \sqrt{\frac{1}{C_{\text{monotone}}} \left(\frac{1}{L} \left\| \begin{pmatrix} x^{(0)} - x_{z,v}^* \\ w^{(0)} - w_{z,v}^* \end{pmatrix} \right\|_M^2 + \tau^2 R^2 \right)} \\ & \leq C_{\text{Lip}} + \frac{C_{\text{grad}} (\|K\| + \|\Sigma^{-1}\|)}{C_{\text{monotone}}} \left(\frac{1}{\sqrt{L}} \left\| \begin{pmatrix} x^{(0)} - x_{z,v}^* \\ w^{(0)} - w_{z,v}^* \end{pmatrix} \right\|_M + \tau R \right). \end{aligned}$$

Combining everything, this proves that

$$\begin{aligned} & \left\| \begin{pmatrix} \tilde{x} - x_{z,v}^* \\ \tilde{w} - w_{z,v}^* \end{pmatrix} \right\|_2 \leq C_{\text{Lip}} + \frac{1 + C_{\text{grad}} (\|K\| + \|\Sigma^{-1}\|)}{C_{\text{monotone}}} \left(\frac{1}{\sqrt{L}} \left\| \begin{pmatrix} x^{(0)} - x_{z,v}^* \\ w^{(0)} - w_{z,v}^* \end{pmatrix} \right\|_M + \tau R \right) \\ & \leq C_{\text{Lip}} + \max\{1, \sqrt{\|M\|}\} \frac{1 + C_{\text{grad}} (\|K\| + \|\Sigma^{-1}\|)}{C_{\text{monotone}}} \left(\frac{1}{\sqrt{L}} \left\| \begin{pmatrix} x^{(0)} - x_{z,v}^* \\ w^{(0)} - w_{z,v}^* \end{pmatrix} \right\|_2 + \tau R \right) \\ & \leq C_{\text{Lip}} + \max\{1, \sqrt{C_{\text{matrix}}}\} \frac{1 + 2C_{\text{grad}}C_{\text{matrix}}}{C_{\text{monotone}}} \left(\frac{1}{\sqrt{L}} \left\| \begin{pmatrix} x^{(0)} - x_{z,v}^* \\ w^{(0)} - w_{z,v}^* \end{pmatrix} \right\|_2 + \tau R \right). \end{aligned}$$

Finally, defining

$$C_{\text{iter}} = \max \left\{ \sqrt{\frac{C_{\text{matrix}}(1 + 3C_{\text{matrix}})^2}{C_{\text{monotone}}}}, \max\{1, \sqrt{C_{\text{matrix}}}\} \frac{1 + 2C_{\text{grad}}C_{\text{matrix}}}{C_{\text{monotone}}} \right\},$$

we have proved the lemma.

References

- Rina Foygel Barber, Emil Y. Sidky, Taly Gilat Schmidt, and Xiaochuan Pan. An algorithm for constrained one-step inversion of spectral CT data. *Physics in Medicine and Biology*, 61(10):3784–3818, 2016.
- Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.
- Jérôme Bolte, Shoham Sabach, and Marc Teboulle. Proximal alternating linearized minimization for nonconvex and nonsmooth problems. *Mathematical Programming*, 146(1-2):459–494, 2014.
- Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.
- Emmanuel J Candès, Michael B Wakin, and Stephen P Boyd. Enhancing sparsity by reweighted ℓ_1 minimization. *Journal of Fourier Analysis and Applications*, 14(5-6):877–905, 2008.
- Antonin Chambolle and Jérôme Darbon. On total variation minimization and surface evolution using parametric maximum flows. *International Journal of Computer Vision*, 84(3):288–307, 2009.
- Antonin Chambolle and Thomas Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2011.
- Antonin Chambolle and Thomas Pock. A remark on accelerated block coordinate descent for computing the proximity operators of a sum of convex functions. *SMAI Journal of Computational Mathematics*, 1:29–54, 2015.
- Rick Chartrand. Exact reconstruction of sparse signals via nonconvex minimization. *IEEE Signal Processing Letters*, 14(10):707–710, 2007.
- Aditya Chopra and Heng Lian. Total variation, adaptive total variation and nonconvex smoothly clipped absolute deviation penalty for denoising blocky images. *Pattern Recognition*, 43(8):2609–2619, 2010.
- Ernie Esser, Xiaoqun Zhang, and Tony Chan. A general framework for a class of first order primal-dual algorithms for TV minimization. *UCLA CAM Report*, pages 09–67, 2009.
- Jianqing Fan and Runze Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96(456):1348–1360, 2001.
- Saeed Ghadimi and Guanghui Lan. Accelerated gradient methods for nonconvex nonlinear and stochastic programming. *Mathematical Programming*, 156(1):59–99, 2016.

- Bingsheng He and Xiaoming Yuan. Convergence analysis of primal-dual algorithms for a saddle-point problem: from contraction perspective. *SIAM Journal on Imaging Sciences*, 5(1):119–149, 2012.
- Mingyi Hong, Zhi-Quan Luo, and Meisam Razaviyayn. Convergence analysis of alternating direction method of multipliers for a family of nonconvex problems. *SIAM Journal on Optimization*, 26(1):337–364, 2016.
- David R Hunter and Kenneth Lange. Quantile regression via an MM algorithm. *Journal of Computational and Graphical Statistics*, 9(1):60–77, 2000.
- Nicholas A Johnson. A dynamic programming algorithm for the fused lasso and ℓ_0 -segmentation. *Journal of Computational and Graphical Statistics*, 22(2):246–260, 2013.
- Keith Knight and Wenjiang Fu. Asymptotics for lasso-type estimators. *Annals of Statistics*, 28(5):1356–1378, 2000.
- Guoyin Li and Ting Kei Pong. Global convergence of splitting methods for nonconvex composite optimization. *SIAM Journal on Optimization*, 25(4):2434–2460, 2015.
- Huan Li and Zhouchen Lin. Accelerated proximal gradient methods for nonconvex programming. In *Advances in Neural Information Processing Systems*, pages 379–387, 2015.
- Po-Ling Loh and Martin J Wainwright. High-dimensional regression with noisy and missing data: Provable guarantees with non-convexity. In *Advances in Neural Information Processing Systems*, pages 2726–2734, 2011.
- Po-Ling Loh and Martin J Wainwright. Regularized M-estimators with nonconvexity: Statistical and algorithmic theory for local optima. In *Advances in Neural Information Processing Systems*, pages 476–484, 2013.
- Chengwu Lu and Hua Huang. $TV + TV_2$ regularization with nonconvex sparseness-inducing penalty for image restoration. *Mathematical Problems in Engineering*, 2014:790547, 2014.
- Sindri Magnússon, Pradeep Chathuranga Weeraddana, Michael G Rabbat, and Carlo Fischione. On the convergence of alternating direction Lagrangian methods for nonconvex structured optimization problems. *IEEE Transactions on Control of Network Systems*, 2015.
- MATLAB. *Version 8.6.0 (R2015b)*. The MathWorks Inc., Natick, Massachusetts, 2015.
- Sahand Negahban, Bin Yu, Martin J Wainwright, and Pradeep K Ravikumar. A unified framework for high-dimensional analysis of M-estimators with decomposable regularizers. In *Advances in Neural Information Processing Systems*, pages 1348–1356, 2009.
- Yu Nesterov. Gradient methods for minimizing composite functions, 2013.
- Peter Ochs, Yunjin Chen, Thomas Brox, and Thomas Pock. iPiano: inertial proximal algorithm for nonconvex optimization. *SIAM Journal on Imaging Sciences*, 7(2):1388–1419, 2014.

- Peter Ochs, Alexey Dosovitskiy, Thomas Brox, and Thomas Pock. On iteratively reweighted algorithms for nonsmooth nonconvex optimization in computer vision. *SIAM Journal on Imaging Sciences*, 8:331–372, 2015.
- James M Ortega and Werner C Rheinboldt. *Iterative Solution of Nonlinear Equations in Several Variables*. SIAM, 1970.
- Ankit Parekh and Ivan W Selesnick. Convex fused lasso denoising with non-convex regularization and its use for pulse detection. In *2015 IEEE Signal Processing in Medicine and Biology Symposium (SPMB)*, pages 1–6. IEEE, 2015.
- Thomas Pock and Antonin Chambolle. Diagonal preconditioning for first order primal-dual algorithms in convex optimization. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1762–1769. IEEE, 2011.
- R Tyrrell Rockafellar. *Convex analysis*, 1997.
- Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1):259–268, 1992.
- Ivan W Selesnick, Ankit Parekh, and Ilker Bayram. Convex 1-d total variation denoising with non-convex regularization. *IEEE Signal Processing Letters*, 22(2):141–144, 2015.
- Emil Y Sidky, Rick Chartrand, John M Boone, and Xiaochuan Pan. Constrained minimization for enhanced exploitation of gradient sparsity: application to CT image reconstruction. *IEEE Journal of Translational Engineering in Health and Medicine*, 2:1800418, 2014.
- Robert Tibshirani, Michael Saunders, Saharon Rosset, Ji Zhu, and Keith Knight. Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(1):91–108, 2005.
- Tuomo Valkonen. A primal–dual hybrid gradient method for nonlinear operators with applications to MRI. *Inverse Problems*, 30(5):055012, 2014.
- Fenghui Wang, Zongben Xu, and Hong-Kun Xu. Convergence of Bregman alternating direction method with multipliers for nonconvex composite problems. *arXiv preprint arXiv:1410.8625*, 2014a.
- Fenghui Wang, Wenfei Cao, and Zongben Xu. Convergence of multi-block Bregman ADMM for nonconvex composite problems. *arXiv preprint arXiv:1505.03063*, 2015a.
- Huahua Wang and Arindam Banerjee. Bregman alternating direction method of multipliers. In *Advances in Neural Information Processing Systems*, pages 2816–2824, 2014.
- Huahua Wang, Arindam Banerjee, and Zhi-Quan Luo. Parallel direction method of multipliers. In *Advances in Neural Information Processing Systems*, pages 181–189, 2014b.
- Yu-Xiang Wang, James Sharpnack, Alexander J Smola, and Ryan J Tibshirani. Trend filtering on graphs. In *Conference on Artificial Intelligence and Statistics (AISTATS)*, 2015b.

Cun-Hui Zhang. Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics*, 38(2):894–942, 2010.