

Solving variational inequalities with monotone operators on domains given by Linear Minimization Oracles

Anatoli Juditsky · Arkadi Nemirovski

Received: 17 December 2013 / Accepted: 14 February 2015
© Springer-Verlag Berlin Heidelberg and Mathematical Optimization Society 2015

Abstract The standard algorithms for solving large-scale convex–concave saddle point problems, or, more generally, variational inequalities with monotone operators, are *proximal type* algorithms which at every iteration need to compute a *prox-mapping*, that is, to minimize over problem’s domain X the sum of a linear form and the specific convex *distance-generating function* underlying the algorithms in question. (Relative) computational simplicity of prox-mappings, which is the standard requirement when implementing proximal algorithms, clearly implies the possibility to equip X with a relatively computationally cheap *Linear Minimization Oracle* (LMO) able to minimize over X linear forms. There are, however, important situations where a cheap LMO indeed is available, but where no proximal setup with easy-to-compute prox-mappings is known. This fact motivates our goal in this paper, which is to develop techniques for solving variational inequalities with monotone operators on domains given by LMO. The techniques we discuss can be viewed as a substantial extension of the proposed in Cox et al. (Math Program Ser B 148(1–2):143–180, 2014) method of nonsmooth convex minimization over an LMO-represented domain.

Mathematics Subject Classification 65K15 · 90C25 · 90C47 · 68T10

Research of the first author was supported by the CNRS-Mastodons project GARGANTUA, and the LabEx PERSYVAL-Lab (ANR-11-LABX-0025). Research of the second author was supported by the NSF grants CMMI 1232623 and CCF 1415498.

A. Juditsky (✉)
LJK, Université Grenoble Alpes, B.P. 53, 38041 Grenoble Cedex 9, France
e-mail: juditsky@imag.fr

A. Nemirovski
Georgia Institute of Technology, Atlanta, GA 30332, USA
e-mail: nemirovs@isye.gatech.edu

1 Introduction

The majority of First order methods (FOM's) for large-scale convex minimization (and *all* known to us FOM's for large-scale convex-concave saddle point problems and variational inequalities with monotone operators) are of *proximal* type: at a step of the algorithm, one needs to compute *prox-mapping*—to minimize over problem's domain the sum of a linear function and a specific for the algorithm strongly convex *distance generating* function (d.-g.f.), in the simplest case, just squared Euclidean norm. As a result, the practical scope of proximal algorithms is restricted to *proximal-friendly* domains—those allowing for d.-g.f.'s with not too expensive computationally prox-mappings. What follows is motivated by the desire to develop FOM's for solving convex–concave saddle point problems on bounded domains with “difficult geometry”—those for which no d.-g.f.'s resulting in nonexpensive prox-mappings (and thus no “implementable” proximal methods) are known. In what follows, we relax the assumption on problem's domain to be proximal-friendly to the weaker assumption to admit computationally nonexpensive *Linear Minimization Oracle* (LMO)—a routine capable to minimize a linear function over the domain. This indeed is a relaxation: to minimize within a desired, whatever high, accuracy a linear form over a bounded proximal-friendly domain is the same as to minimize over the domain the sum of large multiple of the form and the d.-g.f. Thus, proximal friendliness implies existence of a nonexpensive LMO, but not vice versa. For example, when the domain is the ball B_n of nuclear norm in $\mathbf{R}^{n \times n}$, computing prox-mapping, for all known proximal setups, requires full singular value decomposition of an $n \times n$ matrix, which can be prohibitively time consuming when n is large. In contrast to this, minimizing a linear form over B_n only requires finding the leading singular vectors of an $n \times n$ matrix, which is much easier than full-fledged singular value decomposition.

Recently, there was significant interest in solving convex minimization problems on domains given by LMO's. The emphasis in this line of research is on smooth/smooth norm-regularized convex minimization [11, 13, 14, 16, 17, 31], where the main “working horse” is the classical Conditional Gradient (a.k.a. Frank-Wolfe) algorithm originating from [12] and intensively studied in 1970's (see [9, 10, 28] and references therein). Essentially, Conditional Gradient is the only traditional convex optimization technique capable to handle convex minimization problems on LMO-represented domains. In its standard form, Conditional Gradient algorithm, to the best of our knowledge, is not applicable beyond the smooth minimization setting; we are not aware of any attempt to apply this algorithm even to the simplest–bilinear–saddle point problems. The approach proposed in this paper is different and is inspired by our recent paper [7], where a method for *nonsmooth* convex minimization over an LMO-represented convex domain was developed. The latter method utilizes *Fenchel-type* representations of the objective in order to pass from the problem of interest to its special dual. In many important cases the domain of the dual problem is proximal-friendly, so that the dual problem can be solved by proximal FOM's. We then use the machinery of *accuracy certificates* originating from [26] allowing to recover a good solution to the problem of interest from the information accumulated when solving the dual problem. In this paper we follow the same strategy in the context of variational inequalities (v.i.'s) with monotone operators (this covers, in particular, convex–concave saddle

point problems). Specifically, we introduce the notion of a *Fenchel-type representation* of a monotone operator, allowing to associate with the v.i. of interest its dual, which is again a v.i. with monotone operator with the values readily given by the representation and the *LMO representing the domain of the original v.i.* Then we solve the dual v.i. (e.g., by a proximal-type algorithm) and use the machinery of accuracy certificates to recover a good solution to the v.i. of interest from the information gathered when solving the dual v.i. Note that by themselves extensions of Fenchel duality beyond the scope of convex functions is not new; it originates with the seminal papers of Rockafellar [29] on conjugate saddle functions and of Mosco [22] on dual variational inequalities; for further developments, see [4, 8, 30] and references therein. To the best of our knowledge, the notion of *Fenchel-type representation* and based upon this notion algorithmic constructions proposed in this paper are new.

The main body of the paper is organized as follows. Section 2 outlines the background of convex–concave saddle point problems, variational inequalities with monotone operators and accuracy certificates. In Sect. 3, we introduce the notion of a Fenchel-type representation of a monotone operator and the induced by this notion concept of v.i. dual to a given v.i. This section also contains a simple fully algorithmic “calculus” of Fenchel-type representations of monotone operators: it turns out that basic monotonicity-preserving operations with these operators (summation, affine substitution of argument, etc.) as applied to operands given by Fenchel-type representations yield similar representation for the result of the operation. As a consequence, our abilities to operate numerically with Fenchel-type representations of monotone operators are comparable with our abilities to evaluate the operators themselves. Section 4 contains our main result—Theorem 1. It shows how information collected when solving the dual v.i. to some accuracy, can be used to build an approximate solution of the same accuracy to the primal v.i. In Sect. 4 we present a self-contained description of two well known proximal type algorithms for v.i.’s with monotone operators—Mirror Descent (MD) and Mirror Prox (MP)—which indeed are capable to collect the required information. Section 5 is devoted to some modifications of our approach. In the concluding Sect. 6, we illustrate the proposed approach by applying it to the “matrix completion problem with spectral norm fit”—to the problem

$$\min_{\substack{u \in \mathbf{R}^{n \times n}, \\ \|u\|_{\text{nuc}} \leq 1}} \|Au - b\|_{2,2},$$

where $\|x\|_{\text{nuc}} = \sum_i \sigma_i(x)$ is the nuclear norm, $\sigma(x)$ being the singular spectrum of x , $\|x\|_{2,2} = \max_i \sigma_i(x)$ is the spectral norm, and $u \mapsto Au$ is a linear mapping from $\mathbf{R}^{n \times n}$ to $\mathbf{R}^{m \times m}$.

2 Preliminaries

2.1 Variational inequalities and related accuracy measures

Let Y be a nonempty closed convex set in Euclidean space E_y and $H(\cdot) : Y \rightarrow E_y$ be a monotone operator:

$$\langle H(y) - H(y'), y - y' \rangle \geq 0 \quad \forall y, y' \in Y.$$

The variational inequality (v.i.) associated with $(H(\cdot), Y)$ is

$$\text{find } y_* \in Y : \langle H(z), z - y_* \rangle \geq 0 \quad \forall z \in Y; \quad \text{VI}(H, Y)$$

(every) $y_* \in Y$ satisfying the target relation in $\text{VI}(H, Y)$ is called a weak solution to the v.i.; when Y is convex and compact, and $H(\cdot)$ is monotone on Y , weak solutions always exist. A strong solution to v.i. is a point $y_* \in Y$ such that $\langle H(y_*), y - y_* \rangle \geq 0$ for all $y \in Y$; from the monotonicity of $H(\cdot)$ it follows that a strong solution is a weak one as well. Note that when $H(\cdot)$ is monotone and continuous on Y (this is the only case we will be interested in), *weak solutions are exactly the strong solutions*.¹

The accuracy measure naturally quantifying the inaccuracy of a candidate solution $y \in Y$ to $\text{VI}(H, Y)$ is the *dual gap function*

$$\epsilon_{\text{vi}}(y|H, Y) = \sup_{z \in Y} \langle H(z), y - z \rangle;$$

this (clearly nonnegative for $y \in Y$) quantity is zero if and only if y is a weak solution to the v.i.

We will be interested also in the *saddle point case* where $Y = V \times W$ is the direct product of nonempty convex compact subsets $V \subset E_v$ and $W \subset E_w$ of Euclidean spaces E_v, E_w , and $H(\cdot)$ is associated with Lipschitz continuous function $f(v, w) : Y = V \times W \rightarrow \mathbf{R}$ convex in $v \in V$ and concave in $w \in W$:

$$H(y = [v; w]) = [H_v(v, w); H_w(v, w)] \quad \text{with} \quad H_v(v, w) \in \partial_v f(v, w), \\ H_w(v, w) \in \partial_w [-f(v, w)].$$

Here, for a convex function $g : G \rightarrow \mathbf{R}$ defined on a convex set G , and for $u \in G$, $\partial g(x) = \{h : g(x) + h^T(y - x) \leq g(y) \quad \forall y \in G\}$ is the taken at x subdifferential of the function which coincides with g on G and is $+\infty$ outside of G . It is well-known that when g is Lipschitz continuous, $\partial g(x)$ is nonempty for every $x \in G$, so that the above H does exist, although is not uniquely defined by f . When f is continuously differentiable, we can set $H_v(v, w) = \nabla_v f(v, w)$, $H_w(v, w) = -\nabla_w f(v, w)$, and this is the option we use, in the smooth case, in the sequel.

We can associate with the saddle point case two optimization problems

$$\text{Opt}(P) = \min_{v \in V} \left[\overline{f}(v) = \sup_{w \in W} f(v, w) \right] \quad (P) \\ \text{Opt}(D) = \max_{w \in W} \left[\underline{f}(w) = \inf_{v \in V} f(v, w) \right] \quad (D);$$

under our assumptions (V, W are convex and compact, f is continuous convex-concave) these problems are solvable with equal optimal values. We associate with a

¹ In literature on v.i.'s, a variational inequality where a weak solution is sought is often called *Minty's v.i.*, and one where a strong solution is sought—*Stampacchia's v.i.* Equivalence of the notions of weak and strong solutions in the case of continuous monotone operator is the finite dimensional version of the classical Minty's Lemma (1967), see [6].

pair $(v, w) \in V \times W$ the *saddle point inaccuracy*

$$\epsilon_{\text{sad}}(v, w|f, V, W) = \overline{f}(v) - \underline{f}(w) = [\overline{f}(v) - \text{Opt}(P)] + [\text{Opt}(D) - \underline{f}(w)].$$

2.2 Accuracy certificates

Given Y, H , let us call a collection $\mathcal{C}^N = \{y_t \in Y, \lambda_t \geq 0, H(y_t)\}_{t=1}^N$ with $\sum_t \lambda_t = 1$, an *N -step accuracy certificate*. For $Z \subset Y$, we call the quantity

$$\text{Res}(\mathcal{C}^N|Z) = \sup_{y \in Z} \sum_{t=1}^N \lambda_t \langle H(y_t), y_t - y \rangle$$

the *resolution* of the certificate \mathcal{C}^N w.r.t. Z .

Let us make two observations coming back to [26]:

Lemma 1 *Let Y be a closed convex set in Euclidean space E_y , $H(\cdot)$ be a monotone operator on Y , and $\mathcal{C}^N = \{y_t \in Y, \lambda_i \geq 0, H(y_t)\}_{t=1}^N$ be an accuracy certificate. Setting*

$$\widehat{y} = \sum_{t=1}^N \lambda_t y_t,$$

we have $\widehat{y} \in Y$, and for every nonempty closed convex subset Y' of Y it holds

$$\epsilon_{\text{vi}}(\widehat{y}|H, Y') \leq \text{Res}(\mathcal{C}^N|Y'). \tag{1}$$

In the saddle point case we have also

$$\epsilon_{\text{sad}}(\widehat{y}|f, V, W) \leq \text{Res}(\mathcal{C}^N|Y). \tag{2}$$

Proof For $z \in Y'$ we have

$$\begin{aligned} \langle H(z), z - \sum_t \lambda_t y_t \rangle &= \sum_t \lambda_t \langle H(z), z - y_t \rangle \quad [\text{since } \sum_t \lambda_t = 1] \\ &\geq \sum_t \lambda_t \langle H(y_t), z - y_t \rangle \quad [\text{since } H \text{ is monotone and } \lambda_t \geq 0] \\ &\geq -\text{Res}(\mathcal{C}^N|Y') \quad [\text{by definition of resolution}]. \end{aligned}$$

Thus, $\langle H(z), \widehat{y} - z \rangle \leq \text{Res}(\mathcal{C}^N|Y')$ for all $z \in Y'$, and (1) follows. In the saddle point case, setting $y_t = [v_t; w_t]$, $\widehat{y} = [\widehat{v}; \widehat{w}]$, for every $y = [v; w] \in Y = V \times W$ we have

$$\begin{aligned} \text{Res}(\mathcal{C}^N|Y) &\geq \sum_t \lambda_t \langle H(y_t), y_t - y \rangle \quad [\text{by definition of resolution}] \\ &= \sum_t \lambda_t [\langle H_v(v_t, w_t), v_t - v \rangle + \langle H_w(v_t, w_t), w_t - w \rangle] \end{aligned}$$

$$\begin{aligned} &\geq \sum_t \lambda_t [[f(v_t, w_t) - f(v, w_t)] + [f(v_t, w) - f(v, w_t)]] \\ &\quad \text{[by origin of } H \text{ and since } f(v, w) \text{ is convex in } v \text{ and concave in } w] \\ &= \sum_t \lambda_t [f(v_t, w) - f(v, w_t)] \\ &\geq f(\widehat{v}, w) - f(v, \widehat{w}) \quad \text{[since } f \text{ is convex - concave].} \end{aligned}$$

Since the resulting inequality holds true for all $v \in V, w \in W$, we get $\overline{f}(\widehat{v}) - \underline{f}(\widehat{w}) \leq \text{Res}(\mathcal{C}^N | Y)$, and (2) follows. \square

Lemma 1 can be partially inverted in the case of *skew-symmetric* operator $H(\cdot)$, that is,

$$H(y) = a + Sy \tag{3}$$

with *skew-symmetric* ($S = -S^*$)² linear operator S . A skew-symmetric $H(\cdot)$ clearly satisfies the identity

$$\langle H(y), y - y' \rangle = \langle H(y'), y - y' \rangle, \quad y, y' \in E_y,$$

(since for a skew-symmetric S it holds $\langle Sx, x \rangle \equiv 0$).

Lemma 2 *Let Y be a convex compact set in Euclidean space E_y , $H(y) = a + Sy$ be skew-symmetric, and let $\mathcal{C}^N = \{y_t \in Y, \lambda_t \geq 0, H(y_t)\}_{t=1}^N$ be an accuracy certificate. Then for $\widehat{y} = \sum_t \lambda_t y_t$ it holds*

$$\epsilon_{\text{vi}}(\widehat{y} | H, Y) = \text{Res}(\mathcal{C}^N | Y). \tag{4}$$

Proof We already know that $\epsilon_{\text{vi}}(\widehat{y} | H, Y) \leq \text{Res}(\mathcal{C}^N | Y)$. To prove the inverse inequality, note that for every $y \in Y$ we have

$$\begin{aligned} \epsilon_{\text{vi}}(\widehat{y} | H, Y) &\geq \langle H(y), \widehat{y} - y \rangle = \langle H(\widehat{y}), \widehat{y} - y \rangle \\ &\quad \text{[since } H \text{ is skew-symmetric]} \\ &= \langle a, \widehat{y} - y \rangle - \langle S\widehat{y}, y \rangle + \langle S\widehat{y}, \widehat{y} \rangle \\ &= \langle a, \widehat{y} - y \rangle - \langle S\widehat{y}, y \rangle = \sum_t \lambda_t [\langle a, y_t - y \rangle - \langle Sy_t, y \rangle] \\ &\quad \text{[due to } S^* = -S] \qquad \qquad \qquad \text{[due to } \widehat{y} = \sum_t \lambda_t y_t \text{ and } \sum_t \lambda_t = 1] \end{aligned}$$

² From now on, for a linear mapping $x \mapsto Bx : E \rightarrow F$, where E, F are Euclidean spaces, B^* denotes the conjugate of B , that is, a linear mapping $y \mapsto B^*y : F \rightarrow E$ uniquely defined by the identity $\langle Bx, y \rangle = \langle x, B^*y \rangle$ for all $x \in E, y \in F$.

$$\begin{aligned}
 &= \sum_t \lambda_t [\langle a, y_t - y \rangle + \langle S y_t, y_t - y \rangle] \\
 &= \sum_t \lambda_t \langle H(y_t), y_t - y \rangle. \\
 &\quad \text{[due to } S^* = -S]
 \end{aligned}$$

Thus, $\sum_t \lambda_t \langle H(y_t), y_t - y \rangle \leq \epsilon_{\text{vi}}(\hat{y}|H, Y)$ for all $y \in Y$, so that $\text{Res}(\mathcal{C}^N|Y) \leq \epsilon_{\text{vi}}(\hat{y}|H, Y)$. \square

Corollary 1 *Assume we are in the saddle point case, so that $Y = V \times W$ is a direct product of two convex compact sets, and the monotone operator $H(\cdot)$ is associated with a convex–concave function $f(v, w)$. Assume also that f is bilinear: $f(v, w) = \langle a, v \rangle + \langle b, w \rangle + \langle w, Av \rangle$, so that $H(\cdot)$ is affine and skew-symmetric. Then for every $y \in Y$ it holds*

$$\epsilon_{\text{sad}}(y|f, V, W) \leq \epsilon_{\text{vi}}(y|H, Y). \tag{5}$$

Proof Consider accuracy certificate $\mathcal{C}^1 = \{y_1 = y, \lambda_1 = 1, H(y_1)\}$; for this certificate, \hat{y} as defined in Lemma 2 is just y . Therefore, by Lemma 2, $\text{Res}(\mathcal{C}^1|Y) = \epsilon_{\text{vi}}(y|H, Y)$. This equality combines with Lemma 1 to imply (5). \square

3 Representations of monotone operators

3.1 Outline

To explain the origin of the developments to follow, let us summarize the approach to solving convex minimization problems on domains given by LMOs, developed in [7]. The principal ingredient of this approach is a *Fenchel-type representation* of a convex function $f : X \rightarrow \mathbf{R}$ defined on a convex subset X of Euclidean space E ; by definition, such a representation is

$$f(x) = \max_{y \in Y} [\langle x, Ay + a \rangle - \psi(y)], \tag{6}$$

where Y is a convex subset of Euclidean space F and $\psi : Y \rightarrow \mathbf{R}$ is convex. Assuming for the sake of simplicity that X, Y are compact and ψ is continuously differentiable on Y , representation (6) allows to associate with the *primal problem*

$$\text{Opt}(P) = \min_{x \in X} f(x) \tag{P}$$

its dual

$$\text{Opt}(D) = \max_{y \in Y} \left[f_*(y) = \min_{x \in X} \langle x, Ay + a \rangle - \psi(y) \right] \tag{D}$$

with the same optimal value. Observe that the first order information on the (concave) objective of (D) is readily given by the first order information on ψ and the information

provided by an LMO for X . As a result, we can solve (D) by, say, a proximal type FOM, provided that Y is proximal-friendly. The crucial in this approach question is how to recover a good approximate solution to the problem of interest (P) from the information collected when solving (D) . This question is addressed via the machinery of accuracy certificates [7,26].

In the sequel, we intend to apply a similar scheme to the situation where the role of (P) is played by a variational inequality with monotone operator on a convex compact domain X given by an LMO. Our immediate task is to outline *informally* what a Fenchel-type *representation* of a monotone operator is and how we intend to use such a representation. To this end note that (P) and (D) can be reduced to variational inequalities with monotone operators, specifically

- the “primal” v.i. stemming from (P) . The domain of this v.i. is X , and the operator is $f'(x) = Ay(x) + a$, where $y(x)$ is a maximizer of the function $\langle x, Ay \rangle - \psi(y)$ over $y \in Y$, or, which is the same, a (strong) solution to the v.i. given by the domain Y and the monotone operator $y \mapsto G(y) - A^*x$, where $G(y) = \psi'(y)$;
- the “dual” v.i. stemming from (D) . The domain of this v.i. is Y , and the operator is $y \mapsto G(y) - A^*x(y)$, where $x(y)$ is a minimizer of $\langle x, Ay + a \rangle$ over $x \in X$.

Observe that both operators in question are described in terms of a monotone operator G on Y and affine mapping $y \mapsto Ay + a : F \rightarrow E$; in the above construction G was the gradient field of ψ , but the construction of the primal and the dual v.i.’s makes sense whenever G is a monotone operator on Y satisfying minimal regularity assumptions. The idea of the approach we are about to develop is as follows: in order to solve a v.i. with a monotone operator Φ and domain X given by an LMO,

- (A) We represent Φ in the form of $\Phi(x) = Ay(x) + a$, where $y(x)$ is a strong solution to the v.i. on Y given by the operator $G(y) - A^*x$, G being an appropriate monotone operator on Y .

It can be shown that a desired representation always exists, but by itself existence does not help much—we need the representation to be suitable for numerical treatment, to be available in a “closed computation-friendly form.” We show that “computation-friendly” representations of monotone operators admit a kind of *fully algorithmic* calculus which, for all basic monotonicity-preserving operations, allows to get straightforwardly a desired representation of the result of an operation from the representations of the operands. In view of this calculus, “closed analytic form” representations, allowing to compute efficiently the values of monotone operators, automatically lead to required computation-friendly representations.

- (B) We use the representation from A to build the “dual” v.i. with domain Y and the operator $\Theta(y) = G(y) - A^*x(y)$, with exactly the same $x(y)$ as above, that is, $x(y) \in \text{Argmin}_{x \in X} \langle x, Ay + a \rangle$. We shall see that Θ is monotone, and that usually there is a significant freedom in choosing Y ; in particular, we typically can choose Y to be proximal-friendly. The latter is the case, e.g., when Φ is affine, see Remark 2.
- (C) We solve the dual v.i. by an algorithm, like MD or MP, which produce necessary accuracy certificates. We will see—and this is our main result—that such a certificate \mathcal{C}^N can be converted straightforwardly into a feasible solution x^N

to the v.i. of interest such that $\epsilon_{vi}(x^N | \Phi, X) \leq \text{Res}(\mathcal{C}^N | Y)$. As a result, if the certificates in question are *good*, meaning that the resolution of \mathcal{C}^N as a function of N obeys the standard efficiency estimates of the algorithm used to solve the dual v.i., we solve the v.i. of interest with the same efficiency estimate as the one for the dual v.i. It remains to note that most of the existing first order algorithms for solving v.i.'s with monotone operators (various versions of polynomial time cutting plane algorithms, like the Ellipsoid method, Subgradient/MD, and different bundle-level versions of MD) indeed produce good accuracy certificates, see [7,26].

3.2 The construction

3.2.1 Situation

Consider the situation where we are given

- an affine mapping

$$y \mapsto Ay + a : F \rightarrow E,$$

where E, F are Euclidean spaces;

- a nonempty closed convex set $Y \subset F$;
- a continuous monotone operator

$$G(y) : Y \rightarrow F$$

which is *good* w.r.t. A, Y , goodness meaning that the variational inequality $\text{VI}(G(\cdot) - A^*x, Y)$ has a strong solution for every $x \in E$. Note that when Y is convex compact, every continuous monotone operator on Y is good, whatever be A ;

- a nonempty convex compact set X in E .

These data give rise to two operators: “primal” $\Phi : X \rightarrow E$ which is monotone, and “dual” $\Psi : Y \rightarrow F$ which is antimonotone (that is, $-\Psi$ is monotone).

3.2.2 Primal monotone operator

The *primal operator* $\Phi : E \rightarrow E$ is defined by

$$\Phi(x) = Ay(x) + a, \tag{7}$$

where $y(\cdot) : E \rightarrow F$ is a mapping satisfying

$$y(x) \in Y, \langle A^*x - G(y(x)), y(x) - y \rangle \geq 0, \forall y \in Y. \tag{8}$$

Observe that mappings $y(x)$ satisfying (8) do exist: given x , we can take as $y(x)$ (any) strong solution to the variational inequality given by the monotone operator

$G(y) - A^*x$ and the domain Y . Note that $y(\cdot)$ is not necessarily uniquely defined by the data $A, G(\cdot), Y$, so that $y(\cdot)$ should be considered as additional to $A, a, G(\cdot), Y$ data element participating in the description (7) of Φ .

Now, with $x', x'' \in E$, setting $y(x') = y', y(x'') = y''$, so that $y', y'' \in Y$, we have

$$\begin{aligned} \langle \Phi(x') - \Phi(x''), x' - x'' \rangle &= \langle Ay' - Ay'', x' - x'' \rangle = \langle y' - y'', A^*x' - A^*x'' \rangle \\ &= \langle y' - y'', A^*x' \rangle + \langle y'' - y', A^*x'' \rangle \\ &= \langle y' - y'', A^*x' - G(y') \rangle + \langle y'' - y', A^*x'' - G(y'') \rangle \\ &\quad + \langle G(y'), y' - y'' \rangle + \langle G(y''), y'' - y' \rangle \\ &= \underbrace{\langle y' - y'', A^*x' - G(y') \rangle}_{\geq 0 \text{ due to } y'=y(x')} + \underbrace{\langle y'' - y', A^*x'' - G(y'') \rangle}_{\geq 0 \text{ due to } y''=y(x'')} \\ &\quad + \underbrace{\langle G(y') - G(y''), y' - y'' \rangle}_{\geq 0 \text{ since } G \text{ is monotone}} \geq 0. \end{aligned}$$

Thus, $\Phi(x)$ is monotone. We call (7) a *representation* of the monotone operator Φ , and the data $F, A, a, y(\cdot), G(\cdot), Y$ – the *data* of the representation. We also say that these data *represent* Φ .

Given a convex domain $X \subset E$ and a monotone operator $\bar{\Phi}$ on this domain, we say that data $F, A, a, y(\cdot), G(\cdot), Y$ of the above type *represent* $\bar{\Phi}$ on X , if the monotone operator Φ represented by these data coincides with $\bar{\Phi}$ on X ; note that in this case the same data represent the restriction of $\bar{\Phi}$ on any convex subset of X .

3.2.3 Dual operator

Operator $\Psi : Y \rightarrow F$ is given by

$$\Psi(y) = A^*x(y) - G(y) : x(y) \in X, \langle Ay + a, x(y) - x \rangle \leq 0, \forall x \in X \quad (9)$$

(in words: $\Psi(y) = A^*x(y) - G(y)$, where $x(y)$ minimizes $\langle Ay + a, x \rangle$ over $x \in X$). This operator clearly is antimonotone, as the sum of two antimonotone operators $-G(y)$ and $A^*x(y)$; antimonotonicity of the latter operator stems from the fact that it is obtained from the antimonotone operator $\psi(z)$ – a section of the superdifferential $\text{Argmin}_{x \in X} \langle z, x \rangle$ of a concave function – by affine substitution of variables: $A^*x(y) = A^*\psi(Ay + a)$, and this substitution preserves antimonotonicity.

Remark 1 Note that computing the value of Ψ at a point y reduces to computing $G(y)$, $Ay + a$, a single call to the Linear Minimization Oracle for X to get $x(y)$, and computing $A^*x(y)$.

3.3 Calculus of representations

3.3.1 Multiplication by nonnegative constants

Let $F, A, a, y(\cdot), G(\cdot), Y$ represent a monotone operator $\Phi : E \rightarrow E$:

$$\Phi(x) = Ay(x) + a : y(x) \in Y \text{ and } \langle A^*x - G(y(x)), y(x) - y \rangle \geq 0 \forall y \in Y.$$

For $\lambda \geq 0$, we clearly have

$$\lambda\Phi(x) = [\lambda A]y(x) + [\lambda a] : \langle [\lambda A]^*x - [\lambda G(y(x))], y(x) - y \rangle \geq 0 \forall y \in Y,$$

that is, a representation of $\lambda\Phi$ is given by $F, \lambda A, \lambda a, y(\cdot), \lambda G(\cdot), Y$; note that the operator λG clearly is good w.r.t. $\lambda A, Y$, since G is good w.r.t. A, Y .

3.3.2 Summation

Let $F_i, A_i, a_i, y_i(\cdot), G_i(\cdot), Y_i, 1 \leq i \leq m$, represent monotone operators $\Phi_i(x) : E \rightarrow E$:

$$\Phi_i(x) = A_i y_i(x) + a_i : y_i(x) \in Y_i \text{ and } \langle A_i^*x - G_i(y_i(x)), y_i(x) - y_i \rangle \geq 0 \forall y_i \in Y_i.$$

Then

$$\begin{aligned} \sum_i \Phi_i(x) &= [A_1, \dots, A_m][y_1(x); \dots; y_m(x)] + [a_1 + \dots + a_m], \\ y(x) &:= [y_1(x); \dots; y_m(x)] \in Y := Y_1 \times \dots \times Y_m, \\ \langle [A_1, \dots, A_m]^*x - [G_1(y_1(x)); \dots; G_m(y_m(x))], [y_1(x); \dots; y_m(x)] \\ &\quad - [y_1; \dots; y_m] \rangle = \sum_i \langle A_i^*x - G_i(y_i(x)), y_i(x) - y_i \rangle \geq 0 \forall y = [y_1; \dots; y_m] \in Y, \end{aligned}$$

so that the data

$$\begin{aligned} F &= F_1 \times \dots \times F_m, \quad A = [A_1, \dots, A_m], \quad a = a_1 + \dots + a_m, \\ y(x) &= [y_1(x); \dots; y_m(x)], \quad G(y) = [G_1(y_1); \dots; G_m(y_m)], \quad Y = Y_1 \times \dots \times Y_m \end{aligned}$$

represent $\sum_i \Phi_i(x)$. Note that the operator $G(\cdot)$ clearly is good since G_1, \dots, G_m are so.

3.3.3 Affine substitution of argument

Let $F, A, a, y(\cdot), G(\cdot), Y$ represent $\Phi : E \rightarrow E$, let H be a Euclidean space and $h \mapsto Qh + q$ be an affine mapping from H to E . We have

$$\begin{aligned} \widehat{\Phi}(h) &:= Q^*\Phi(Qh + q) = Q^*(Ay(Qh + q) + a) : \\ y(Qh + q) &\in Y \text{ and } \langle A^*[Qh + q] - G(y(Qh + q)), y(Qh + q) - y \rangle \geq 0 \forall y \in Y \\ \Rightarrow \text{with } \widehat{A} &= Q^*A, \widehat{a} = Q^*a, \widehat{G}(y) = G(y) - A^*q, \widehat{y}(h) = y(Qh + q) \text{ we have} \\ \widehat{\Phi}(h) &= \widehat{A}\widehat{y}(h) + \widehat{a} : \widehat{y}(h) \in Y \text{ and } \langle \widehat{A}^*h - \widehat{G}(\widehat{y}(h)), \widehat{y}(h) - y \rangle \geq 0 \forall y \in Y, \end{aligned}$$

that is, $F, \widehat{A}, \widehat{a}, \widehat{y}(\cdot), \widehat{G}(\cdot), Y$ represent $\widehat{\Phi}$. Note that \widehat{G} clearly is good since G is so.

3.3.4 Direct sum

Let $F_i, A_i, a_i, y_i(\cdot), G_i(\cdot), Y_i, 1 \leq i \leq m$, represent monotone operators $\Phi_i(x_i) : E_i \rightarrow E_i$. Denoting by $\text{Diag}\{A_1, \dots, A_m\}$ the block-diagonal matrix with diagonal blocks A_1, \dots, A_m , we have

$$\begin{aligned} \Phi(x := [x_1; \dots; x_m]) &:= [\Phi_1(x_1); \dots; \Phi_m(x_m)] \\ &= \text{Diag}\{A_1, \dots, A_m\}y(x) + [a_1; \dots; a_m] : \\ y(x) &:= [y_1(x_1); \dots; y_m(x_m)] \in Y := Y_1 \times \dots \times Y_m \text{ and} \\ &\langle \text{Diag}\{A_1^*, \dots, A_m^*\}[x_1; \dots; x_m] - [G_1(y_1(x_1)); \dots; G_m(y_m(x_m))], \\ &\quad y(x) - [y_1; \dots; y_m] \rangle \\ &= \sum_i \langle A_i^*x - G_i(y_i(x_i)), y_i(x_i) - y_i \rangle \geq 0 \quad \forall y = [y_1; \dots; y_m] \in Y, \end{aligned}$$

so that

$$\begin{aligned} F &= F_1 \times \dots \times F_m, \quad A = \text{Diag}\{A_1, \dots, A_m\}, \\ a &= [a_1; \dots; a_m], \\ y(x) &= [y_1(x_1); \dots; y_m(x_m)], \\ G(y = [y_1; \dots; y_m]) &= [G_1(y_1); \dots; G_m(y_m)], \quad Y = Y_1 \times \dots \times Y_m \end{aligned}$$

represent $\Phi : E_1 \times \dots \times E_m \rightarrow E_1 \times \dots \times E_m$. Note that G clearly is good since G_1, \dots, G_m are so.

3.3.5 Representing affine monotone operators

Consider an affine monotone operator on a Euclidean space E :

$$\begin{aligned} \Phi(x) &= Sx + a : E \rightarrow E \\ [S : \langle x, Sx \rangle \geq 0 \forall x \in E] \end{aligned} \tag{10}$$

Its Fenchel-type representation on a convex compact set $X \subset E$ is readily given by the data $F = E, A = S, G(y) = S^*y : F \rightarrow F$ (this operator indeed is monotone), $y(x) = x$ and Y being either the entire F , or (any) compact convex subset of $E = F$ which contains X ; note that G clearly is good w.r.t. A, F , same as is good w.r.t. A, Y when Y is compact. To check that the just defined $F, A, a, y(\cdot), G(\cdot), Y$ indeed represent Φ on X , observe that when $x \in X, y(x) = x$ belongs to $Y \supset X$ and clearly satisfies the relation $0 \leq \langle A^*x - G(y(x)), y(x) - y \rangle \geq 0$ for all $y \in Y$ (see (7)), since $A^*x - G(y(x)) = S^*x - S^*x = 0$. Besides this, for $x \in X$ we have

$$Ay(x) + a = Sx + a = \Phi(x),$$

as required for a representation. The dual antimonotone operator associated with this representation of Φ on X is

$$\Psi(y) = S^*[x(y) - y], \quad x(y) \in \underset{x \in X}{\text{Argmin}} \langle x, Sy + a \rangle. \tag{11}$$

Remark 2 With the above representation of an affine monotone operator, the only restriction on Y , aside from convexity, is to contain X . In particular, when X is compact, Y can definitely be chosen to be compact and “proximal-friendly” – it suffices, e.g., to take as Y an Euclidean ball containing X .

3.3.6 Representing gradient fields

Let $f(x)$ be a convex function given by Fenchel-type representation

$$f(x) = \max_{y \in Y} \{ \langle x, Ay + a \rangle - \psi(y) \}, \tag{12}$$

where Y is a convex compact set in Euclidean space F , and $\psi(\cdot) : F \rightarrow \mathbf{R}$ is a continuously differentiable convex function. Denoting by $y(x)$ a maximizer of $\langle x, Ay \rangle - \psi(y)$ over y , observe that

$$\Phi(x) := Ay(x) + a$$

is a subgradient field of f , and that this monotone operator is given by a representation with the data $F, A, a, y(\cdot), G(\cdot) := \nabla\psi(\cdot), Y$; G is good since Y is compact.

4 Main result

Consider the situation described in Sect. 3.2. Thus, we are given Euclidean space E , a convex compact set $X \subset E$ and a monotone operator $\Phi : E \rightarrow E$ represented according to (7), the data being $F, A, a, y(\cdot), G(\cdot), Y$. We denote by $\Psi : Y \rightarrow F$ the dual (antimonotone) operator induced by the data $X, A, a, y(\cdot), G(\cdot)$, see (9). Our goal is to solve variational inequality given by Φ, X , and our main observation is that *a good accuracy certificate for the variational inequality given by $(-\Psi, Y)$ induces an equally good solution to the variational inequality*

$$\text{find } x_* \in X : \langle \Phi(z), z - x_* \rangle \geq 0 \quad \forall z \in X, \quad \text{VI}(\Phi, X)$$

given by (Φ, X) . The exact statement is as follows.

Theorem 1 *Let $X \subset E$ be a convex compact set and $\Phi : X \rightarrow E$ be a monotone operator represented on X , in the sense of Sect. 3.2.2, by data $F, A, a, y(\cdot), G(\cdot), Y$. Let also $\Psi : Y \rightarrow F$ be the antimonotone operator as defined by the data $X, A, a, y(\cdot), G(\cdot)$, see (9). Let, finally,*

$$C^N = \{y_t, \lambda_t, -\Psi(y_t)\}_{t=1}^N$$

be an accuracy certificate associated with the monotone operator $[-\Psi]$ and Y . Setting

$$x_t = x(y_t)$$

[these points are byproducts of computing $\Psi(y_t)$, $1 \leq t \leq N$, see (9)] and

$$\widehat{x} = \sum_{t=1}^N \lambda_t x_t \ (\in X),$$

we ensure that

$$\epsilon_{\text{vi}}(\widehat{x}|\Phi, X) \leq \text{Res}\left(\mathcal{C}^N|Y(X)\right), \quad Y(X) := \{y(x) : x \in X\} \subset Y. \tag{13}$$

When $\Phi(x) = a + Sx$, $x \in X$, with skew-symmetric S , we have also

$$\text{Res}(\{x_t, \lambda_t, \Phi(x_t)\}_{t=1}^N|X) \leq \text{Res}\left(\mathcal{C}^N|Y(X)\right). \tag{14}$$

In view of Theorem 1, given a representation of the monotone operator Φ participating in the v.i. of interest $\text{VI}(\Phi, X)$, we can reduce solving the v.i. to solving the dual v.i. $\text{VI}(-\Psi, Y)$ by an algorithm producing good accuracy certificates. Below we discuss in details the situation when the latter algorithm is either MD [18, Chapter 5], or MP ([25], see also [18, Chapter 6] and [27]).

Theorem 1 can be extended to the situation where the relations (9), defining the dual operator Ψ hold only approximately. We present here the following slight extension of the main result:

Theorem 2 *Let $X \subset E$ be a convex compact set and $\Phi : X \rightarrow E$ be a monotone operator represented on X , in the sense of Sect. 3.2.2, by data $F, A, a, y(\cdot), G(\cdot), Y$. Given a positive integer N , sequences $y_t \in Y, x_t \in X, 1 \leq t \leq N$, and nonnegative reals $\lambda_t, 1 \leq t \leq N$, summing up to 1, let us set*

$$\begin{aligned} \epsilon = \text{Res}(\{y_t, \lambda_t, G(y_t) - A^*x_t\}_{t=1}^N|Y(X)) &= \sup_{z \in Y(X)} \sum_{t=1}^N \lambda_t \langle G(y_t) - A^*x_t, y_t - z \rangle, \\ [Y(X) := \{y(x) : x \in X\} \subset Y], \end{aligned} \tag{15}$$

and

$$\widehat{x} = \sum_{t=1}^N \lambda_t x_t \ (\in X).$$

Then

$$\epsilon_{\text{vi}}(\widehat{x}|\Phi, X) \leq \epsilon + \sup_{x \in X} \sum_{t=1}^N \lambda_t \langle Ay_t + a, x_t - x \rangle. \tag{16}$$

Proofs of Theorems 1 and 2 are given in Sect. 7.1.

4.1 Mirror Descent and Mirror Prox algorithms

4.1.1 Preliminaries

Saddle Point MD and MP are algorithms for solving convex–concave saddle point problems and variational inequalities with monotone operators.³ The algorithms are of *proximal type*, meaning that in order to apply the algorithm to a v.i. $VI(H, Y)$, where Y is a nonempty closed convex set in Euclidean space E_y and $H(\cdot)$ is a monotone operator on Y , one needs to equip E_y with a norm $\|\cdot\|$, and Y —with a continuously differentiable *distance generating function* (d.-g.f.) $\omega(\cdot) : Y \rightarrow \mathbf{R}$ compatible with $\|\cdot\|$, meaning that ω is strongly convex, modulus 1, w.r.t. $\|\cdot\|$. We call $\|\cdot\|, \omega(\cdot)$ *proximal setup* for Y . This setup gives rise to

- ω -center $y_\omega = \operatorname{argmin}_{y \in Y} \omega(y)$ of Y ,
- *Bregman distance*

$$V_y(z) = \omega(z) - \omega(y) - \langle \omega'(y), z - y \rangle \geq \frac{1}{2} \|z - y\|^2$$

where the concluding inequality is due to strong convexity of ω ,

- ω -size of a nonempty subset $Y' \subset Y$

$$\Omega[Y'] = \sqrt{2 \left[\max_{y' \in Y'} \omega(y') - \min_{y \in Y} \omega(y) \right]}.$$

Due to the definition of y_ω , we have $V_{y_\omega}(y) \leq \frac{1}{2} \Omega^2[Y']$ for all $y \in Y'$, implying that $\|y - y_\omega\| \leq \Omega[Y']$ for all $y \in Y'$.

Given $y \in Y$, the *prox-mapping* with center y is defined as

$$\operatorname{Prox}_y(\xi) = \operatorname{argmin}_{z \in Y} [V_y(z) + \langle \xi, z \rangle] = \operatorname{argmin}_{z \in Y} [\omega(z) + \langle \xi - \omega'(y), z \rangle] : E_y \rightarrow Y.$$

4.1.2 The algorithms

Let Y be a nonempty closed convex set in Euclidean space E_y , $\mathcal{H} = \{H_t(\cdot) : Y \rightarrow E_y\}_{t=1}^\infty$ be a sequence of vector fields, and $\|\cdot\|, \omega(\cdot)$ be a proximal setup for Y . As applied to (\mathcal{H}, Y) , MD is the recurrence

$$\begin{aligned} y_1 &= y_\omega; \\ y_t &\mapsto y_{t+1} = \operatorname{Prox}_{y_t}(\gamma_t H_t(y_t)), \quad t = 1, 2, \dots \end{aligned} \tag{17}$$

MP is the recurrence

$$\begin{aligned} y_1 &= y_\omega; \\ y_t &\mapsto z_t = \operatorname{Prox}_{y_t}(\gamma_t H_t(y_t)) \mapsto y_{t+1} = \operatorname{Prox}_{y_t}(\gamma_t H_t(z_t)), \quad t = 1, 2, \dots \end{aligned} \tag{18}$$

³ MD algorithm originates from [23,24]; its modern proximal form was developed in [1]. MP was proposed in [25]. For the most present exposition of the algorithms, see [18, Chapters 5,6] and [27].

In both MD and MP, $\gamma_t > 0$ are stepsizes. The most important to us properties of these recurrences are as follows.

Proposition 1 For $N = 1, 2, \dots$, consider the accuracy certificate

$$\mathcal{C}^N = \left\{ y_t \in Y, \lambda_t^N := \gamma_t \left[\sum_{\tau=1}^N \gamma_\tau \right]^{-1}, H_t(y_t) \right\}_{t=1}^N,$$

associated with (17). Then for every $Y' \subset Y$ one has

$$\text{Res}(\mathcal{C}^N | Y') \leq \frac{\Omega^2[Y'] + \sum_{t=1}^N \gamma_t^2 \|H_t(y_t)\|_*^2}{2 \sum_{t=1}^N \gamma_t}, \tag{19}$$

where $\|\cdot\|_*$ is the norm conjugate to $\|\cdot\|$: $\|\xi\|_* = \max_{\|x\| \leq 1} \langle \xi, x \rangle$. In particular, if

$$\forall (y \in Y, t) : \|H_t(y)\|_* \leq M \tag{20}$$

with some finite $M \geq 0$, then, given $Y' \subset Y$, N and setting

$$(a) : \gamma_t = \frac{\Omega[Y']}{M\sqrt{N}}, 1 \leq t \leq N, \text{ or } (b) : \gamma_t = \frac{\Omega[Y']}{\|H_t(y_t)\|_*\sqrt{N}}, 1 \leq t \leq N, \tag{21}$$

one has

$$\text{Res}(\mathcal{C}^N | Y') \leq \frac{\Omega[Y']M}{\sqrt{N}}. \tag{22}$$

Proposition 2 For $N = 1, 2, \dots$, consider the accuracy certificate

$$\mathcal{C}^N = \left\{ z_t \in Y, \lambda_t^N := \gamma_t \left[\sum_{\tau=1}^N \gamma_\tau \right]^{-1}, H_t(z_t) \right\}_{t=1}^N,$$

associated with (18). Then, setting

$$d_t = \gamma_t \langle H_t(z_t), z_t - y_{t+1} \rangle - V_{y_t}(y_{t+1}), \tag{23}$$

we have for every $Y' \subset Y$

$$\text{Res}(\mathcal{C}^N | Y') \leq \frac{\frac{1}{2}\Omega^2[Y'] + \sum_{t=1}^N d_t}{\sum_{t=1}^N \gamma_t} \tag{24}$$

$$d_t \leq \frac{1}{2} \left[\gamma_t^2 \|H_t(z_t) - H_t(y_t)\|_*^2 - \|y_t - z_t\|^2 \right], \tag{25}$$

where $\|\cdot\|_*$ is the norm conjugate to $\|\cdot\|$. In particular, if

$$\forall (y, y' \in Y, t) : \|H_t(y) - H_t(y')\|_* \leq L\|y - y'\| + M \tag{26}$$

with some finite $L \geq 0, M \geq 0$, then given $Y' \subset Y, N$ and setting

$$\gamma_t = \frac{1}{\sqrt{2}} \min \left[\frac{1}{L}, \frac{\Omega[Y']}{M\sqrt{N}} \right], 1 \leq t \leq N, \tag{27}$$

one has

$$\text{Res}(\mathcal{C}^N | Y') \leq \frac{1}{\sqrt{2}} \max \left[\frac{\Omega^2[Y']L}{N}, \frac{\Omega[Y']M}{\sqrt{N}} \right]. \tag{28}$$

To make the text self-contained, we provide the proofs of these known results in the appendix.

4.2 Intermediate summary

Theorem 1 combines with Proposition 1 to imply the following claim:

Corollary 2 *In the situation of Theorem 1, let y_1, \dots, y_N be the trajectory of N -step MD as applied to the stationary sequence $\mathcal{H} = \{H_t(\cdot) = -\Psi(\cdot)\}_{t=1}^\infty$ of vector fields, and let $x_t = x(y_t), t = 1, \dots, N$. Then, setting $\lambda_t = \frac{\gamma_t}{\sum_{\tau=1}^N \gamma_\tau}, 1 \leq t \leq N$, we ensure that*

$$\begin{aligned} \epsilon_{\text{vi}} \left(\underbrace{\sum_{t=1}^N \lambda_t x_t}_{\hat{x}} | \Phi, X \right) &\leq \text{Res} \left(\underbrace{\{y_t, \lambda_t, -\Psi(y_t)\}_{t=1}^N}_{\mathcal{C}^N} | Y(X) \right) \\ &\leq \frac{\Omega^2[Y(X)] + \sum_{t=1}^N \gamma_t^2 \|\Psi(y_t)\|_*^2}{2 \sum_{t=1}^N \gamma_t}. \end{aligned} \tag{29}$$

In particular, assuming

$$M = \sup_{y \in Y} \|\Psi(y)\|_*$$

finite and specifying $\gamma_t, 1 \leq t \leq N$, according to (21) with $Y' = Y(X)$, we ensure that

$$\epsilon_{\text{vi}}(\hat{x} | \Phi, X) \leq \text{Res}(\mathcal{C}^N | Y(X)) \leq \frac{\Omega[Y(X)]M}{\sqrt{N}}. \tag{30}$$

When $\Phi(x) = Sx + a, x \in X$, with a skew-symmetric $S, \epsilon_{\text{vi}}(\hat{x} | \Phi, X)$ in the latter relation can be replaced with $\text{Res}(\{x_t, \lambda_t, \Phi(x_t)\}_{t=1}^N | X)$.

In the sequel, we shall refer to the implementation of our approach presented in Corollary 2 as to our *basic* scheme.

5 Modified algorithmic scheme

In this section, we present some modifications of the proposed approach as applied to the case of v.i. VI(Φ, X) with LMO-represented convex compact domain X and monotone operator Φ satisfying some additional conditions to be specified below. While the worst-case complexity bounds for the modified scheme are similar to the ones stated in Corollary 2, there are reasons to believe that in practice the modified scheme could outperform the basic one.

5.1 Situation

In this section, we consider the case when Φ is a monotone operator represented on X by a Fenchel-type representation with data

$$F, A, a, G(\cdot), y(\cdot), Y, \quad (31)$$

where $G(\cdot)$ is Lipschitz continuous. Thus,

1. F is a Euclidean space, $y \mapsto Ay + a$ is an affine mapping from F to E , X is a nonempty compact convex set in E , Y is a closed convex set in F ;
2. $y \mapsto G(y) : Y \rightarrow F$ is a monotone mapping such that

$$\forall (y, y' \in Y) : \|G(y) - G(y')\|_* \leq L\|y - y'\|. \quad (32)$$

for some $L < \infty$;

3. $x \mapsto y(x) : E \rightarrow Y$ is such that $\langle G(y(x)) - A^*x, z - y(x) \rangle \geq 0$ for all $z \in Y$ and all $x \in E$, and $\Phi(x) = Ay(x) + a, x \in X$.

Remark The above assumptions are satisfied in the case of *affine* monotone operator

$$\Phi(x) = Sx + a : E \rightarrow E \quad [S : \langle Sx, x \rangle \geq 0 \forall x \in E]$$

given by the *affine* Fenchel-type representation as follows. Given S, a , we specify an Euclidean space F and linear mappings

$$y \mapsto Ay : F \rightarrow E, \quad y \mapsto Gy : F \rightarrow F, \quad x \mapsto Bx : E \rightarrow F,$$

in such a way that the mapping $y \mapsto Gy$ is monotone and

$$(a) GB = A^*, \quad (b) AB = S \quad (33)$$

(for example, one can set $F = E, A = S, G = S^*$ and define B as the identity mapping; note that the mapping $y \mapsto Gy = S^*y$ is monotone along with Φ). Given just outlined F, A, B, G , let us set

$$y(x) = Bx, G(y) = Gy, Y = F$$

It is immediately seen from (33) that the data $F, A, a, G(y) := Gy, y(x) := Bx, Y$ represent Φ on X , and, of course, $G(\cdot)$ is Lipschitz continuous.⁴

5.2 Hybrid Mirror Prox algorithm

5.2.1 Construction

In the situation of Sect. 5.1 we aim to get an approximate solution to $VI(\Phi, X)$ by applying MP to a properly built sequence $\mathcal{H} = \{H_t(\cdot)\}$ of vector fields on F . Let us fix a proximal setup $\|\cdot\|, \omega(\cdot)$ for Y , and a constant $L = L_G$ satisfying (32). In the sequel, we set

$$\gamma = L_G^{-1}. \tag{34}$$

Instead of choosing $H_t(\cdot) = -\Psi(\cdot)$ as required in the basic scheme, the fields $H_t(\cdot)$ are built recursively, according to the Mirror Prox recurrence

$$\begin{aligned} y_1 &= y_\omega \\ y_t \mapsto x_t \in X &\mapsto H_t(\cdot) = G(\cdot) - A^*x_t \mapsto \\ z_t = \text{Prox}_{y_t}(\gamma H_t(y_t)) &\mapsto y_{t+1} = \text{Prox}_{y_t}(\gamma H_t(z_t)) \end{aligned} \tag{35}$$

with somehow “on-line” chosen x_t . Note that independently of the choice of $x_t \in X$, we have

$$\|H_t(y) - H_t(y')\|_* \leq L\|y - y'\| \quad \forall y, y' \in Y.$$

As a result, (34) and (35) ensure [see (55) and (56)] that

$$\forall z \in Y : \gamma \langle H_t(z_t), z_t - z \rangle \leq V_{y_t}(z) - V_{y_{t+1}}(z). \tag{36}$$

We are about to show that with properly updated x_t 's, (36) yields nice convergence results.

Functions $f_y(\cdot)$. Given $y \in Y$, let us set

$$f_y(x) = \gamma \langle a, x \rangle + \max_{z \in Y} [\langle z, \gamma[A^*x - G(y)] \rangle - V_y(z)] \tag{37}$$

Since $\omega(\cdot)$ is strongly convex on Y , the function $f_y(\cdot)$ is well defined on E ; f_y is convex as the supremum of a family of affine functions of x . Moreover, it is well known that in fact $f_y(\cdot)$ possesses Lipschitz continuous gradient. Specifically, let $\|\cdot\|_E$ be a norm on E , $\|\cdot\|_{E,*}$ be the norm conjugate to $\|\cdot\|_E$, and let L_A be the operator norm of the

⁴ We could also take as Y a convex compact subset of F containing Bx and define $y(x)$ as Bx for $x \in X$ and as (any) strong solution to the $VI(G(\cdot) - A^*x, Y)$ when $x \notin X$.

linear mapping $y \rightarrow Ay : F \rightarrow E$ from the norm $\|\cdot\|$ on F to the norm $\|\cdot\|_{E,*}$ on E , so that

$$\|Ay\|_{E,*} \leq L_A \|y\| \quad \forall y \in F, \quad \|A^*x\|_* \leq L_A \|x\|_E \quad \forall x \in E, \tag{38}$$

or, what is the same,

$$\langle Ay, x \rangle \leq L_A \|y\| \|x\|_E \quad \forall (y \in F, x \in E).$$

Lemma 3 *Function $f_y(\cdot)$ is continuously differentiable with the gradient*

$$\nabla f_y(x) = \gamma A \text{Prox}_y(\gamma[G(y) - A^*x]) + \gamma a, \tag{39}$$

and this gradient is Lipschitz continuous:

$$\|\nabla f_y(x') - \nabla f_y(x'')\|_{E,*} \leq (L_A/L_G)^2 \|x' - x''\|_E \quad \forall x', x'' \in E. \tag{40}$$

For proof, see Sect. 7.4.

Updating x_t 's, preliminaries Observe, first, that when summing up inequalities (36), we get for $Y(X) = \{y(x) : x \in X\} \subset Y$:

$$\begin{aligned} \text{Res} \left(\{z_t, \lambda_t = N^{-1}, H_t(z_t)\}_{t=1}^N | Y(X) \right) &\leq \sup_{z \in Y(X)} \frac{V_{y_1}(z)}{\gamma N} \leq \frac{1}{2\gamma N} \Omega^2[Y(X)] \\ &= \frac{\Omega^2[Y(X)]L_G}{2N} \end{aligned} \tag{41}$$

(recall that $y_1 = y_\omega$). Second, for any $x_t \in X, 1 \leq t \leq N$, we have $\widehat{x} = \frac{1}{N} \sum_{t=1}^N x_t \in X$. Further, invoking (16) with $\lambda_t = N^{-1}, 1 \leq t \leq N$, and z_t in the role of y_t (which by (41) allows to set $\epsilon = \frac{\Omega^2[Y(X)]L_G}{2N}$), we get

$$\begin{aligned} \epsilon_{\text{vi}}(\widehat{x} | \Phi, X) &= \max_{x \in X} \langle \Phi(x), \widehat{x} - x \rangle \\ &\leq \frac{L_G \Omega^2[Y(X)]}{2N} + \max_{x \in X} \frac{1}{N} \sum_{t=1}^N \langle Az_t + a, x_t - x \rangle \\ &= \frac{L_G \Omega^2[Y(X)]}{2N} + \max_{x \in X} \frac{L_G}{N} \sum_{t=1}^N \langle \nabla f_{y_t}(x_t), x_t - x \rangle, \end{aligned} \tag{42}$$

where the last equality is due to (39) along with the fact that $z_t = \text{Prox}_{y_t}(\gamma[G(y_t) - A^*x_t])$ with $\gamma = L_G^{-1}$, see (35). Note that so far our conclusions were independent on how $x_t \in X$ are selected.

Relation (42) implies that when x_t is a minimizer of $f_{y_t}(\cdot)$ on X , we have $\langle \nabla f_{y_t}(x_t), x_t - x \rangle \leq 0$ for all $x \in X$, and with this “ideal” for our purposes choice of x_t , (42) would imply

$$\epsilon_{\text{vi}}(\widehat{x} | \Phi, X) \leq \frac{L_G \Omega^2[Y(X)]}{2N},$$

which is an $O(1/N)$ efficiency estimate, much better than the $O(1/\sqrt{N})$ -efficiency estimate (30).

Updating x_t 's by Conditional Gradient algorithm. Of course, we cannot simply specify x_t as a point from $\text{Argmin}_X f_{y_t}(x)$, since this would require solving precisely at every step of the MP recurrence (35) a large-scale convex optimization problem. What we indeed intend to do, is to solve this problem *approximately*. Specifically, given y_t (so that $f_{y_t}(\cdot)$ is identified), we can apply the classical Conditional Gradient Algorithm (CGA) (which, as was explained in Introduction, is, basically, the only traditional algorithm capable to minimize a smooth convex function over an LMO-represented convex compact set) in order to generate an approximate solution x_t to the problem $\min_X f_{y_t}(x)$ satisfying, for some prescribed $\epsilon > 0$, the relation

$$\delta_t := \max_{x \in X} \langle \nabla f_{y_t}(x_t), x_t - x \rangle \leq \epsilon. \tag{43}$$

By (42), this course of actions implies the efficiency estimate

$$\epsilon_{\text{vi}}(\widehat{x}|\Phi, X) \leq \frac{L_G \Omega^2[Y(X)]}{2N} + L_G \epsilon. \tag{44}$$

5.2.2 Complexity analysis

Let us equip E with a norm $\|\cdot\|_E$, the conjugate norm being $\|\cdot\|_{E,*}$. Let, further, $R_E(X)$ be the radius of the smallest $\|\cdot\|_E$ -ball which contains X . Taking into account (40) and applying the standard results on CGA (see Sect. 7.5), for every $\epsilon > 0$ it takes at most

$$O(1) \left[\frac{L_A^2 R_E^2(X)}{L_G^2 \epsilon} + 1 \right]$$

CGA steps to generate a point x_t with $\delta_t \leq \epsilon$; here and below $O(1)$'s are some absolute constants. Specifying ϵ as $\frac{\Omega^2[Y(X)]}{2N}$, (44) becomes

$$\epsilon_{\text{vi}}(\widehat{x}|\Phi, X) \leq \frac{L_G \Omega^2[Y(X)]}{N},$$

while the computational effort to generate \widehat{x} is dominated by the necessity to generate x_1, \dots, x_N , which amounts to the total of

$$\mathcal{N}(N) = O(1) \left[\frac{L_A^2 R_E^2(X) N^2}{L_G^2 \Omega^2[Y(X)]} + N \right]$$

CGA steps. The computational effort per step is dominated by the necessity to compute the vector $g = \nabla f_y(x)$, given $y \in Y, x \in E$, and to minimize the linear form $\langle g, u \rangle$ over $u \in X$. In particular, for every $\epsilon \in (0, L_G \Omega^2[Y(X)])$, to ensure $\epsilon_{\text{vi}}(\widehat{x}|\Phi, X) \leq \epsilon$, the total number of CGA steps can be upper-bounded by

$$O(1) \left[\left(\frac{L_A R_E(X) \Omega[Y(X)]}{\epsilon} \right)^2 + \frac{L_G \Omega^2[Y(X)]}{\epsilon} \right].$$

We see that in terms of the theoretical upper bound on the number of calls to the LMO for X needed to get an ϵ -solution, our current scheme has no advantages as compared to the MD-based approach analyzed in Corollary 2—in both cases, the bound is proportional to ϵ^{-2} . We, however, may hope that in practice the outlined MP-based scheme can be better than our basic MD-based one, provided that we apply CGA in a “smart” way, for instance, when using CGA with memory, see [15].

6 Illustration

6.1 The problem

We apply our construction to the following problem (“matrix completion with spectral norm fit”):

$$\text{Opt}(P) = \min_{v \in \mathbf{R}^{p_v \times q_v} : \|v\|_{\text{nuc}} \leq 1} [\bar{f}(v) := \|\mathcal{A}v - b\|_{2,2}] \tag{45}$$

where $\mathbf{R}^{p \times q}$ is the space of $p \times q$ real matrices, $\|x\|_{\text{nuc}} = \sum_i \sigma_i(x)$ is the nuclear norm on this space (sum of the singular values $\sigma_i(x)$ of x), $\|x\|_{2,2} = \max_i \sigma_i(x)$ is the spectral norm of x (which is exactly the conjugate of the nuclear norm), and \mathcal{A} is a linear mapping from $\mathbf{R}^{p_v \times q_v}$ to $\mathbf{R}^{p_b \times q_b}$. Note that a more interesting for applications problem (cf. [2,3,20]) would be

$$\text{Opt} = \min_{v \in \mathbf{R}^{p_v \times q_v}} \{ \|v\|_{\text{nuc}} : \|\mathcal{A}v - b\|_{2,2} \leq \delta \};$$

applying the approach from [19], this problem can be reduced to a “small series” of problems (45). We are interested in the “large-scale” case, where the sizes of p_v, q_v of v are large enough to make the full singular value decomposition of a $p_v \times q_v$ matrix prohibitively time consuming, what seemingly rules out the possibility to solve (45) by proximal type First Order algorithms. We assume, at the same time, that computing the leading singular vectors and the leading singular value of a $p_v \times q_v$ or a $p_b \times q_b$ matrix (which, computationally, is by far easier task than finding full singular value decomposition) still can be carried out in reasonable time.

6.1.1 Processing the problem

We rewrite (45) as a bilinear saddle point problem

$$\begin{aligned} \text{Opt}(P) &= \min_{v \in V} \max_{w \in W} \underbrace{\langle w, [\mathcal{A}v - b] \rangle_{\text{Fro}}}_{f(v,w)} \\ V &= \{v \in \mathbf{R}^{p_v \times q_v} : \|v\|_{\text{nuc}} \leq 1\}, \quad W = \{w \in \mathbf{R}^{p_b \times q_b} : \|w\|_{\text{nuc}} \leq 1\} \end{aligned} \tag{46}$$

(from now on $\langle \cdot, \cdot \rangle_{\text{Fro}}$ stands for Frobenius inner product, and $\|\cdot\|_{\text{Fro}}$ – for the Frobenius norm on the space(s) of matrices). The domain X of the problem is the direct product of two unit nuclear norm balls; minimizing a linear form over this domain reduces to minimizing, given ξ and η , the linear forms $\text{Tr}(v\xi^T)$, $\text{Tr}(w\eta^T)$ over $\{v \in \mathbf{R}^{p_v \times q_v} : \|v\|_{\text{nuc}} \leq 1\}$, resp., $\{w \in \mathbf{R}^{p_b \times q_b} : \|w\|_{\text{nuc}} \leq 1\}$, which, in turn, reduces to computing the leading singular vectors and singular values of ξ and η .

The monotone operator associated with (46) is affine and skew-symmetric:

$$\Phi(v, w) = [\nabla_v f(v, w); -\nabla_w f(v, w)] = [\mathcal{A}^*w; -\mathcal{A}v] + [0; b] : \underbrace{\mathbf{R}^{p_v \times q_v} \times \mathbf{R}^{p_b \times q_b}}_E \rightarrow E.$$

From now on we assume that \mathcal{A} is of spectral norm at most 1, i.e.,

$$\|\mathcal{A}v\|_{\text{Fro}} \leq \|v\|_{\text{Fro}}, \quad \forall v$$

(this always can be achieved by scaling).

Representing Φ We can represent the restriction of Φ on X by the data

$$\begin{aligned} F &= \mathbf{R}^{p_v \times q_v} \times \mathbf{R}^{p_b \times q_b} \\ Ay + a &= [\xi; \mathcal{A}\eta + b], \quad y = [\xi; \eta] \in F \quad (\xi \in \mathbf{R}^{p_v \times q_v}, \eta \in \mathbf{R}^{p_b \times q_b}), \\ G(\underbrace{[\xi; \eta]}_y) &= [-\eta; \xi] : F \rightarrow F \\ Y &= \{y = [\xi; \eta] \in F : \|\xi\|_{\text{Fro}} \leq 1, \|\eta\|_{\text{Fro}} \leq 1\} \end{aligned} \tag{47}$$

Indeed, in the notation from Sect. 3.2, for $x = [v; w] \in X = \{[v; w] \in \mathbf{R}^{p_v \times q_v} \times \mathbf{R}^{p_b \times q_b} : \|v\|_{\text{nuc}} \leq 1, \|w\|_{\text{nuc}} \leq 1\}$, the solution $y(x) = [\xi(x); \eta(x)]$ to the linear system $A^*x = G(y)$ is given by $\eta(x) = -v, \xi(x) = \mathcal{A}^*w$, so that both components of $y(x)$ are of Frobenius norm ≤ 1 (recall that spectral norm of \mathcal{A} is ≤ 1), and therefore $y(x) \in Y$. Besides this,

$$Ay(x = [v; w]) + a = [\xi(x); \mathcal{A}\eta(x) + b] = [\mathcal{A}^*w; b - \mathcal{A}v] = \Phi(v, w).$$

We conclude that when $x = [v; w] \in X$, the just defined $y(x)$ meets all requirements from (7), and thus the data $F, A, a, y(\cdot), G(\cdot), Y$ given by (47) indeed represent the monotone operator Φ on X .

The dual operator Ψ given by the data $F, A, a, y(\cdot), G(\cdot), Y$ is

$$\begin{aligned} \Psi(\underbrace{[\xi; \eta]}_y) &= A^*x(y) - G(y) = [v(y) + \eta; \mathcal{A}^*w(y) - \xi], \\ v(y) &\in \underset{\|v\|_{\text{nuc}} \leq 1}{\text{Argmin}} \langle v, \xi \rangle, \quad w(y) \in \underset{\|w\| \leq 1}{\text{Argmin}} \langle w, \mathcal{A}\eta + b \rangle. \end{aligned} \tag{48}$$

Proximal setup We use the Euclidean proximal setup for Y , i.e., we equip the space F embedding Y with the Frobenius norm and take, as the d.-g.f. for Y , the function

$$\omega(\xi, \eta) = \frac{1}{2} \left[\|\xi\|_{\text{Fro}}^2 + \|\eta\|_{\text{Fro}}^2 \right] : F := \mathbf{R}^{p_v \times q_v} \times \mathbf{R}^{p_v \times q_v} \rightarrow \mathbf{R},$$

resulting in $\Omega[Y] = \sqrt{2}$. Furthermore, from (48) and the fact that the spectral norm of \mathcal{A} is bounded by 1 it follows that the monotone operator $\Theta(y) = -\Psi(y) : Y \rightarrow F$ satisfies (20) with $M = 2\sqrt{2}$ and (26) with $L = 0$ and $M = 4\sqrt{2}$.

Remark Theorem 1 combines with Corollary 1 to imply that when converting an accuracy certificate \mathcal{C}^N for the dual v.i. $\text{VI}(-\Psi, Y)$ into a feasible solution \hat{x}^N to the primal v.i. $\text{VI}(\Phi, X)$, we ensure that

$$\epsilon_{\text{sad}}(\hat{x}^N | f, V, W) \leq \text{Res}(\mathcal{C}^N | Y(X)) \leq \text{Res}(\mathcal{C}^N | Y), \tag{49}$$

with f, V, W given by (46). In other words, in the representation $\hat{x}^N = [\hat{v}^N; \hat{w}^N]$, \hat{v}^N is a feasible solution to problem (45) (which is the primal problem associated with (46)), and \hat{w}^N is a feasible solution to the problem

$$\text{Opt}(D) = \max_{w \in W} \min_{v \in V} \langle w, \mathcal{A}v - b \rangle = \max_{w \in W} \left\{ \underline{f}(w) := -\|A^*w\|_{2,2} - \langle b, w \rangle \right\},$$

(which is the dual problem associated with (46)) with the sum of non-optimality, in terms of respective objectives, $\leq \text{Res}(\mathcal{C}^N | Y)$. Computing $\underline{f}(\hat{w})$ (which, together with computing $\bar{f}(\hat{v})$, takes a single call to LMO for X), we get a lower bound on $\text{Opt}(P) = \text{Opt}(D)$ which certifies that $\bar{f}(\hat{v}) - \text{Opt}(P) \leq \text{Res}(\mathcal{C}^N | Y)$.

6.2 Numerical illustration

Here we report on some numerical experiments with problem (45). In these experiments, we used $p_b = q_b =: m$, $p_v = q_v =: n$, with $n = 2m$, and the mapping \mathcal{A} given by

$$\mathcal{A}v = \sum_{i=1}^k \ell_i v r_i^T, \tag{50}$$

with generated at random $m \times n$ factors ℓ_i, r_i scaled to get $\|\mathcal{A}\|_* \approx 1$. In all our experiments, we used $k = 2$. Matrix b in (45) was built as follows: we generated at random $n \times n$ matrix \bar{v} with $\|\bar{v}\|_{\text{nuc}}$ less than (and close to) 1 and $\text{Rank}(\bar{v}) \approx \sqrt{n}$, and took $b = \mathcal{A}\bar{v} + \delta$, with randomly generated $m \times m$ matrix δ of spectral norm about 0.01.

6.2.1 Experiments with the MD-based scheme

Implementing the MD-based scheme In the first series of experiments, the dual v.i. $\text{VI}(-\Psi, Y)$ is solved by the MD algorithm with $N = 512$ steps for all but the largest instance, where $N = 257$ is used. The MD is applied to the stationary sequence $H_t(\cdot) \equiv -\Psi(\cdot)$, $t = 1, 2, \dots$, of vector fields. The stepsizes γ_t are proportional, with

coefficient of order of 1, to those given by (21.b) with $\| \cdot \| \equiv \| \cdot \|_* = \| \cdot \|_{\text{Fro}}$ and $\Omega[Y] = \sqrt{2}$;⁵ the coefficient was tuned empirically in pilot runs on small instances and is never changed afterwards. We also use two straightforward “tricks”:

- Instead of considering one accuracy certificate, $C^N = \{y_t, \lambda_t^N = 1/N, -\Psi(y_t)\}_{t=1}^N$, we build a “bunch” of certificates

$$C_\mu^v = \left\{ y_t, \lambda_t = \frac{1}{v - \mu + 1}, -\Psi(y_t) \right\}_{t=\mu}^v,$$

where μ runs through a grid in $\{1, \dots, N\}$ (in this implementation, a 16-element equidistant grid), and $v \in \{\mu, \mu + 1, \dots, N\}$ runs through another equidistant grid (e.g., for the largest problem instance, the grid $\{1, 9, 17, \dots, 257\}$). We compute the resolutions of these certificates and identify the best (with the smallest resolution) certificate obtained so far. Every 8 steps, the best certificate is used to compute the current approximate solution to (46) along with the saddle point inaccuracy of this solution.

- When applying MD to problem (46), the “dual iterates” $y_t = [\xi_t; \eta_t]$ and the “primal iterates” $x_t := x(y_t) = [v_t; w_t]$ are pairs of matrices, with $n \times n$ matrices ξ_t, η_t, v_t and $m \times m$ matrices w_t (recall that we are in the case of $p_v = q_v = n, p_b = q_b = m$). It is easily seen that with \mathcal{A} given by (50), the matrices ξ_t, η_t, v_t are linear combinations of rank 1 matrices $\alpha_i \beta_i^T, 1 \leq i \leq (k + 1)t$, and w_t are linear combinations of rank 1 matrices $\delta_i \epsilon_i^T, 1 \leq i \leq t$, with on-line computable vectors $\alpha_i, \beta_i, \delta_i, \epsilon_i$. Every step of MD adds $k + 1$ new α - and $k + 1$ new β -vectors, and a pair of new δ - and ϵ -vectors. Our matrix iterates were represented by the vectors of coefficients in the above rank 1 decompositions (let us call this representation *incremental*), so that the computations performed at a step of MD, including computing the leading singular vectors by straightforward power iterations, are as if the standard representations of matrices were used, but all these matrices were of the size (at most) $n \times [(k + 1)N]$, and not $n \times n$ and $m \times m$, as they actually are. In our experiments, for $k = 2$ and $N \leq 512$, this incremental representation of iterates yields meaningful computational savings (e.g., by factor of 6 for $n = 8192$) as compared to the plain representation of iterates by 2D arrays.

Typical results of our preliminary experiments are presented in Table 1. There C^t stands for the best certificate found in course of t steps, and $\text{Gap}(C^t)$ denotes the saddle point inaccuracy of the solution to (46) induced by this certificate (so that $\text{Gap}(C^t)$ is a valid upper bound on the inaccuracy, in terms of the objective, to which the problem of interest (45) was solved in course of t steps). The comments are as follows:

1. The computational results demonstrate “nearly linear”, and not quadratic, growth of running time with m, n ; this is due to the incremental representation of iterates.
2. When evaluating the “convergence patterns” presented in the table, one should keep in mind that we are dealing with a method with slow $O(1/\sqrt{N})$ convergence

⁵ As we have already mentioned, with our proximal setup, the ω -size of Y is $\leq \sqrt{2}$, and (20) is satisfied with $M = 2\sqrt{2}$.

rate, and from this perspective, 50-fold reduction in resolution in 512 steps is not that bad.

3. A natural alternative to the proposed approach would be to solve the saddle point problem (46) “as it is,” by applying to the associated primal v.i. (where the domain is the product of two nuclear norm balls and the operator is Lipschitz continuous and even skew symmetric) a proximal type saddle point algorithm and computing the required prox-mappings via full singular value decompositions. The state-of-the-art MP algorithm when applied to this problem exhibits $O(1/N)$ convergence rate;⁶ yet, every step of this method would require 2 SVD’s of $n \times n$, and 2 SVD’s of $m \times m$ matrices. As applied to the primal v.i., MD exhibits $O(1/\sqrt{N})$ convergence rate, but the steps are cheaper—we need one SVD of $n \times n$, and one SVD of an $m \times m$ matrix, and we are unaware of a proximal type algorithm for the primal v.i. with cheaper iterations. For the sizes m, n, k we are interested in, the computational effort required by the outlined SVD’s is, for all practical purposes, the same as the overall effort per step. Taking into account the actual SVD cpu times on the platform used in our experiments, the overall running times presented in Table 1, i.e., times required by 512 steps of MD as applied to the dual v.i., allow for the following iteration counts N for MP as applied to the primal v.i.:

n	1024	2048	4096	8192
N	406	72	17	4

and for twice larger iteration counts for MD. From our experience, for $n = 1024$ (and perhaps for $n = 2048$ as well), MP algorithm as applied to the primal v.i. would yield solutions of better quality than those obtained with our approach. It, however, would hardly be the case, for both MP and MD, when $n = 4096$, and definitely would not be the case for $n = 8192$. Finally, with $n = 16384$, CPU time used by the 257-step MD as applied to the dual v.i. is hardly enough to complete *just one* iteration of MD as applied to the primal v.i. We believe these data demonstrate that the approach developed in this paper has certain practical potential.

6.2.2 Experiments with the MP-based scheme

In this section we briefly report on the results obtained with the modified MP-based scheme presented in Sect. 5. Same as above, we use the test problems and representation (47) of the monotone operator of interest (with the only difference that now $Y = F$), and the Euclidean proximal setup. Using the Euclidean setup on $Y = F$ makes prox-mappings and functions $f_y(\cdot)$, defined in (37), extremely simple:

$$\text{Prox}_{[\xi; \eta]}([d\xi; d\eta]) = [\xi - d\xi; \eta - d\eta] \quad [\xi, d\xi, \eta, d\eta \in \mathbf{R}^{p_v \times q_v}]$$

$$f_y(x) = \frac{1}{2} \langle y - \gamma[Gy - A^*x], y - \gamma[Gy - A^*x] \rangle + \gamma \langle a, x \rangle$$

⁶ For the primal v.i., (26) holds true for some $L > 0$ and $M = 0$. Moreover, with properly selected proximal setup for (45) the complexity bound (28) becomes $\text{Res}(C^N|Y) \leq O(1)\sqrt{\ln(n)\ln(m)}/N$.

Table 1 MD on problem (45)

	Iteration count t									
	1	65	129	193	257	321	385	449	512	
$n = 1024 \ m = 512 \ k = 2$										
Res($C^t Y$)	1.5402	0.1535	0.0886	0.0621	0.0487	0.0389	0.03288	0.0293	0.0278	
Res($C^{\perp} Y$)/Res($C^t Y$)	1.00	10.04	17.38	24.79	31.61	39.61	46.84	52.64	55.41	
Gap(C^t)	0.1269	0.0239	0.0145	0.0103	0.0075	0.0063	0.0042	0.0040	0.0040	
Gap(C^{\perp})/Gap(C^t)	1.00	5.31	8.78	12.38	17.03	20.20	29.98	31.41	31.66	
cpu, sec	0.2	9.5	27.6	69.1	112.6	218.1	326.2	432.6	536.4	
$n = 2048 \ m = 1024 \ k = 2$										
Res($C^t Y$)	1.4809	0.1559	0.0842	0.0607	0.0471	0.0391	0.0337	0.0306	0.0285	
Res($C^{\perp} Y$)/Res($C^t Y$)	1.00	9.50	17.59	24.38	31.43	37.88	43.89	48.36	51.96	
Gap(C^t)	0.1329	0.0196	0.0119	0.0075	0.0053	0.0041	0.0036	0.0034	0.0027	
Gap(C^{\perp})/Gap(C^t)	1.00	6.79	11.21	17.81	25.09	32.29	37.23	38.70	50.06	
cpu, sec	0.7	38.0	101.1	206.3	314.1	508.9	699.0	884.9	1070.0	
$n = 4096 \ m = 2048 \ k = 2$										
Res($C^t Y$)	1.4845	0.1476	0.0891	0.0605	0.0491	0.0395	0.0329	0.0292	0.0275	
Res($C^{\perp} Y$)/Res($C^t Y$)	1.00	10.06	16.66	24.53	30.25	37.60	45.17	50.85	53.95	
Gap(C^t)	0.1239	0.0222	0.0139	0.0108	0.0086	0.0041	0.0037	0.0035	0.0035	
Gap(C^{\perp})/Gap(C^t)	1.00	5.57	8.93	11.48	14.40	30.48	33.14	35.76	35.77	
cpu, sec	2.2	103.5	257.6	496.9	742.5	1147.8	1564.4	1981.4	2401.0	

Table 1 continued

		Iteration count t									
		1	65	129	193	257	321	385	449	512	
$n = 8192$ $m = 4096$ $k = 2$											
Res($\mathcal{C}^t Y$)	1.4778	0.1391	0.0888	0.0590	0.0469	0.0386	0.0324	0.0289	0.0270		
Res($\mathcal{C}^\perp Y$)/Res($\mathcal{C}^t Y$)	1.00	10.63	16.64	25.06	31.53	38.29	45.68	51.10	54.76		
Gap(\mathcal{C}^t)	0.1193	0.0232	0.0134	0.0108	0.0054	0.0040	0.0035	0.0034	0.0034		
Gap(\mathcal{C}^t)/Gap(\mathcal{C}^t)	1.00	5.14	8.90	11.08	22.00	29.83	33.93	34.85	35.14		
cpu, sec	6.5	289.9	683.8	1238.1	1816.0	2724.5	3648.3	4572.2	5490.8		
$n = 16384$ $m = 8192$ $k = 2$											
Res(\mathcal{C}^t)	1.4566	0.1154	0.0767	0.0556	0.0447						
Res($\mathcal{C}^\perp Y$)/Res($\mathcal{C}^t Y$)	1.00	12.62	19.00	26.22	32.60						
Gap(\mathcal{C}^t)	0.11959	0.02136	0.01460	0.01011	0.00853						
Gap(\mathcal{C}^t)/Gap(\mathcal{C}^t)	1.00	5.60	8.19	11.82	14.01						
cpu, sec	21.7	920.4	2050.2	3492.4	4902.2						

Platform: 3.40 GHz i7-3770 desktop with 16 GB RAM, MATLAB Windows 7 64

$$\begin{aligned}
 &= \frac{1}{2} \left[\|\xi + \gamma\eta + \gamma\mathcal{A}^*w\|_{\text{Fro}}^2 + \|\eta - \gamma\xi + \gamma v\|_{\text{Fro}}^2 \right] + \gamma \langle b, w \rangle_{\text{Fro}} \\
 y &= [\xi; \eta], x = [v; w].
 \end{aligned}$$

When choosing $\|\cdot\|_E$ to be the Frobenius norm,

$$\underbrace{\| [v; w] \|_E}_x = \|x\|_{E,*} = \sqrt{\|v\|_{\text{Fro}}^2 + \|w\|_{\text{Fro}}^2}$$

and taking into account that the spectral norm of \mathcal{A} is ≤ 1 , it is immediately seen that the quantities L_G, L_A and γ introduced in Sect. 5.2, can be set to 1, and what was called $R_E(X)$ in Sect. 5.2.2, can be set to $\sqrt{2}$. As a result, by the complexity analysis of Sect. 5.2.2, in order to find an ϵ -solution to the problem of interest, we need $O(1)\epsilon^{-1}$ iterations of the recurrence (35), with $O(1)\epsilon^{-1}$ CGA steps of minimizing $f_{y_t}(\cdot)$ over X per iteration, that is, the total of at most $O(1)\epsilon^{-2}$ calls to the LMO for $X = V \times W$. In fact, in our implementation ϵ is not fixed in advance; instead, we fix the total number $N = 256$ of calls to LMO, and terminate CGA at iteration t of the recurrence (35) when either a solution $x_t \in x$ with $\delta_t \leq 0.1/t$ is achieved, or the number of CGA steps reaches a prescribed limit (set to 32 in the experiment to be reported).

Same as in the first series of experiments, “incremental” representation of matrix iterates is used in the experiments with the MP-based scheme. In these experiments we also use a special *post-processing* of the solution we explain next.

Post-processing Recall that in the situation in question the step $\#i$ of the CGA at iteration $\#t$ of the MP-based recurrence produces a pair $[v_{t,i}; w_{t,i}]$ of rank 1 $n \times n$ and $m \times m$ matrices of unit spectral norm—the minimizers of the linear form $\langle \nabla f_{y_t}(x_{t,i}), x \rangle$ over $x \in X$; here $x_{t,i}$ is i -th step of CGA minimization of $f_{y_t}(\cdot)$ over X . As a result, upon termination, we have at our disposal $N = 256$ pairs of rank one matrices $[v_j; w_j]$, $1 \leq j \leq N$, known to belong to X . Note that the approximate solution \hat{x} , as defined in Sect. 5.2.1, is a certain convex combination of these matrices. A natural way to get a better solution is to solve the optimization problem

$$\text{Opt} = \min_{\lambda, v} \left\{ f(\lambda) = \|Av - b\|_{2,2} : v = \sum_{j=1}^N \lambda_j v_j, \sum_{j=1}^N |\lambda_j| \leq 1 \right\}. \quad (51)$$

Indeed, note that the v -components of feasible solutions to this problem are of nuclear norm ≤ 1 , i.e., are feasible solutions to the problem of interest (45), and that in terms of the objective of (45), the v -component of an optimal solution to (51) can be only better than the v -component of \hat{x} . On the other hand, (51) is a low-dimensional convex optimization problem on a simple domain, and the first order information on f can be obtained, at a relatively low cost, by Power Method, so that (51) is well suited for solving by proximal first order algorithms, e.g., the Bundle Level algorithm [21] we use in our experiments.

Numerical illustration Here we present just one (in fact, quite representative) numerical example. In this example $n = 4096$ and $m = 2048$ (i.e., in (45) the variable matrix u is of size 4096×4096 , and the data matrix b is of size 2048×2048); the mapping \mathcal{A} is given by (50) with $k = 2$. The data are generated in the same way as in the experiments described in Sect. 6.2.1 except for the fact that we used $b = \mathcal{A}\bar{u}$ to

ensure zero optimal value in (45). As a result, the value of the objective of (45) at an approximate solution coincides with the inaccuracy of this solution in terms of the objective of (45). In the experiment we report on here, the objective of (45) evaluated at the initial – zero – solution, i.e., $\|b\|_{2,2}$, is equal to 0.751. After the total of 256 calls to the LMO for X [just 11 steps of recurrence (35)] and post-processing which took 24% of the overall CPU time, the value of the objective is reduced to 0.013—by factor 57.3. For comparison, when processing the same instance by the basic MD scheme, augmented by the just outlined post-processing, after 256 MD iterations (i.e., after the same as above 256 calls to the LMO), the value of the objective at the resulting feasible solution to (45) was 0.071, meaning the progress in accuracy by factor 10.6 (5 times worse than the progress in accuracy for the MP-based scheme). Keeping the instance intact and increasing the number of MD iterations in the basic scheme from 256 to 512, the objective at the approximate solution yielded by the post-processing reduces from 0.071 to 0.047, which still is 3.6 times worse than that achieved with the MP-based scheme after 256 calls to LMO.

7 Proofs

7.1 Proof of Theorems 1 and 2

We start with proving Theorem 2. In the notation of the theorem, we have

$$\begin{aligned} & \forall x \in X : \Phi(x) = Ay(x) + a, \\ (a) : & y(x) \in Y, \\ (b) : & \langle y(x) - y, A^*x - G(y(x)) \rangle \geq 0 \forall y \in Y. \end{aligned} \tag{52}$$

For $\bar{x} \in X$, let $\bar{y} = y(\bar{x})$, and let $\hat{y} = \sum_t \lambda_t y_t$, so that $\bar{y}, \hat{y} \in Y$ by (52.a). Since G is monotone, for all $t \in \{1, \dots, N\}$ we have

$$\begin{aligned} & \langle \bar{y} - y_t, G(\bar{y}) - G(y_t) \rangle \geq 0 \\ \Rightarrow & \langle \bar{y}, G(\bar{y}) \rangle \geq \langle y_t, G(\bar{y}) \rangle + \langle \bar{y}, G(y_t) \rangle - \langle y_t, G(y_t) \rangle \quad \forall t \\ \Rightarrow & \langle \bar{y}, G(\bar{y}) \rangle \geq \sum_t \lambda_t [\langle y_t, G(\bar{y}) \rangle + \langle \bar{y}, G(y_t) \rangle - \langle y_t, G(y_t) \rangle] \\ & [\text{since } \lambda_t \geq 0 \text{ and } \sum_t \lambda_t = 1], \end{aligned}$$

and we conclude that

$$\langle \bar{y}, G(\bar{y}) \rangle - \langle \hat{y}, G(\bar{y}) \rangle \geq \sum_{t=1}^N \lambda_t [\langle \bar{y}, G(y_t) \rangle - \langle y_t, G(y_t) \rangle]. \tag{53}$$

We now have

$$\begin{aligned} & \langle \Phi(\bar{x}), \bar{x} - \sum_t \lambda_t x_t \rangle \\ & = \langle A\bar{y} + a, \bar{x} - \sum_t \lambda_t x_t \rangle = \langle \bar{y}, A^*\bar{x} - \sum_t \lambda_t A^*x_t \rangle + \langle a, \bar{x} - \sum_t \lambda_t x_t \rangle \\ & = \langle \bar{y}, A^*\bar{x} - G(\bar{y}) \rangle + \langle \bar{y}, G(\bar{y}) - \sum_t \lambda_t A^*x_t \rangle + \langle a, \bar{x} - \sum_t \lambda_t x_t \rangle \end{aligned}$$

$$\begin{aligned}
 &\geq \langle \widehat{y}, A^* \bar{x} - G(\bar{y}) \rangle + \langle \bar{y}, G(\bar{y}) - \sum_t \lambda_t A^* x_t \rangle + \langle a, \bar{x} - \sum_t \lambda_t x_t \rangle \\
 &\quad [\text{by (52.b) with } y = \widehat{y} \text{ and due to } \bar{y} = y(\bar{x})] \\
 &= \langle \widehat{y}, A^* \bar{x} \rangle + [\langle G(\bar{y}), \bar{y} \rangle - \langle G(\bar{y}), \widehat{y} \rangle] - \langle \bar{y}, \sum_t \lambda_t A^* x_t \rangle + \langle a, \bar{x} - \sum_t \lambda_t x_t \rangle \\
 &\geq \langle \widehat{y}, A^* \bar{x} \rangle + \sum_t \lambda_t [\langle \bar{y}, G(y_t) \rangle - \langle y_t, G(y_t) \rangle] - \langle \bar{y}, \sum_t \lambda_t A^* x_t \rangle \\
 &\quad + \langle a, \bar{x} - \sum_t \lambda_t x_t \rangle \quad [\text{by (53)}] \\
 &= \sum_t \lambda_t \langle y_t, A^* \bar{x} \rangle + \sum_t \lambda_t [\langle \bar{y}, G(y_t) \rangle - \langle y_t, G(y_t) \rangle - \langle \bar{y}, A^* x_t \rangle + \langle a, \bar{x} - x_t \rangle] \\
 &\quad [\text{since } \widehat{y} = \sum_t \lambda_t y_t \text{ and } \sum_t \lambda_t = 1] \\
 &= \sum_t \lambda_t [\langle Ay_t, \bar{x} - x_t \rangle + \langle Ay_t, x_t \rangle + \langle \bar{y}, G(y_t) \rangle - \langle y_t, G(y_t) \rangle \\
 &\quad - \langle \bar{y}, A^* x_t \rangle + \langle a, \bar{x} - x_t \rangle] \\
 &= \sum_t \lambda_t [\langle y_t, A^* x_t \rangle + \langle \bar{y}, G(y_t) \rangle - \langle y_t, G(y_t) \rangle - \langle \bar{y}, A^* x_t \rangle] \\
 &\quad + \sum_t \lambda_t \langle Ay_t + a, \bar{x} - x_t \rangle \\
 &= \sum_t \lambda_t \langle A^* x_t - G(y_t), y_t - \bar{y} \rangle + \sum_t \lambda_t \langle Ay_t + a, \bar{x} - x_t \rangle \geq -\epsilon \\
 &\quad + \sum_t \lambda_t \langle Ay_t + a, \bar{x} - x_t \rangle \\
 &\quad [\text{by (15) due to } \bar{y} = y(\bar{x}) \in Y(X)].
 \end{aligned}$$

The bottom line is that

$$\langle \Phi(\bar{x}), \widehat{x} - \bar{x} \rangle \leq \epsilon + \sum_{t=1}^N \lambda_t \langle Ay_t + a, x_t - \bar{x} \rangle \forall \bar{x} \in X,$$

as stated in (16). Theorem 2 is proved.

To prove Theorem 1, let $y_t \in Y, 1 \leq t \leq N$, and $\lambda_1, \dots, \lambda_N$ be from the premise of the theorem, and let $x_t, 1 \leq t \leq N$, be specified as $x_t = x(y_t)$, so that x_t is the minimizer of the linear form $\langle Ay_t + a, x \rangle$ over $x \in X$. Due to the latter choice, we have $\sum_{t=1}^N \lambda_t \langle Ay_t + a, x_t - \bar{x} \rangle \leq 0$ for all $\bar{x} \in X$, while ϵ as defined by (15) is nothing but $\text{Res}(\{y_t, \lambda_t, -\Psi(x_t)\}_{t=1}^N | Y(X))$. Thus, (16) in the case in question implies that

$$\forall \bar{x} \in X : \langle \Phi(\bar{x}), \sum_{t=1}^N \lambda_t x_t - \bar{x} \rangle \leq \text{Res}(\{y_t, \lambda_t, -\Psi(x_t)\}_{t=1}^N | Y(X)),$$

and (13) follows. Relation (14) is an immediate corollary of (13) and Lemma 2 as applied to X in the role of Y , Φ in the role of $H(\cdot)$, and $\{x_t, \lambda_t, \Phi(x_t)\}_{t=1}^N$ in the role of \mathcal{C}^N . □

7.2 Proof of Proposition 1

Observe that the optimality conditions in the optimization problem specifying $v = \text{Prox}_y(\xi)$ imply that

$$\langle \xi - \omega'(y) + \omega'(v), z - v \rangle \geq 0, \quad \forall z \in Y,$$

or

$$\langle \xi, v - z \rangle \leq \langle \omega'(v) - \omega'(y), z - v \rangle = \langle V'_y(v), z - v \rangle, \quad \forall z \in Y,$$

which, using a remarkable identity [5]

$$\langle V'_y(v), z - v \rangle = V_y(z) - V_v(z) - V_y(v),$$

can be rewritten equivalently as

$$v = \text{Prox}_y(\xi) \Rightarrow \langle \xi, v - z \rangle \leq V_y(z) - V_v(z) - V_y(v) \quad \forall z \in Y. \tag{54}$$

Setting $y = y_t$, $\xi = \gamma_t H_t(y_t)$, which results in $v = y_{t+1}$, we get

$$\forall z \in Y : \gamma_t \langle H_t(y_t), y_{t+1} - z \rangle \leq V_{y_t}(z) - V_{y_{t+1}}(z) - V_{y_t}(y_{t+1}),$$

whence,

$$\begin{aligned} \forall z \in Y : \gamma_t \langle H_t(y_t), y_t - z \rangle &\leq V_{y_t}(z) - V_{y_{t+1}}(z) + \underbrace{\left[\gamma_t \langle H_t(y_t), y_t - y_{t+1} \rangle - V_{y_t}(y_{t+1}) \right]}_{\leq \gamma_t \|H_t(y_t)\|_* \|y_t - y_{t+1}\| - \frac{1}{2} \|y_t - y_{t+1}\|^2} \\ &\leq V_{y_t}(z) - V_{y_{t+1}}(z) + \frac{1}{2} \gamma_t^2 \|H_t(y_t)\|_*^2. \end{aligned}$$

Summing up these inequalities over $t = 1, \dots, N$ and taking into account that for $z \in Y'$, we have $V_{y_1}(z) \leq \frac{1}{2} \Omega^2 [Y']$ and that $V_{y_{N+1}}(z) \geq 0$, we get (19). \square

7.3 Proof of Proposition 2

Applying (54) to $y = y_t$, $\xi = \gamma_t H_t(z_t)$, which results in $v = y_{t+1}$, we get

$$\forall z \in Y : \gamma_t \langle H_t(z_t), y_{t+1} - z \rangle \leq V_{y_t}(z) - V_{y_{t+1}}(z) - V_{y_t}(y_{t+1}),$$

whence, by the definition (23) of d_t ,

$$\forall z \in Y : \gamma_t \langle H_t(z_t), z_t - z \rangle \leq V_{y_t}(z) - V_{y_{t+1}}(z) + d_t. \tag{55}$$

Summing up the resulting inequalities over $t = 1, \dots, N$ and taking into account that $V_{y_1}(z) \leq \frac{1}{2}\Omega^2[Y']$ for all $z \in Y'$ and $V_{y_{N+1}}(z) \geq 0$, we get

$$\forall z \in Y' : \sum_{t=1}^n \lambda_t^N \langle H_t(z_t), z_t - z \rangle \leq \frac{\frac{1}{2}\Omega^2[Y'] + \sum_{t=1}^N d_t}{\sum_{t=1}^N \gamma_t}.$$

The right hand side in the latter inequality is independent of $z \in Y'$. Taking supremum of the left hand side over $z \in Y'$, we arrive at (24).

Moreover, invoking (54) with $y = y_t, \xi = \gamma_t H_t(y_t)$ and specifying z as y_{t+1} , we get

$$\gamma_t \langle H_t(y_t), z_t - y_{t+1} \rangle \leq V_{y_t}(y_{t+1}) - V_{z_t}(y_{t+1}) - V_{y_t}(z_t),$$

whence

$$\begin{aligned} d_t &= \gamma_t \langle H_t(z_t), z_t - y_{t+1} \rangle - V_{y_t}(y_{t+1}) \leq \gamma_t \langle H_t(y_t), z_t - y_{t+1} \rangle \\ &\quad + \gamma_t \langle H_t(z_t) - H_t(y_t), z_t - y_{t+1} \rangle - V_{y_t}(y_{t+1}) \\ &\leq -V_{z_t}(y_{t+1}) - V_{y_t}(z_t) + \gamma_t \langle H_t(z_t) - H_t(y_t), z_t - y_{t+1} \rangle \\ &\leq \gamma_t \|H_t(z_t) - H_t(y_t)\|_* \|z_t - y_{t+1}\| - \frac{1}{2} \|z_t - y_{t+1}\|^2 - \frac{1}{2} \|y_t - z_t\|^2 \\ &\leq \frac{1}{2} \left[\gamma_t^2 \|H_t(z_t) - H_t(y_t)\|_*^2 - \|y_t - z_t\|^2 \right], \end{aligned} \tag{56}$$

as required in (25). □

7.4 Proof of Lemma 3

1⁰. We start with the following standard fact:

Lemma 4 *Let Y be a nonempty closed convex set in Euclidean space F , $\|\cdot\|$ be a norm on F , and $\omega(\cdot)$ be a continuously differentiable function on Y which is strongly convex, modulus 1 , w.r.t. $\|\cdot\|$. Given $b \in F$ and $y \in Y$, let us set*

$$\begin{aligned} g_y(\xi) &= \max_{z \in Y} [\langle z, \omega'(y) - \xi \rangle - \omega(z)] : F \rightarrow \mathbf{R}, \\ \text{Prox}_y(\xi) &= \underset{z \in Y}{\text{argmax}} [\langle z, \omega'(y) - \xi \rangle - \omega(z)]. \end{aligned}$$

The function g_y is convex with Lipschitz continuous gradient $\nabla g_y(\xi) = -\text{Prox}_y(\xi)$:

$$\|\nabla g_y(\xi) - \nabla g_y(\xi')\| \leq \|\xi - \xi'\|_* \quad \forall \xi, \xi', \tag{57}$$

where $\|\cdot\|_*$ is the norm conjugate to $\|\cdot\|$.

Indeed, since ω is strongly convex and continuously differentiable on Y , $\text{Prox}_Y(\cdot)$ is well defined, and from optimality conditions it holds

$$\langle \omega'(\text{Prox}_Y(\xi)) + \xi - \omega'(y), \text{Prox}_Y(\xi) - z \rangle \leq 0 \quad \forall z \in Y. \tag{58}$$

Consequently, $g_Y(\cdot)$ is well defined; this function clearly is convex, and the vector $-\text{Prox}_Y(\xi)$ clearly is a subgradient of g_Y at ξ . If now $\xi', \xi'' \in F$, then, setting $z' = \text{Prox}_Y(\xi')$, $z'' = \text{Prox}_Y(\xi'')$ and invoking (58), we get

$$\langle \omega'(z') + \xi' - \omega'(y), z' - z'' \rangle \leq 0, \quad \langle -\omega'(z'') - \xi'' + \omega'(y), z' - z'' \rangle \leq 0$$

whence, summing the inequalities up,

$$\langle \xi' - \xi'', z' - z'' \rangle \leq \langle \omega'(z') - \omega'(z''), z'' - z' \rangle \leq -\|z' - z''\|^2,$$

implying that $\|z' - z''\| \leq \|\xi' - \xi''\|_*$. Thus, a subgradient field $-\text{Prox}_Y(\cdot)$ of $g_Y(\cdot)$ is Lipschitz continuous with constant 1 from $\|\cdot\|_*$ into $\|\cdot\|$, whence g_Y is continuously differentiable and (57) takes place. \square

2⁰. To derive Lemma 3 from Lemma 4, note that $f_Y(x)$ is obtained from $g_Y(\cdot)$ by affine substitution of variables and adding linear form:

$$f_Y(x) = g_Y(\gamma[G(y) - A^*x]) + \gamma \langle a, x \rangle + \omega(y) - \langle \omega'(y), y \rangle.$$

whence $\nabla f_Y(x) = -\gamma A \nabla g_Y(\gamma[G(y) - A^*x]) + \gamma a = \gamma A \text{Prox}_Y(\gamma[G(y) - A^*x]) + \gamma a$, as required in (39), and

$$\begin{aligned} & \|\nabla f_Y(x') - \nabla f_Y(x'')\|_{E,*} \\ &= \gamma \|A [\nabla g_Y(\gamma[G(y) - A^*x']) - \nabla g_Y(\gamma[G(y) - A^*x''])]\|_{E,*} \\ &\leq (\gamma L_A) \|\nabla g_Y(\gamma[G(y) - A^*x']) - \nabla g_Y(\gamma[G(y) - A^*x''])\| \\ &\leq (\gamma L_A) \|\gamma[G(y) - A^*x'] - \gamma[G(y) - A^*x'']\|_* \\ &\leq (\gamma L_A)^2 \|x' - x''\|_E = (L_A/L_G)^2 \|x' - x''\|_E \end{aligned}$$

[we have used (57) and equivalences in (38)], as required in (40). \square

7.5 Review of Conditional Gradient algorithm

The required description of CGA and its complexity analysis are as follows.

As applied to minimizing a smooth – with Lipschitz continuous gradient

$$\|\nabla f(u) - \nabla f(u')\|_{E,*} \leq \mathcal{L} \|u - u'\|_E, \quad \forall u, u' \in X,$$

convex function f over a convex compact set $X \subset E$, the generic CGA is the recurrence of the form

$$\begin{aligned} u_1 &\in X \\ u_{s+1} &\in X \text{ satisfies } f(u_{s+1}) \leq f(u_s + \gamma_s[u_s^+ - u_s]), \quad s = 1, 2, \dots \\ \gamma_s &= \frac{2}{s+1}, \quad u_s^+ \in \operatorname{Argmin}_{u \in X} \langle f'(u_s), u \rangle. \end{aligned}$$

The standard results on this recurrence (see, e.g., proof of Theorem 1 in [15]) state that if $f_* = \min_X f$, then

$$\begin{aligned} (a) \quad \epsilon_{t+1} &:= f(u_{t+1}) - f_* \leq \epsilon_t - \gamma_t \delta_t + 2\mathcal{L}R^2 \gamma_t^2, \quad t = 1, 2, \dots \\ \delta_t &:= \max_{u \in X} \langle \nabla f(u_t), u_t - u \rangle; \\ (b) \quad \epsilon_t &\leq \frac{2\mathcal{L}R^2}{t+1}, \quad t = 2, 3, \dots \end{aligned} \tag{59}$$

where R is the smallest of the radii of $\|\cdot\|_E$ -balls containing X . From (59.a) it follows that

$$\gamma_\tau \delta_\tau \leq \epsilon_\tau - \epsilon_{\tau+1} + 2\mathcal{L}R^2 \gamma_\tau^2, \quad \tau = 1, 2, \dots;$$

summing up these inequalities over $\tau = t, t + 1, \dots, 2t$, where $t > 1$, we get

$$\left[\min_{\tau \leq 2t} \delta_\tau \right] \sum_{\tau=t}^{2t} \gamma_\tau \leq \epsilon_t + 2\mathcal{L}R^2 \sum_{\tau=t}^{2t} \gamma_\tau^2,$$

which combines with (59.b) to imply that

$$\min_{\tau \leq 2t} \delta_\tau \leq O(1)\mathcal{L}R^2 \frac{\frac{1}{t} + \sum_{\tau=t}^{2t} \frac{1}{\tau^2}}{\sum_{\tau=t}^{2t} \frac{1}{\tau}} \leq O(1) \frac{\mathcal{L}R^2}{t}.$$

It follows that given $\epsilon < \mathcal{L}R^2$, it takes at most $O(1) \frac{\mathcal{L}R^2}{\epsilon}$ steps of CGA to generate a point $u^\epsilon \in X$ with $\max_{u \in X} \langle \nabla f(u^\epsilon), u^\epsilon - u \rangle \leq \epsilon$.

References

1. Beck, A., Teboulle, M.: Mirror descent and nonlinear projected subgradient methods for convex optimization. *Oper. Res. Lett.* **31**(3), 167–175 (2003)
2. Candes, E.J., Plan, Y.: Matrix completion with noise. *Proc. IEEE* **98**(6), 925–936 (2010)
3. Candes, E.J., Plan, Y.: Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements. *Inf. Theory IEEE Trans.* **57**(4), 2342–2359 (2011)
4. Castellani, M., Mastroeni, G.: On the duality theory for finite dimensional variational inequalities. In: Giannessi, F., Maugeri, A. (eds.) *Variational Inequalities Network Equilibrium Problems*, pp. 21–31. Plenum Publishing, New York (1995)
5. Chen, G., Teboulle, M.: Convergence analysis of a proximal-like minimization algorithm using Bregman functions. *SIAM. J. Optim.* **3**(3), 538–543 (1993)
6. Chipot, M.: *Variational Inequalities and Flow in Porous Media*. Springer, New York (1984)
7. Cox, B., Juditsky, A., Nemirovski, A.: Dual subgradient algorithms for large-scale nonsmooth learning problems. *Math. Program. Ser. B* **148**(1–2), 143–180 (2014)

8. Daniele, P.: Dual variational inequality and applications to asymmetric traffic equilibrium problems with capacity constraints. *Le Matematiche* **XLIX**(2), 111–211 (1994)
9. Demyanov, V., Rubinov, A.: *Approximate Methods in Optimization Problems*. Elsevier, Amsterdam (1970)
10. Dunn, J.C., Harshbarger, S.: Conditional gradient algorithms with open loop step size rules. *J. Math. Anal. Appl.* **62**(2), 432–444 (1978)
11. Dudik, M., Harchaoui, Z., Malick, J.: Lifted coordinate descent for learning with trace-norm regularization. In: *AISTATS-Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics*, vol. 22, pp. 327–336 (2012)
12. Frank, M., Wolfe, P.: An algorithm for quadratic programming. *Naval Res. Logist. Q.* **3**(1–2), 95–110 (1956)
13. Freund, R., Grigasy, P.: New analysis and results for the Conditional Gradient method (2013) submitted to *Mathematical Programming*, E-print: <http://web.mit.edu/rfreund/www/FW-paper-final>
14. Harchaoui, Z., Douze, M., Paulin, M., Dudik, M., Malick, J.: Large-scale image classification with trace-norm regularization. In: *2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3386–3393 (2012)
15. Harchaoui, Z., Juditsky, A., Nemirovski, A.: Conditional Gradient algorithms for norm-regularized smooth convex optimization. *Mathematical Programming*. Online First (2014). doi:[10.1007/s10107-014-0778-9](https://doi.org/10.1007/s10107-014-0778-9). E-print: [arXiv:1302.2325](https://arxiv.org/abs/1302.2325)
16. Jaggi, M.: Revisiting Frank–Wolfe: projection-free sparse convex optimization. In: *Proceedings of the 30th International Conference on Machine Learning (ICML 2013)*, pp. 427–435 (2013)
17. Jaggi, M., Sulovsky, M.: A simple algorithm for nuclear norm regularized problems. In: *Proceedings of the 27th International Conference on Machine Learning (ICML 2010)*, pp. 471–478 (2010)
18. Juditsky, A., Nemirovski, A.: First order methods for nonsmooth large-scale convex minimization, I: general purpose methods; II: utilizing problem’s structure. In: Sra, S., Nowozin, S., Wright, S. (eds.) *Optimization for Machine Learning*, pp. 121–184. The MIT Press, Cambridge (2012)
19. Juditsky, A., Kiliç Karzan, F., Nemirovski, A.: Randomized first order algorithms with applications to ℓ_1 -minimization. *Math. Program. Ser. A* **142**(1–2), 269–310 (2013)
20. Juditsky, A., Kiliç Karzan, F., Nemirovski, A.: On unified view of nullspace-type conditions for recoveries associated with general sparsity structures. *Linear Algebra Appl.* **441**(1), 124–151 (2014)
21. Lemarechal, C., Nemirovski, A., Nesterov, Yu.: New variants of bundle methods. *Math. Program.* **69**(1), 111–148 (1995)
22. Mosco, U.: Dual variational inequalities. *J. Math. Anal. Appl.* **40**, 202–206 (1972)
23. Nemirovski, A., Yudin, D.: *Problem Complexity and Method Efficiency in Optimization*. Nauka Publishers, Moscow (1978) (in Russian); John Wiley, New York (1983) (in English)
24. Nemirovskii, A.: Efficient iterative algorithms for variational inequalities with monotone operators. *Ekonomika i Matematicheskie Metody* **17**(2), 344–359 (1981). (in Russian; English translation: *Matekon*)
25. Nemirovski, A.: Prox-method with rate of convergence $O(1/t)$ for variational inequalities with Lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM J. Optim.* **15**, 229–251 (2004)
26. Nemirovski, A., Onn, S., Rothblum, U.: Accuracy certificates for computational problems with convex structure. *Math. Oper. Res.* **35**, 52–78 (2010)
27. Nesterov, Yu., Nemirovski, A.: On first order algorithms for ℓ_1 /nuclear norm minimization. *Acta Numer.* **22**, 509–575 (2013)
28. Pshenichnyi, B.N., Danilin, Y.M.: *Numerical Methods in Extremal Problems*. Mir Publishers, Moscow (1978)
29. Rockafellar, R.T.: Minimax theorems and conjugate saddle functions. *Math. Scand.* **14**, 151–173 (1964)
30. Rockafellar, R.T.: *Convex Analysis*. Princeton University Press, Princeton (1970)
31. Shalev-Shwartz, S., Gonen, A., Shamir, O.: Large-scale convex minimization with a low-rank constraint (2011). E-print: [arXiv:1106.1622](https://arxiv.org/abs/1106.1622)