

Generalized conditional subgradient and generalized mirror descent: duality, convergence, and symmetry

Javier F. Peña*

February 28, 2019

Abstract

We provide new insight into a *generalized conditional subgradient* algorithm and a *generalized mirror descent* algorithm for the convex minimization problem

$$\min_x \{f(Ax) + h(x)\}.$$

As Bach showed in [*SIAM J. Optim.*, 25 (2015), pp. 115–129], applying either of these two algorithms to this problem is equivalent to applying the other one to its Fenchel dual. We leverage this duality relationship to develop new upper bounds and convergence results for the gap between the primal and dual iterates generated by these two algorithms. We also propose a new *primal-dual hybrid* algorithm that combines features of the conditional subgradient and mirror descent algorithms to solve the primal and dual problems in a symmetric fashion. Our algorithms and main results rely only on the availability of computable oracles for ∂f and ∂h^* , and for A and A^* .

AMS Subject Classification: 90C25, 90C46

Keywords: conditional gradient, mirror descent, Fenchel duality, symmetry

*Tepper School of Business, Carnegie Mellon University, USA, jfp@andrew.cmu.edu

1 Introduction

Consider the convex minimization problem

$$\min_{x \in X} \{f(Ax) + h(x)\} \quad (1)$$

where $A : X \rightarrow Y$ is a linear mapping between finite dimensional real vector spaces X, Y , and $f : Y \rightarrow \mathbb{R} \cup \{\infty\}$ and $h : X \rightarrow \mathbb{R} \cup \{\infty\}$ are closed convex functions.

We provide new insight into two natural algorithms for (1), namely a *generalized conditional subgradient* algorithm and a *generalized mirror descent* algorithm. These algorithms rely only on the availability of computable oracles for ∂f and ∂h^* . Our approach hinges on the interesting pairing established by Bach [1] between these algorithms and the problem (1) and its Fenchel dual

$$\max_{u \in Y^*} \{-f^*(u) - h^*(-A^*u)\}. \quad (2)$$

Here $A^* : Y^* \rightarrow X^*$ is the adjoint of A , and $f^* : Y^* \rightarrow \mathbb{R} \cup \{\infty\}$ and $h^* : X^* \rightarrow \mathbb{R} \cup \{\infty\}$ are the Fenchel conjugates of f and h respectively [4, 14, 21].

As Bach [1] showed, and as we detail in Section 4 below, applying the generalized conditional subgradient algorithm to (1) is equivalent to applying the generalized mirror descent algorithm to (2). Alternatively, applying the generalized mirror descent algorithm to (1) is equivalent to applying the generalized conditional subgradient algorithm to (2).

Our central results (Theorem 1 and Theorem 2) take this duality relationship further. These results give generic upper bounds on the gap between the primal and dual iterates generated by the generalized conditional subgradient, the generalized mirror descent, and a new *primal-dual hybrid* algorithm. We subsequently leverage these results to obtain an interesting generalization of the classic $\mathcal{O}(1/k)$ convergence rate of the conditional gradient algorithm [11, 15]. More precisely, we show that when the step sizes are properly chosen, the duality gap between the primal and an average of dual iterates generated by the generalized conditional subgradient algorithm converges to zero at a rate $\mathcal{O}(1/k^{\gamma-1})$ provided f satisfies a suitable γ -curvature condition relative to h for some constant $\gamma > 1$. We obtain analogous results for the generalized mirror descent and for the primal-dual hybrid algorithms. The classic $\mathcal{O}(1/k)$ rate corresponds to the special case $\gamma = 2$.

1.1 Positioning of the paper and related work

This paper sheds new light into the close duality connection between generalizations of two popular algorithmic schemes for problems of the form (1), namely the conditional gradient (also known as Frank-Wolfe) algorithm [10, 11, 15] and the mirror descent algorithm [3, 9, 19]. The conditional gradient and the mirror descent algorithms share the feature of not requiring any orthogonal projections. This feature makes them

attractive in a variety of applications where orthogonal projections are too costly or impractical but where subgradient oracles and Bregman projections are viable. Both the conditional gradient and mirror descent algorithms as well as numerous variants of them have been subjects of active research for several years. Some of the many articles in this rapidly evolving literature include [2, 5, 8, 11–13, 15–18, 20, 22–24] as well as the many references therein.

Our work is inspired by and extends the ideas and results introduced by Bach [1] who showed the correspondence between a generalized conditional subgradient algorithm and a generalized mirror descent algorithm for (1) and (2). Bach [1] also showed convergence rates for both algorithms under certain strong convexity and Lipschitz assumptions. In contrast to the approach followed by Bach [1], a main feature of our work is our focus on the gap between the primal and dual iterates. Another main feature of our approach is the lack of reliance on any strong convexity or Lipschitz conditions. Indeed, our algorithms and results make no references to any norms in X or Y at all. Our approach enables us to give tighter and more general analyses of the generalized conditional subgradient, generalized mirror descent, and a new primal-dual hybrid algorithms under very general and mild assumptions. Our duality gap approach also suggests a novel line-search strategy for selecting the step sizes in these three algorithms. This strategy in turn gives an interesting generalization of the classical $\mathcal{O}(1/k)$ convergence rate of the conditional gradient algorithm [11, 15]. Our generalization relies on a new concept of *relative γ -curvature*. This concept is a natural extension of the curvature constant introduced by Jaggi [15]. It is also similar in spirit to the concepts of *relative continuity* and *relative smoothness* as defined by Lu [17], Lu et al. [18], Bauschke et al. [2], and Teboulle [22].

Our algorithms and main results rely solely on the minimal conditions stated in Assumption 1 in Section 2. This assumption concerns the availability of computable oracles for ∂f and ∂h^* and the compatibility of the ranges and domains of these oracles. Our convergence results rely on the additional mild Assumption 2 concerning the computability of some generalized Bregman distances. Our algorithms and results are readily invariant under invertible affine transformations of the spaces X and Y .

1.2 Organization of the paper

The main sections of the paper are organized as follows. Section 2 describes some technical assumptions that we make throughout the paper. This section also introduces the concept of *generalized Bregman distance* that plays a central role in this paper. Section 3 through Section 5 present our main developments. For exposition purposes, these sections consider the special case when $X = Y$ and $A : X \rightarrow X$ is the identity. This simplification enables us to convey the gist of our developments more easily. Section 6 describes how all of our developments extend to the more general problem (1).

In the special case when $X = Y$ and $A : X \rightarrow X$ is the identity, problem (1)

becomes

$$\min_{x \in X} \{f(x) + h(x)\}. \quad (3)$$

Section 3 motivates and presents Algorithm 1 and Algorithm 2 which give descriptions of the generalized conditional subgradient and generalized mirror descent algorithms for problem (3). Section 4 presents our core developments. First, we detail the equivalence between applying Algorithm 1 to (3) and applying Algorithm 2 to the Fenchel dual of (3). We then introduce a new *primal-dual hybrid* algorithm, namely Algorithm 3, that combines features of Algorithm 1 and Algorithm 2 in a perfectly symmetric fashion. Section 4 also presents our main results, namely Theorem 1 and Theorem 2. Theorem 1 establishes a bound on the duality gap between the $(k + 1)$ -th primal iterate and a convex combination of the first k dual iterates generated when Algorithm 1 is applied to (3). As Corollary 1 states, an equivalent dual result readily follows when Algorithm 2 is applied to (3). Theorem 2 gives a similar bound on the gap between the primal and dual iterates generated by Algorithm 3.

Section 5 leverages the results of Section 4 to bound the rate convergence to zero of the gap between primal and dual iterates generated when Algorithm 1, Algorithm 2, or Algorithm 3 is applied to (3) and the step sizes are chosen judiciously. Theorem 3 shows that for Algorithm 1 this gap converges to zero at a rate $\mathcal{O}(1/k^{\gamma-1})$ when f satisfies a suitable γ -curvature condition relative to h for $\gamma > 1$. Corollary 3 and Theorem 4 state analogous results for Algorithm 2 and Algorithm 3.

Finally, Section 6 shows how all of our developments extend to the more general problem (1).

2 Technical background

We will rely on basic convex analysis machinery concerning convex functions, subgradients, Fenchel conjugate, and Fenchel duality as presented in the textbooks [4, 14, 21].

2.1 Technical assumptions

Throughout the paper we will make the following blanket assumption about the tuple (X, Y, A, f, h) .

Assumption 1. X, Y are finite dimensional real vector spaces, $A : X \rightarrow Y$ is a linear mapping, and $f : Y \rightarrow \mathbb{R} \cup \{\infty\}$ and $h : X \rightarrow \mathbb{R} \cup \{\infty\}$ are closed convex functions. There are available oracles that compute $x \mapsto Ax$ and $y \mapsto A^*u$ for all $x \in X, u \in Y^*$. Furthermore, for all $y \in \text{dom}(\partial f)$ and $u \in \text{dom}(\partial h^*)$ the following two conditions hold:

(i) there are available oracles that compute

$$y \mapsto \partial f(y) := \operatorname{argmax}_{v \in Y^*} \{\langle v, y \rangle - f^*(v)\}$$

and

$$u \mapsto \partial h^*(u) := \operatorname{argmax}_{x \in X} \{\langle u, x \rangle - h(x)\},$$

(ii) $-A^* \partial f(y) \in \operatorname{dom}(\partial h^*)$ and $A \partial h^*(u) \in \operatorname{dom}(\partial f)$.

The line-search procedures in Section 5 and Section 6 will require the following additional mild assumption.

Assumption 2. The following *generalized Bregman distances* [6, 7] are computable for all $x, y \in \operatorname{dom}(\partial f)$ and $u, v \in \operatorname{dom}(\partial h^*)$:

$$D_f(y, x) = f(y) - f(x) - \langle \partial f(x), y - x \rangle,$$

and

$$D_{h^*}(v, u) = h^*(v) - h^*(u) - \langle v - u, \partial h^*(u) \rangle.$$

Assumption 1 ensures the well-posedness of problem (1) and also ensures that (1) is amenable to the algorithms introduced in Section 3 and Section 6 below. In particular, the technical assumptions imply that if $x \in \operatorname{dom}(\partial f \circ A)$ and $u \in -\operatorname{dom}(\partial h^* \circ A^*)$ then the points $x_+ := \partial h^*(-A^*u)$ and $u_+ := \partial f(Ax)$ are feasible for (1) and (2) respectively, that is, $x_+ \in \operatorname{dom}(f \circ A) \cap \operatorname{dom}(h)$ and $u_+ \in \operatorname{dom}(f^*) \cap (-\operatorname{dom}(h^* \circ A^*))$. The mapping

$$(x, u) \mapsto (\partial h^*(-A^*u), \partial f(Ax))$$

is at the heart of the algorithms in Section 3 and Section 6 below.

2.2 Some notational convention

We will rely on the following convenient notational convention. When $x = \partial h^*(v)$ for some $v \in \operatorname{dom}(\partial h^*)$, we will write $\partial h(x)$ to denote v . In this case, we will also write $D_h(y, x)$ to denote

$$D_h(y, x) := h(y) - h(x) - \langle v, y - x \rangle = h(y) - h(x) - \langle \partial h(x), y - x \rangle.$$

In a symmetric fashion, when $u = \partial f(y)$ for some $y \in \operatorname{dom}(\partial f)$, we will write $\partial f^*(u)$ to denote y and $D_{f^*}(v, u)$ to denote

$$D_{f^*}(v, u) := f^*(v) - f^*(u) - \langle v - u, y \rangle = f^*(v) - f^*(u) - \langle v - u, \partial f^*(u) \rangle.$$

2.3 Traditional conditional gradient context

Assumption 1 and Assumption 2 readily hold in the usual set up of the conditional gradient algorithm [11, 15]. Consider the problem

$$\min_{x \in Q} f(x), \tag{4}$$

where f is differentiable on the compact convex set $Q \subseteq X$ and there is a linear oracle that computes the support function

$$u \mapsto \operatorname{argmax}_{x \in Q} \langle u, x \rangle.$$

Problem (4) can be written in the form (1) by taking $h := \delta_Q$, the indicator function of the set Q , and $A : X \rightarrow X$ equal to the identity. The linear oracle for Q corresponds to an oracle for ∂h^* and the compactness of Q implies that $\operatorname{dom}(\partial h^*) = X$. Thus Assumption 1 and Assumption 2 readily hold in this case provided oracles for f and ∇f are available.

3 Generalized conditional subgradient and generalized mirror descent

Suppose that $X = Y$, $A : X \rightarrow X$ is the identity mapping and suppose that the tuple (X, Y, A, f, h) satisfies Assumption 1.

The thrust for this paper are two natural algorithmic schemes for solving (3) via the available oracles for ∂f and ∂h^* . The first scheme is a generalized version of the conditional gradient (also known as Frank-Wolfe) algorithm [10, 11, 15] based on the following update. Given a current trial solution x for (3), compute a new trial solution x_+ via

$$x_+ := (1 - \alpha)x + \alpha s$$

where $\alpha \in [0, 1]$ and

$$s := \operatorname{argmin}_{y \in X} \{ \langle \partial f(x), y \rangle + h(y) \} = \partial h^*(-\partial f(x)).$$

The second scheme is a generalized mirror descent algorithm [3, 9, 19] based on the following update. Given a current trial solution y that satisfies $y = \partial h^*(v)$ for some $v \in \operatorname{dom}(h^*)$, compute a new trial solution y_+ via

$$\begin{aligned} y_+ &= \operatorname{argmin}_{z \in X} \{ \alpha \langle \partial f(y) + v, z \rangle + D_h(z, y) \} \\ &= \operatorname{argmin}_{z \in X} \{ \alpha \langle \partial f(y) + \partial h(y), z \rangle + D_h(z, y) \}. \end{aligned}$$

The generalized conditional subgradient update can be written as

$$x_+ = (1 - \alpha)x + \alpha \partial h^*(-\partial f(x)). \tag{5}$$

On the other hand, the generalized mirror descent update can be written as

$$\begin{aligned} y_+ &= \operatorname{argmin}_{z \in X} \{ \langle \alpha \partial f(y) + (\alpha - 1) \partial h(y), z \rangle + h(z) \} \\ &= \partial h^*((1 - \alpha) \partial h(y) - \alpha \partial f(y)). \end{aligned} \tag{6}$$

Observe the striking similarity between (5) and (6). This similarity is at the heart of our main developments. As we detail below, the similarity between (5) and (6) underlies the duality between the generalized conditional subgradient and the generalized mirror descent algorithms [1].

Algorithm 1 and Algorithm 2 give descriptions of a generalized conditional subgradient and a generalized mirror descent algorithm for (3) respectively. It is easy to see that the iterates generated by Algorithm 1 and by Algorithm 2 implement the update rules (5) and (6) respectively. We chose the descriptions in Algorithm 1 and Algorithm 2 to make the resemblance between the two algorithms more salient and to highlight that the algorithms only rely on the oracles for ∂f and ∂h^* .

Algorithm 1 Generalized conditional subgradient

```

1: input:  $(f, h)$  and  $x_0 \in \text{dom}(\partial f)$ 
2: for  $k = 0, 1, 2, \dots$  do
3:    $u_k := \partial f(x_k)$ 
4:    $s_k := \partial h^*(-u_k)$ 
5:   pick  $\alpha_k \in [0, 1]$ 
6:    $x_{k+1} := (1 - \alpha_k)x_k + \alpha_k s_k$ 
7: end for

```

Algorithm 2 Generalized mirror descent

```

1: input:  $(f, h)$  and  $v_0 \in \text{dom}(\partial h^*)$ 
2: for  $k = 0, 1, 2, \dots$  do
3:    $y_k := \partial h^*(v_k)$ 
4:    $z_k := \partial f(y_k)$ 
5:   pick  $\alpha_k \in [0, 1]$ 
6:    $v_{k+1} := (1 - \alpha_k)v_k - \alpha_k z_k$ 
7: end for

```

4 Duality and main results

Again suppose that $X = Y$, $A : X \rightarrow X$ is the identity mapping and suppose that the tuple (X, Y, A, f, h) satisfies Assumption 1. Consider the Fenchel dual [4, 14, 21] of (3)

$$\max_{u \in X^*} \{-f^*(u) - h^*(-u)\}, \quad (7)$$

which can be written as

$$\min_{v \in X^*} \{h^*(v) + \tilde{f}^*(v)\} \quad (8)$$

for $\tilde{f} : X \rightarrow \mathbb{R} \cup \{\infty\}$ is defined via $\tilde{f}(y) := f(-y)$.

Notice the nice symmetry between (3) and (8). The subgradient oracles for the pair (f, h^*) are in one-to-one correspondence with subgradient oracles for the pair (h^*, \tilde{f}^{**}) respectively. Thus, like problem (3), problem (8) is amenable to both Algorithm 1 and Algorithm 2.

Furthermore, as shown by Bach [1], the following interesting duality between these two algorithms holds. Running Algorithm 1 on (3) is identical to running Algorithm 2 on (8). More precisely, the iterates generated by Algorithm 1 applied to (f, h) and started from $x_0 \in \text{dom}(\partial f)$ are the same, modulo some flipped signs, as those generated by Algorithm 2 applied to (h^*, \tilde{f}^*) and started from $v_0 = -x_0 \in \text{dom}(\partial \tilde{f}^{**})$. Indeed, by letting $(v_k, y_k, z_k) := (-x_k, -u_k, s_k)$, the update at each iteration of Algorithm 1

$$x_+ = (1 - \alpha)x + \alpha s = (1 - \alpha)x + \alpha \partial h^*(-u), \quad u = \partial f(x)$$

can be written as

$$v_+ = (1 - \alpha)v - \alpha z = (1 - \alpha)v - \alpha \partial h^*(y), \quad y = -\partial f(-v) = \partial \tilde{f}^{**}(v),$$

which is exactly the update at each iteration of Algorithm 2 applied to (h^*, \tilde{f}^*) .

Based on the above duality relationship, we propose a new *primal-dual hybrid* algorithm described in Algorithm 3 below. Unlike Algorithm 1 and Algorithm 2, Algorithm 3 is symmetric when applied to (3) and (8). That is, Algorithm 3 applied to (f, h) and started from $x_0 \in \text{dom}(\partial f)$, $u_0 \in -\text{dom}(\partial h^*)$ generates the same iterates as it does when applied to (h^*, \tilde{f}^*) and started from $-u_0 \in \text{dom}(\partial h^*)$, $x_0 \in -\text{dom}(\partial \tilde{f}^{**})$.

Algorithm 3 Primal-dual hybrid

- 1: **input:** (f, h) and $x_0 \in \text{dom}(\partial f)$, $u_0 \in -\text{dom}(\partial h^*)$
 - 2: **for** $k = 0, 1, 2, \dots$ **do**
 - 3: $(s_k, z_k) := (\partial h^*(-u_k), \partial f(x_k))$
 - 4: pick $\alpha_k \in [0, 1]$
 - 5: $(x_{k+1}, u_{k+1}) := (1 - \alpha_k)(x_k, u_k) + \alpha_k(s_k, z_k)$
 - 6: **end for**
-

The weak duality between (3) and (7), which is well-known and easy to show [4, 14, 21], is equivalent to the non-negativity of the duality gap:

$$f(x) + h(x) + f^*(u) + h^*(-u) \geq 0 \quad \text{for all } x \in X, u \in X^*.$$

The next two theorems, which are the central results of this paper, provide *upper* bounds on the duality gap for the primal and dual iterates generated by Algorithm 1, Algorithm 2, and Algorithm 3. As we detail in Section 5 below, Theorem 1 and Theorem 2 in turn imply that the duality gap converges to zero at a rate $\mathcal{O}(1/k^{\gamma-1})$ for $\gamma > 1$ when the pair of functions (f, h) satisfies a suitable relative γ -curvature conditions and the step sizes are judiciously chosen.

The statements below will rely on the following double sequences λ_i^k, μ_i^k for $k = 1, 2, \dots$ and $i = 0, 1, \dots, k-1$ determined by a sequence of step sizes $\alpha_k \in [0, 1]$, $k = 0, 1, \dots$. For $k = 0, 1, 2, \dots$ let

$$\begin{aligned}\lambda_k^{k+1} &= \alpha_k, \quad \mu_k^{k+1} = 1 \\ \lambda_i^{k+1} &= (1 - \alpha_k)\lambda_i^k, \quad i = 0, \dots, k-1 \\ \mu_i^{k+1} &= (1 - \alpha_k)\mu_i^k, \quad i = 0, \dots, k-1.\end{aligned}\tag{9}$$

Observe that if $\alpha_0 = 1$ then for each $k = 1, 2, \dots$ we have $\lambda_i^k \geq 0$, $i = 0, 1, \dots, k-1$ and $\sum_{i=0}^{k-1} \lambda_i^k = 1$.

Theorem 1. *Let (x_k, u_k, s_k) , $k = 0, 1, 2, \dots$ be the sequence of iterates generated by Algorithm 1 applied to (3). If $\alpha_0 = 1$ then for $k = 1, 2, \dots$*

$$\sum_{i=0}^{k-1} \lambda_i^k (f^*(u_i) + h^*(-u_i)) - \sum_{i=0}^{k-1} \mu_i^k D_f(x_{i+1}, x_i) \leq -f(x_k) - h(x_k),\tag{10}$$

where λ_i^k, μ_i^k are as in (9).

Proof. First, $u_k = \partial f(x_k)$ and $s_k = \partial h^*(-u_k)$ imply that

$$f^*(u_k) + h^*(-u_k) = \langle u_k, x_k - s_k \rangle - f(x_k) - h(s_k).\tag{11}$$

We now prove (10) by induction. For $k = 1$ we have $x_1 = s_0 = \partial h^*(-u_0)$ because $\alpha_0 = 1$. Thus (11) implies

$$\begin{aligned}f^*(u_0) + h^*(-u_0) &= \langle u_0, x_0 - x_1 \rangle - f(x_0) - h(x_1) \\ &= -\langle \partial f(x_0), x_1 - x_0 \rangle - f(x_0) - h(x_1) \\ &= D_f(x_1, x_0) - f(x_1) - h(x_1).\end{aligned}$$

Hence (10) holds for $k = 1$ since $\lambda_0^1 = \mu_0^1 = 1$.

Suppose (10) holds for $k \geq 1$. Adding up $(1 - \alpha_k)$ times (10) plus α_k times (11), and using (9), the convexity of h , and $x_{k+1} = (1 - \alpha_k)x_k + \alpha_k s_k$ we obtain

$$\begin{aligned}&\sum_{i=0}^k \lambda_i^{k+1} (f^*(u_i) + h^*(-u_i)) - \sum_{i=0}^k \mu_i^{k+1} D_f(x_{i+1}, x_i) \\ &\leq -f(x_k) - (1 - \alpha_k)h(x_k) - \alpha_k h(s_k) + \alpha_k \langle u_k, x_k - s_k \rangle - D_f(x_{k+1}, x_k) \\ &\leq -f(x_k) - h(x_{k+1}) - \langle \partial f(x_k), x_{k+1} - x_k \rangle - D_f(x_{k+1}, x_k) \\ &= -f(x_{k+1}) - h(x_{k+1}).\end{aligned}$$

Therefore (10) holds for $k + 1$ as well. \square

The duality between (3) and (7) and between Algorithm 1 and Algorithm 2 automatically yields the following dual counterpart of Theorem 1.

Corollary 1. Let (v_k, y_k, z_k) , $k = 0, 1, 2, \dots$ be the sequence of iterates generated by Algorithm 2 applied to (3). If $\alpha_0 = 1$ then for $k = 1, 2, \dots$

$$\sum_{i=0}^{k-1} \lambda_i^k (f(y_i) + h(y_i)) - \sum_{i=0}^{k-1} \mu_i^k D_{h^*}(v_{i+1}, v_i) \leq -f^*(-v_k) - h^*(v_k), \quad (12)$$

where λ_i^k, μ_i^k are as in (9).

We also have the following symmetric analogue of Theorem 1.

Theorem 2. Let (x_k, u_k, s_k, z_k) , $k = 0, 1, 2, \dots$ be the sequence of iterates generated by Algorithm 3. If $\alpha_0 = 1$ then for $k = 1, 2, \dots$

$$f(x_k) + h(x_k) + f^*(u_k) + h^*(-u_k) \leq \sum_{i=0}^{k-1} \mu_i^k (D_f(x_{i+1}, x_i) + D_{h^*}(-u_{i+1}, -u_i)), \quad (13)$$

where μ_i^k are as in (9).

Proof. We proceed by induction. For $k = 1$ we have $x_1 = s_0 = \partial h^*(-u_0)$ and $u_1 = z_0 = \partial f(x_0)$ because $\alpha_0 = 1$. Thus

$$\begin{aligned} & f(x_1) + h(x_1) + f^*(u_1) + h^*(-u_1) \\ &= f(x_1) - \langle u_0, x_1 \rangle - h^*(-u_0) + \langle u_1, x_0 \rangle - f(x_0) + h^*(-u_1) \\ &= f(x_1) - f(x_0) - \langle u_1, x_1 - x_0 \rangle + h^*(-u_1) - h^*(-u_0) - \langle u_0 - u_1, x_1 \rangle \\ &= f(x_1) - f(x_0) - \langle \partial f(x_0), x_1 - x_0 \rangle + h^*(-u_1) - h^*(-u_0) - \langle u_0 - u_1, \partial h^*(-u_0) \rangle \\ &= D_f(x_1, x_0) + D_{h^*}(-u_1, -u_0). \end{aligned}$$

Hence (13) holds for $k = 1$ since $\mu_0^1 = 1$.

Suppose (13) holds for $k \geq 1$. Since $z_k = \partial f(x_k)$ and $s_k = \partial h^*(-u_k)$ and $(x_{k+1}, u_{k+1}) = (1 - \alpha_k)(x_k, u_k) + \alpha_k(s_k, z_k)$, it follows that

$$f(x_k) + h(s_k) + f^*(z_k) + h^*(-u_k) = \langle z_k, x_k \rangle - \langle u_k, s_k \rangle, \quad (14)$$

and

$$\begin{aligned} & f(x_{k+1}) - f(x_k) + h^*(-u_{k+1}) - h^*(-u_k) \\ &= D_f(x_{k+1}, x_k) + D_{h^*}(-u_{k+1}, -u_k) - \alpha_k (\langle z_k, x_k \rangle - \langle u_k, s_k \rangle). \end{aligned} \quad (15)$$

Next, adding up $(1 - \alpha_k)$ times (13) plus α_k times (14) plus (15), and using (9) we get

$$\begin{aligned} & f(x_{k+1}) + (1 - \alpha_k)h(x_k) + \alpha_k h(s_k) + (1 - \alpha_k)f^*(u_k) + \alpha_k f^*(z_k) + h^*(-u_{k+1}) \\ & \leq \sum_{i=0}^k \mu_i^{k+1} (D_f(x_{i+1}, x_i) + D_{h^*}(-u_{i+1}, -u_i)). \end{aligned}$$

Since $(x_{k+1}, u_{k+1}) = (1 - \alpha_k)(x_k, u_k) + \alpha_k(s_k, z_k)$, the convexity of h and f^* yields

$$\begin{aligned} f(x_{k+1}) + h(x_{k+1}) + f^*(u_{k+1}) + h^*(-u_{k+1}) \\ \leq \sum_{i=0}^k \mu_i^{k+1} (D_f(x_{i+1}, x_i) + D_{h^*}(-u_{i+1}, -u_i)). \end{aligned}$$

Therefore (13) holds for $k + 1$ as well. \square

5 Convergence results

Once again suppose that $X = Y$, $A : X \rightarrow X$ is the identity mapping and suppose that the tuple (X, Y, A, f, h) satisfies Assumption 1. We next leverage Theorem 1, Corollary 1, and Theorem 2 to obtain some convergence results for Algorithm 1, Algorithm 2, and Algorithm 3 applied to (3) and (7).

Theorem 1 readily implies that if $\alpha_0 = 1$ then the sequence of iterates (x_k, u_k, s_k) , $k = 0, 1, 2, \dots$ generated by Algorithm 1 satisfies

$$f(x_k) + g(x_k) + f^*(\hat{u}_k) + h^*(-\hat{u}_k) \leq \text{CGgap}_k, \quad (16)$$

where

$$\hat{u}_k = \sum_{i=0}^{k-1} \lambda_i^k u_i \quad \text{or} \quad \hat{u}_k = \underset{u_0, \dots, u_{k-1}}{\operatorname{argmin}} \{f^*(u_i) + h^*(-u_i)\}, \quad (17)$$

and CGgap_k , $k = 1, 2, \dots$ is defined via $\text{CGgap}_1 = D_f(x_1, x_0)$ and

$$\text{CGgap}_{k+1} = (1 - \alpha_k)\text{CGgap}_k + D_f(x_{k+1}, x_k), \quad k = 1, 2, \dots \quad (18)$$

Similarly, Corollary 1 implies that if $\alpha_0 = 1$ then the sequence of iterates (v_k, y_k, z_k) , $k = 0, 1, 2, \dots$ generated by Algorithm 2 satisfies

$$f(\hat{y}_k) + g(\hat{y}_k) + f^*(-v_k) + h^*(v_k) \leq \text{MDgap}_k, \quad (19)$$

where

$$\hat{y}_k = \sum_{i=0}^{k-1} \lambda_i^k y_i \quad \text{or} \quad \hat{y}_k = \underset{y_0, \dots, y_{k-1}}{\operatorname{argmin}} \{f(y_i) + h(y_i)\}, \quad (20)$$

and MDgap_k , $k = 1, 2, \dots$ is defined via $\text{MDgap}_1 = D_{h^*}(v_1, v_0)$ and

$$\text{MDgap}_{k+1} = (1 - \alpha_k)\text{MDgap}_k + D_{h^*}(v_{k+1}, v_k), \quad k = 1, 2, \dots \quad (21)$$

On the other hand, Theorem 2 implies that if $\alpha_0 = 1$ then the sequence of iterates (x_k, u_k, s_k, z_k) , $k = 0, 1, 2, \dots$ generated by Algorithm 3 satisfies

$$f(x_k) + h(x_k) + f^*(u_k) + h^*(-u_k) \leq \text{HYBgap}_k, \quad (22)$$

where $\text{HYBgap}_1 = D_f(x_1, x_0) + D_{h^*}(-u_1, -u_0)$ and

$$\text{HYBgap}_{k+1} = (1 - \alpha_k)\text{HYBgap}_k + D_f(x_{k+1}, x_k) + D_{h^*}(-u_{k+1}, -u_k), \quad k = 1, 2, \dots \quad (23)$$

As the propositions below formally show, the above observations yield familiar $\mathcal{O}(1/k)$ convergence results for the popular step size $\alpha_k = 2/(k+2)$, $k = 0, 1, \dots$ provided a suitable *relative quadratic curvature* condition holds.

Definition 1. We say that f has *quadratic curvature relative to h* if there exists a finite constant C such that for all $x \in \text{dom}(\partial f)$ and $v \in \text{dom}(\partial h^*)$ the following inequality holds for $s := \partial h^*(v)$

$$D_f(x + \alpha(s - x), x) \leq \frac{C\alpha^2}{2} \quad \text{for all } \alpha \in [0, 1]. \quad (24)$$

This new concept of relative quadratic curvature condition is a generalization of the curvature constant introduced by Jaggi [15]. Indeed, consider a problem of the form (4) where Q is compact and convex, f is differentiable on Q , and a linear oracle for Q is available. In this context, Jaggi [15] defines the curvature constant of f on Q as follows

$$C_{f,Q} = \sup_{\substack{x, s \in Q \\ \alpha \in (0, 1]}} \frac{D_f(x + \alpha(s - x), x)}{\alpha^2/2}.$$

Observe that for $h = \delta_Q$, inequality (24) holds if $C \geq C_{f,Q}$. We note that the smallest constant C such that (24) holds for all $x \in Q$ and $s = \partial h^*(v)$, $v \in \text{dom}(\partial h^*)$ could be potentially smaller.

The above concept of relative quadratic curvature is inspired by the concepts of relative smoothness and relative continuity introduced in [2, 17, 18, 22].

Proposition 1. *Suppose f has quadratic curvature relative to h with constant C . If $\alpha_k = 2/(k+2)$, $k = 0, 1, \dots$ then the sequence of iterates (x_k, u_k, s_k) , $k = 1, 2, \dots$ generated by Algorithm 1 satisfies*

$$f(x_k) + h(x_k) + f^*(\hat{u}_k) + h^*(-\hat{u}_k) \leq \frac{2C}{k+2} \quad (25)$$

where \hat{u}_k is as in (17).

Proof. By (16), it suffices to show that for $k = 1, 2, \dots$

$$\text{CGgap}_k \leq \frac{2C}{k+2}. \quad (26)$$

We proceed by induction. For $k = 1$ inequality (24) and $\alpha_0 = 1$ imply that

$$\text{CGgap}_0 = D_f(x_1, x_0) = D_f(x_0 + (s_0 - x_0), x_0) \leq \frac{C}{2} \leq \frac{2C}{3}.$$

Hence (26) holds for $k = 1$. Suppose (26) holds for $k \geq 1$. Then (18), (24), and $\alpha_k = 2/(k+2)$ imply that

$$\text{CGgap}_{k+1} \leq \frac{k}{k+2} \cdot \frac{2C}{k+2} + \frac{2C}{(k+2)^2} = \frac{2C(k+1)}{(k+2)^2} \leq \frac{2C}{k+3}.$$

Therefore (26) holds for $k+1$ as well. □

Again the duality between (3) and (7) and between Algorithm 1 and Algorithm 2 automatically yield the following corollary of Proposition 1. Recall that \tilde{f} is defined via $\tilde{f}(y) = f(-y)$.

Corollary 2. *Suppose h^* has quadratic curvature relative to \tilde{f}^* with constant C^* . If $\alpha_k = 2/(k+2)$, $k = 0, 1, \dots$ then the sequence of iterates (v_k, y_k, z_k) , $k = 0, 1, 2, \dots$ generated by Algorithm 2 satisfies*

$$f(\hat{y}_k) + h(\hat{y}_k) + f^*(-v_k) + h^*(v_k) \leq \frac{2C^*}{k+2}$$

where \hat{y}_k is as in (20).

The same inductive argument underlying the proof of Proposition 1 together with (22) and (23) yields the following analogous result for Algorithm 3.

Proposition 2. *Suppose f has quadratic curvature relative to h with constant C and h^* has quadratic curvature relative to \tilde{f}^* with constant C^* . If $\alpha_k = 2/(k+2)$, $k = 0, 1, \dots$ then the sequence of iterates (x_k, u_k, s_k, z_k) , $k = 0, 1, 2, \dots$ generated by Algorithm 3 satisfies*

$$f(x_k) + h(x_k) + f^*(u_k) + h^*(-u_k) \leq \frac{2(C + C^*)}{k+2}.$$

The identity (18) suggests the following *line-search* procedure to select the step size α_k in Algorithm 1:

$$\alpha_k := \operatorname{argmin}_{\alpha \in [0,1]} \{(1 - \alpha)\text{CGgap}_k + D_f(x_k + \alpha(s_k - x_k), x_k)\}. \quad (27)$$

Likewise, the identity (21) suggests the following *line-search* procedure to select the step size α_k in Algorithm 2:

$$\alpha_k := \operatorname{argmin}_{\alpha \in [0,1]} \{(1 - \alpha)\text{MDgap}_k + D_{h^*}(v_k + \alpha(z_k - v_k), v_k)\}. \quad (28)$$

Similarly, the identity (23) suggests the following *line-search* procedure to select the step size α_k in Algorithm 3:

$$\begin{aligned} \alpha_k := \operatorname{argmin}_{\alpha \in [0,1]} \{ & (1 - \alpha)\text{HYBgap}_k + D_f(x_k + \alpha(s_k - x_k), x_k) \\ & + D_{h^*}(-u_k - \alpha(z_k - u_k), -u_k)\}. \end{aligned} \quad (29)$$

The above line-search procedures are computable via binary search provided Assumption 2 holds. These line-search procedures enable us to prove the convergence of Algorithm 1, Algorithm 2, and Algorithm 3 under the following more general γ -curvature condition.

Definition 2. Let $\gamma > 1$. We say that f has γ -curvature relative to h if there exists a finite constant C such that for all $x \in \text{dom}(f)$ and $v \in \text{dom}(\partial h^*)$ the following inequality holds for $s = \partial h^*(v)$

$$D_f(x + \alpha(s - x), x) \leq \frac{C\alpha^\gamma}{\gamma} \text{ for all } \alpha \in [0, 1]. \quad (30)$$

We have the following interesting generalization of Proposition 1.

Theorem 3. Suppose $\gamma > 1$ is such that f has γ -curvature relative to h with constant C . If $\alpha_0 = 1$ and $\alpha_k \in [0, 1]$, $k = 1, 2, \dots$ is chosen via (27) then the sequence of iterates (x_k, u_k, s_k) , $k = 0, 1, 2, \dots$ generated by Algorithm 1 satisfies

$$f(x_k) + g(x_k) + f^*(\hat{u}_k) + h^*(-\hat{u}_k) \leq C \left(\frac{\gamma}{k + \gamma} \right)^{\gamma-1}, \quad (31)$$

where \hat{u}_k is as in (17).

Proof. By (16) it suffices to show that for $k = 1, 2, \dots$

$$\text{CGgap}_k \leq C \left(\frac{\gamma}{k + \gamma} \right)^{\gamma-1}. \quad (32)$$

We prove (32) by induction on k . For $k = 1$ we have

$$\text{CGgap}_1 = D_f(x_1, x_0) \leq \frac{C}{\gamma} \leq C \left(\frac{\gamma}{1 + \gamma} \right)^{\gamma-1}. \quad (33)$$

where the last step follows from the weighted arithmetic mean geometric mean inequality. Hence (32) holds for $k = 1$.

Suppose (32) holds for $k \geq 1$. Then (30) and (27) implies that for all $\alpha \in [0, 1]$

$$\text{CGgap}_{k+1} \leq (1 - \alpha)\text{CGgap}_k + \frac{C\alpha^\gamma}{\gamma}$$

In particular, for $\alpha = \gamma/(k + \gamma)$ we have

$$\begin{aligned} \text{CGgap}_{k+1} &\leq C \frac{k}{k + \gamma} \left(\frac{\gamma}{k + \gamma} \right)^{\gamma-1} + \frac{C\gamma^\gamma}{\gamma(k + \gamma)^\gamma} \\ &= \frac{C(k + 1)\gamma^{\gamma-1}}{(k + \gamma)^\gamma} \\ &\leq C \left(\frac{\gamma}{k + 1 + \gamma} \right)^{\gamma-1}, \end{aligned}$$

where the last step follows from the inequality

$$(k+1)(k+1+\gamma)^{\gamma-1} \leq (k+\gamma)^\gamma,$$

which in turn follows from the weighted arithmetic geometric mean inequality. Therefore (32) holds for $k+1$ as well. \square

Once again, the duality between (3) and (7) and between Algorithm 1 and Algorithm 2 automatically yields the following dual counterpart of Theorem 3.

Corollary 3. *Suppose $\gamma > 1$ is such that h^* has γ -curvature relative to \tilde{f}^* with constant C^* . If $\alpha_0 = 1$ and $\alpha_k \in [0, 1]$, $k = 1, 2, \dots$ is chosen via (28) then the sequence of iterates (v_k, y_k, z_k) , $k = 0, 1, 2, \dots$ generated by Algorithm 2 satisfies*

$$f(\hat{y}_k) + g(\hat{y}_k) + f^*(-v_k) + h^*(v_k) \leq C^* \left(\frac{\gamma}{k+\gamma} \right)^{\gamma-1}$$

where \hat{y}_k is as in (20).

We also have the following analogue of Theorem 3 for Algorithm 3. We omit the proof of Theorem 4 since it is a straightforward extension of the proof of Theorem 3.

Theorem 4. *Suppose $\gamma > 1$ is such that f has γ -curvature relative to h with constant C and h^* has γ -curvature relative to \tilde{f}^* with constant C^* . If $\alpha_k \in [0, 1]$, $k = 0, 1, \dots$ are chosen via (29) then the sequence of iterates x_k, u_k , $k = 0, 1, 2, \dots$ generated by Algorithm 3 satisfies*

$$f(x_k) + g(x_k) + f^*(u_k) + h^*(-u_k) \leq (C + C^*) \left(\frac{\gamma}{k+\gamma} \right)^{\gamma-1}.$$

The proof of Theorem 3 readily shows that (31) holds if α_k is chosen as $\alpha_k = \gamma/(k+\gamma)$. However, this requires knowledge of $\gamma > 1$ which is unrealistic and could be too conservative. A similar bound holds if instead α_k is chosen via the following approximate and more realistic line-search procedure. Let $\delta \in (0, 1)$ be a small fixed constant and choose $\alpha_k = \gamma_k/(k+\gamma_k)$ where γ_k is such that $\gamma_k \geq \gamma - \delta$. This can be easily done via binary search as long as Assumption 2 holds. The proof of Theorem 3 shows that in this case the following modified version of (31) holds

$$f(x_k) + g(x_k) + f^*(\hat{u}_k) + h^*(-\hat{u}_k) \leq C \left(\frac{\gamma - \delta}{k + \gamma - \delta} \right)^{\gamma - \delta - 1}.$$

The same considerations apply to the bounds in Corollary 3 and Theorem 4.

6 Extension to the general format

Suppose that the tuple (X, Y, A, f, h) satisfies Assumption 1. We next discuss how all of our previous developments extend to problems in the more general format

$$\min_{x \in X} \{f(Ax) + h(x)\} \quad (34)$$

Algorithm 1, Algorithm 2, and Algorithm 3 extend to (34) as detailed in Algorithm 4, Algorithm 5, and Algorithm 6 respectively. Furthermore, applying Algorithm 4 (Algorithm 5) to (34) is equivalent to applying Algorithm 5 (Algorithm 4) to its Fenchel dual

$$\max_{u \in Y^*} \{-f^*(u) - h^*(-A^*u)\},$$

which can be written as

$$\min_{v \in Y^*} \{h^*(A^*v) + \tilde{f}^*(v)\}$$

for $\tilde{f} : Y \rightarrow \mathbb{R} \cup \{\infty\}$ defined via $\tilde{f}(y) := f(-y)$.

Algorithm 4 Generalized conditional subgradient, version 2

- 1: **input:** (f, h, A) and $x_0 \in \text{dom}(\partial f \circ A)$
 - 2: **for** $k = 0, 1, 2, \dots$ **do**
 - 3: $u_k := \partial f(Ax_k)$
 - 4: $s_k := \partial h^*(-A^*u_k)$
 - 5: pick $\alpha_k \in [0, 1]$
 - 6: $x_{k+1} := (1 - \alpha_k)x_k + \alpha_k s_k$
 - 7: **end for**
-

Algorithm 5 Generalized mirror descent, version 2

- 1: **input:** (f, h, A) and $v_0 \in \text{dom}(\partial h^* \circ A^*)$
 - 2: **for** $k = 0, 1, 2, \dots$ **do**
 - 3: $y_k := \partial h^*(A^*v_k)$
 - 4: $z_k := \partial f(Ay_k)$
 - 5: pick $\alpha_k \in [0, 1]$
 - 6: $v_{k+1} := (1 - \alpha_k)v_k - \alpha_k z_k$
 - 7: **end for**
-

Again by weak duality [4, 14, 21] the duality gap is non-negative:

$$f(Ax) + h(x) + f^*(u) + h^*(-A^*u) \geq 0 \quad \text{for all } x \in X, u \in Y^*.$$

Proposition 3, Proposition 4, and Proposition 5 below give upper bounds on this duality gap for the primal-dual iterates generated by Algorithm 4, Algorithm 5, and Algorithm 6. These propositions are extensions of Theorem 1, Corollary 1, and Theorem 2 respectively.

Algorithm 6 Primal-dual hybrid, version 2

- 1: **input:** (f, h, A) and $x_0 \in \text{dom}(\partial f \circ A)$, $u_0 \in -\text{dom}(\partial h^* \circ A^*)$
 - 2: **for** $k = 0, 1, 2, \dots$ **do**
 - 3: $(s_k, z_k) := (\partial h^*(-A^*u_k), \partial f(Ax_k))$
 - 4: pick $\alpha_k \in [0, 1]$
 - 5: $(x_{k+1}, u_{k+1}) := (1 - \alpha_k)(x_k, u_k) + \alpha_k(s_k, z_k)$
 - 6: **end for**
-

Proposition 3. *Let (x_k, u_k, s_k) , $k = 0, 1, 2, \dots$ be the sequence of iterates generated by Algorithm 4 applied to (34). If $\alpha_0 = 1$ then for $k = 1, 2, \dots$*

$$\sum_{i=0}^{k-1} \lambda_i^k (f^*(u_i) + h^*(-A^*u_i)) - \sum_{i=0}^{k-1} \mu_i^k D_f(Ax_{i+1}, Ax_i) \leq -f(Ax_k) - h(x_k)$$

where λ_i^k, μ_i^k are as in (9).

Proposition 4. *Let (v_k, y_k, z_k) , $k = 0, 1, 2, \dots$ be the sequence of iterates generated by Algorithm 5 applied to (34). If $\alpha_0 = 1$ then for $k = 1, 2, \dots$*

$$\sum_{i=0}^{k-1} \lambda_i^k (f(Ay_i) + h(y_i)) - \sum_{i=0}^{k-1} \mu_i^k D_{h^*}(A^*v_{i+1}, A^*v_i) \leq -f^*(-v_k) - h^*(A^*v_k),$$

where λ_i^k, μ_i^k are as in (9).

Proposition 5. *Let (x_k, u_k, s_k, z_k) , $k = 0, 1, 2, \dots$ be the sequence of iterates generated by Algorithm 6 applied to (34). If $\alpha_0 = 1$ then for $k = 1, 2, \dots$*

$$\begin{aligned} f(Ax_k) + h(x_k) + f^*(u_k) + h^*(-A^*u_k) \\ \leq \sum_{i=0}^{k-1} \mu_i^k (D_f(Ax_{i+1}, Ax_i) + D_{h^*}(-A^*u_{i+1}, -A^*u_i)), \end{aligned}$$

where μ_i^k are as in (9).

We omit the proofs of Proposition 3, Proposition 4, and Proposition 5 since they are straightforward modifications of the proofs of Theorem 1, Corollary 1, and Theorem 2.

The developments in Section 5 also extend in a similar fashion. Proposition 3 implies that if $\alpha_0 = 1$ then the sequence of iterates (x_k, u_k, s_k) , $k = 0, 1, 2, \dots$ generated by Algorithm 4 satisfies

$$f(Ax_k) + g(x_k) + f^*(\hat{u}_k) + h^*(-A^*\hat{u}_k) \leq \text{CGgap}_k,$$

where

$$\hat{u}_k = \sum_{i=0}^{k-1} \lambda_i^k u_i \quad \text{or} \quad \hat{u}_k = \underset{u_0, \dots, u_{k-1}}{\text{argmin}} \{f^*(u_i) + h^*(-A^*u_i)\}, \quad (35)$$

$\text{CGgap}_1 = D_f(Ax_1, Ax_0)$, and

$$\text{CGgap}_{k+1} = (1 - \alpha_k)\text{CGgap}_k + D_f(Ax_{k+1}, Ax_k), \quad k = 1, 2, \dots$$

Similarly, Proposition 4 implies that if $\alpha_0 = 1$ then the sequence of iterates (v_k, y_k, z_k) , $k = 0, 1, 2, \dots$ generated by Algorithm 2 satisfies

$$f(A\hat{y}_k) + g(\hat{y}_k) + f^*(-v_k) + h^*(A^*v_k) \leq \text{MDgap}_k,$$

where

$$\hat{y}_k = \sum_{i=0}^{k-1} \lambda_i^k y_i \quad \text{or} \quad \hat{y}_k = \underset{y_0, \dots, y_{k-1}}{\operatorname{argmin}} \{f(Ay_i) + h(y_i)\}, \quad (36)$$

$\text{MDgap}_1 = D_{h^*}(A^*v_1, A^*v_0)$, and

$$\text{MDgap}_{k+1} = (1 - \alpha_k)\text{MDgap}_k + D_{h^*}(A^*v_{k+1}, A^*v_k), \quad k = 1, 2, \dots$$

On the other hand, Proposition 5 implies that if $\alpha_0 = 1$ then the sequence of iterates (x_k, u_k, s_k, z_k) , $k = 0, 1, 2, \dots$ generated by Algorithm 6 satisfies

$$f(Ax_k) + h(x_k) + f^*(u_k) + h^*(-A^*u_k) \leq \text{HYBgap}_k, \quad (37)$$

where $\text{HYBgap}_1 := D_f(Ax_1, Ax_0) + D_{h^*}(-A^*u_1, -A^*u_0)$ and

$$\text{HYBgap}_{k+1} := (1 - \alpha_k)\text{HYBgap}_k + D_f(Ax_{k+1}, Ax_k) + D_{h^*}(-A^*u_{k+1}, -A^*u_k).$$

Consider the following procedures for step size selection. For Algorithm 4:

$$\alpha_k := \underset{\alpha \in [0,1]}{\operatorname{argmin}} \{(1 - \alpha)\text{CGgap}_k + D_f(A(x_k + \alpha(s_k - x_k)), Ax_k)\}. \quad (38)$$

For Algorithm 5:

$$\alpha_k := \underset{\alpha \in [0,1]}{\operatorname{argmin}} \{(1 - \alpha)\text{MDgap}_k + D_{h^*}(A^*(v_k + \alpha(z_k - v_k)), A^*v_k)\}. \quad (39)$$

Finally, for Algorithm 6:

$$\begin{aligned} \alpha_k := \underset{\alpha \in [0,1]}{\operatorname{argmin}} \{ & (1 - \alpha)\text{HYBgap}_k + D_f(Ax_k + \alpha(s_k - x_k), Ax_k) \\ & + D_{h^*}(-A^*(u_k + \alpha(z_k - u_k)), -A^*u_k)\}. \end{aligned} \quad (40)$$

Again the above line-search procedures are computable via binary search provided Assumption 2 holds. Theorem 3, Corollary 3, and Theorem 4 extend as follows.

Proposition 6. *Suppose $\gamma > 1$ is such that $f \circ A$ has γ -curvature relative to h with constant C . If $\alpha_0 = 1$ and $\alpha_k \in [0, 1]$, $k = 1, 2, \dots$ is chosen via (38) then the sequence of iterates (x_k, u_k, s_k) , $k = 0, 1, 2, \dots$ generated by Algorithm 4 satisfies*

$$f(Ax_k) + g(x_k) + f^*(\hat{u}_k) + h^*(-A^*\hat{u}_k) \leq C \left(\frac{\gamma}{k + \gamma} \right)^{\gamma-1}$$

where \hat{u}_k is as in (35).

Proposition 7. *Suppose $\gamma > 1$ is such that $h^* \circ A^*$ has γ -curvature relative to \tilde{f}^* with constant C^* . If $\alpha_0 = 1$ and $\alpha_k \in [0, 1]$, $k = 1, 2, \dots$ is chosen via (39) then the sequence of iterates (v_k, y_k, z_k) , $k = 0, 1, 2, \dots$ generated by Algorithm 5 satisfies*

$$f(A\hat{y}_k) + h(\hat{y}_k) + f^*(-v_k) + h^*(A^*v_k) \leq C^* \left(\frac{\gamma}{k + \gamma} \right)^{\gamma-1}$$

where \hat{y}_k is as in (36).

Proposition 8. *Suppose $\gamma > 1$ is such that $f \circ A$ has γ -curvature relative to h with constant C for some $\gamma > 1$ and $h^* \circ A^*$ has γ -curvature relative to \tilde{f}^* with constant C^* . If $\alpha_0 = 1$ and $\alpha_k \in [0, 1]$, $k = 1, 2, \dots$ is chosen via (40) then the sequence of iterates (x_k, u_k, s_k, z_k) , $k = 0, 1, 2, \dots$ generated by Algorithm 4 satisfies*

$$f(Ax_k) + g(x_k) + f^*(u_k) + h^*(-A^*u_k) \leq (C + C^*) \left(\frac{\gamma}{k + \gamma} \right)^{\gamma-1}.$$

Again we omit the proofs of Proposition 6, Proposition 7, and Proposition 8 since they are straightforward modifications of the proofs of Theorem 3, Corollary 3, and Theorem 4.

7 Conclusions

We discussed three algorithms (Algorithm 4, Algorithm 5, and Algorithm 6) for problem (1) and its Fenchel dual (2) that are based on the mapping

$$(x, u) \mapsto (\partial h^*(-A^*u), \partial f(Ax)).$$

Applying either Algorithm 4 or Algorithm 5 to (1) is equivalent to applying the other one to (2). On the other hand, Algorithm 6 treats (1) and (2) in a completely symmetric fashion.

We established new upper bounds (Proposition 3, Proposition 4, and Proposition 5) on the gap between the primal and dual iterates generated by Algorithm 4, Algorithm 5, and Algorithm 6. These bounds in turn imply that the duality gap converges to zero at a rate $\mathcal{O}(1/k^{\gamma-1})$ for $\gamma > 1$ provided the functions f, h satisfy some suitable γ -curvature conditions and the step sizes are judiciously chosen (Proposition 6, Proposition 7, and Proposition 8).

References

- [1] F. Bach. Duality between subgradient and conditional gradient methods. *SIAM Journal on Optimization*, 25(1):115–129, 2015.

- [2] H. Bauschke, J. Bolte, and M. Teboulle. A descent lemma beyond Lipschitz gradient continuity: first-order methods revisited and applications. *Mathematics of Operations Research*, 42(2):330–348, 2016.
- [3] A. Beck and M. Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3):167–175, 2003.
- [4] J. Borwein and A. Lewis. *Convex analysis and nonlinear optimization*. Springer, New York, 2000.
- [5] G. Braun, S. Pokutta, D. Tu, and S. Wright. Blended conditional gradients: the unconditioning of conditional gradients. *arXiv preprint arXiv:1805.07311*, 2018.
- [6] L. Bregman. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR Computational Mathematics and Mathematical Physics*, 7(3):200–217, 1967.
- [7] Y. Censor and A. Lent. An iterative row-action method for interval convex programming. *Journal of Optimization Theory and Applications*, 34(3):321–353, 1981.
- [8] K. Clarkson. Coresets, sparse greedy approximation, and the Frank-Wolfe algorithm. *ACM Transactions on Algorithms (TALG)*, 6(4):63, 2010.
- [9] J. Duchi, S. Shalev-Shwartz, Y. Singer, and A. Tewari. Composite objective mirror descent. In *COLT*, pages 14–26, 2010.
- [10] M. Frank and P. Wolfe. An algorithm for quadratic programming. *Naval Research Quarterly*, 3:95–110, 1956.
- [11] R. Freund and P. Grigas. New analysis and results for the Frank-Wolfe method. *Mathematical Programming*, 155(1-2):199–230, 2016.
- [12] G. Gidel, T. Jebara, and S. Lacoste-Julien. Frank-Wolfe algorithms for saddle point problems. *arXiv preprint arXiv:1610.07797*, 2016.
- [13] Z. Harchaoui, A. Juditsky, and A. Nemirovski. Conditional gradient algorithms for norm-regularized smooth convex optimization. *Mathematical Programming*, 152(1-2):75–112, 2015.
- [14] J. Hiriart-Urruty and C. Lemaréchal. *Convex Analysis and Minimization Algorithms*. Springer-Verlag, Berlin, 1993.
- [15] M. Jaggi. Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In *ICML*, volume 28 of *JMLR Proceedings*, pages 427–435, 2013.
- [16] G. Lan and Y. Zhou. Conditional gradient sliding for convex optimization. *SIAM Journal on Optimization*, 26(2):1379–1409, 2016.

- [17] H. Lu. “Relative-continuity” for non-Lipschitz non-smooth convex optimization using stochastic (or deterministic) mirror descent. *To Appear in INFORMS Journal on Optimization*, 2019.
- [18] H. Lu, R. Freund, and Y. Nesterov. Relatively smooth convex optimization by first-order methods, and applications. *SIAM Journal on Optimization*, 28(1):333–354, 2018.
- [19] A. Nemirovsky and D. Yudin. *Problem Complexity and Method Efficiency in Optimization*. Wiley, 1983.
- [20] N. Rao, P. Shah, and S. Wright. Forward–backward greedy algorithms for atomic norm regularization. *IEEE Transactions on Signal Processing*, 63(21):5798–5811, 2015.
- [21] T. Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, 1970.
- [22] M. Teboulle. A simplified view of first order methods for optimization. *Mathematical Programming*, pages 1–30, 2018.
- [23] Y. Yu, X. Zhang, and D. Schuurmans. Generalized conditional gradient for sparse estimation. *The Journal of Machine Learning Research*, 18(1):5279–5324, 2017.
- [24] A. Yurtsever, M. Udell, J. Tropp, and V. Cevher. Sketchy decisions: Convex low-rank matrix optimization with optimal storage. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1188–1196, 2017.