

One-sided Frank-Wolfe algorithms for saddle problems

Supplementary material: full version of the ICML'21 paper.
Also available as technical report arXiv:2101.12617v3.

Vladimir Kolmogorov *
Institute of Science and Technology Austria
vnk@ist.ac.at

Thomas Pock †
Institute of Computer Graphics and Vision,
Graz University of Technology
pock@icg.tugraz.at

Abstract

We study a class of convex-concave saddle-point problems of the form $\min_x \max_y \langle Kx, y \rangle + f_{\mathcal{P}}(x) - h^*(y)$ where K is a linear operator, $f_{\mathcal{P}}$ is the sum of a convex function f with a Lipschitz-continuous gradient and the indicator function of a bounded convex polytope \mathcal{P} , and h^* is a convex (possibly nonsmooth) function. Such problem arises, for example, as a Lagrangian relaxation of various discrete optimization problems. Our main assumptions are the existence of an efficient *linear minimization oracle* (**lmo**) for $f_{\mathcal{P}}$ and an efficient *proximal map* (**prox**) for h^* which motivate the solution via a blend of proximal primal-dual algorithms and Frank-Wolfe algorithms. In case h^* is the indicator function of a linear constraint and function f is quadratic, we show a $O(1/n^2)$ convergence rate on the dual objective, requiring $O(n \log n)$ calls of **lmo**. If the problem comes from the constrained optimization problem $\min_{x \in \mathbb{R}^d} \{f_{\mathcal{P}}(x) \mid Ax - b = 0\}$ then we additionally get bound $O(1/n^2)$ both on the primal gap and on the infeasibility gap. In the most general case, we show a $O(1/n)$ convergence rate of the primal-dual gap again requiring $O(n \log n)$ calls of **lmo**. To the best of our knowledge, this improves on the known convergence rates for the considered class of saddle-point problems. We show applications to labeling problems frequently appearing in machine learning and computer vision.

1 Introduction

In this paper, we consider the following class of saddle-point problems:

$$\min_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} \mathcal{L}(x, y) := \langle Kx, y \rangle + f_{\mathcal{P}}(x) - h^*(y) \quad (1)$$

where \mathcal{X}, \mathcal{Y} are finite dimensional spaces, equipped with an inner product $\langle \cdot, \cdot \rangle$ and $K : \mathcal{X} \rightarrow \mathcal{Y}$ is a bounded linear operator with operator norm $L_K = \|K\|$. Usually the underlying spaces are the standard Euclidean spaces $\mathcal{X} = \mathbb{R}^n$ and $\mathcal{Y} = \mathbb{R}^m$.

The functions $f_{\mathcal{P}}(x)$ and h^* are convex, lower semicontinuous functions.

For a differentiable convex function f we say that it has a Lipschitz continuous gradient if there exists a constant $L_f \geq 0$ such that

$$\|\nabla f(x) - \nabla f(y)\| \leq L_f \|x - y\|, \quad \forall x, y \in \mathcal{X}.$$

Moreover, the function f is called strongly convex with strong convexity parameter $\mu_f > 0$ if

$$f(y) \geq f(x) + \langle \nabla f(x), y - x \rangle + \frac{\mu_f}{2} \|y - x\|^2, \quad \forall x, y \in \mathcal{X}.$$

We make the following important structural assumptions on the functions $f_{\mathcal{P}}(x)$ and h^* :

*Vladimir Kolmogorov was supported by the European Research Council under the European Unions Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement no 616160.

†Thomas Pock acknowledges support by an ERC grant HOMOVIS, no 640156.

- The function $f_{\mathcal{P}}(x)$ has the following composite form

$$f_{\mathcal{P}}(x) = f(x) + \delta_{\mathcal{P}}(x),$$

where f is a convex function with a L_f -Lipschitz continuous gradient and $\delta_{\mathcal{P}}$ is the indicator function of a convex polytope $\mathcal{P} \subset \mathcal{X}$. For this polytope, we assume the existence of an efficient **linear minimization oracle** (**lmo**), which means that for any $a \in \mathcal{X}^*$, one can efficiently solve

$$\mathbf{lmo}_{\mathcal{P}}(a) \in \arg \min_{x \in \mathcal{P}} \langle a, x \rangle.$$

This is for example the case if \mathcal{P} is the polytope arising from LP relaxations of MAP-MRF problems in a tree-structured graph, where the above problem can be solved efficiently using dynamic programming.

- The function h^* is a convex function which allows to efficiently compute its **proximal map** (**prox**), which for any $\bar{y} \in \mathcal{Y}$ and $\tau > 0$ is defined as

$$\mathbf{prox}_{\tau h^*}(\bar{y}) = \arg \min_{y \in \mathcal{Y}} \frac{1}{2\tau} \|y - \bar{y}\|^2 + h^*(y).$$

Important examples of h^* which allow for an efficient proximal map include quadratic functions and various norms. If $h^* = \delta_C$ i.e. the indicator functions of some convex set C the proximal map reduces the orthogonal projection operator.

An important special case of problem (1) is given by

$$\min_{x \in \mathbb{R}^d} \max_{y \in \mathbb{R}^d} \mathcal{L}(x, y) := f_{\mathcal{P}}(x) + \langle y, Ax - b \rangle \quad (2)$$

where A is a matrix and b is a vector of appropriate dimensions. This corresponds to the problem of minimizing $f_{\mathcal{P}}(x)$ subject to the linear constraint $Ax - b = 0$.

Primal and dual problems We denote by (x^*, y^*) a saddle point of problem (1), which satisfies

$$\mathcal{L}(x^*, y) \leq \mathcal{L}(x^*, y^*) \leq \mathcal{L}(x, y^*), \quad \forall (x, y) \in \mathcal{X} \times \mathcal{Y}.$$

Throughout the paper we denote the primal and dual problems respectively by

$$F(x) = \max_{y \in \mathcal{Y}} \mathcal{L}(x, y), \quad H(y) = \min_{x \in \mathcal{X}} \mathcal{L}(x, y).$$

We assume that strong duality holds, that is

$$H(y^*) = \max_{y \in \mathcal{Y}} H(y) = \max_{y \in \mathcal{Y}} \min_{x \in \mathcal{X}} \mathcal{L}(x, y) = \min_{x \in \mathcal{X}} F(x) = F(x^*)$$

Some of the results will also assume coercivity of a function. We recall that a proper function $\phi(z)$ is called coercive if

$$\lim_{\|z\| \rightarrow \infty} \phi(z) = \infty.$$

Contributions The algorithms we propose here are based on inexact proximal algorithms, which allow for an approximate evaluation of the proximal maps. For this we make use of efficient variants of the Frank-Wolfe algorithm that offer a linear convergence rate on the proximal subproblems. In summary, after $O(n \log n)$ calls to **lmo** we achieve the following guarantees.

- If function f is linear or quadratic, h^* is the indicator function of a linear constraint, and function $-H(y)$ is coercive then we obtain accuracy $O(1/n^2)$ on the dual objective H . If in addition the problem has the form of eq. (2) then we obtain accuracy $O(1/n^2)$ both on the primal gap $f_{\mathcal{P}}(x) - f_{\mathcal{P}}(x^*)$ and on the infeasibility gap $\|Ax - b\|$.

- In the most general case, we obtain accuracy $O(1/n)$ on the dual objective (and also on the primal objective assuming that $\text{dom } h^* \subseteq \mathcal{Y}$ is compact).

To the best of our knowledge these rates improve on the so far known rates for the class of saddle-point problems considered in this paper. In particular, for the problem in eq. (2) previous works (described later in Sec. 1.2) after n calls to `lmo` obtained accuracy $O(1/n)$ on the dual objective and accuracy $O(1/\sqrt{n})$ on both $f_{\mathcal{P}}(x) - f_{\mathcal{P}}(x^*)$ and $\|Ax - b\|$.

1.1 Motivating example

An important application, which also serves as the main motivation for the class of saddle-point problems studied in this paper, is given by the Lagrangian relaxation of discrete optimization problems. To form such relaxation, one needs to first encode discrete variables via Boolean indicator variables $X \in \{0, 1\}^d$, and then express a difficult optimization problem as a sum of tractable subproblems:

$$\min_{X \in \{0,1\}^d} \sum_{t \in T} f_t(X_{A_t}) \quad (3)$$

Here $T = \{1, \dots, m\}$ is the set of terms where each term t is specified by a subset of variables $A_t \subseteq [d]$ and a function $f_t : \{0, 1\}^{A_t} \rightarrow \mathbb{R} \cup \{+\infty\}$ of $|A_t|$ variables. Vector X_{A_t} is the restriction of vector $X \in \mathbb{R}^d$ to A_t . The arity $|A_t|$ of function f_t can be arbitrarily large, however we assume the existence of an efficient *min-oracle* that for a given vector $Y \in \mathbb{R}^{A_t}$ computes $X \in \arg \min_{X \in \{0,1\}^{A_t}} [f_t(X) + \langle X, Y \rangle]$ together with the cost $f_t(X)$. For example, this holds if $f_t(\cdot)$ corresponds to a MAP-MRF inference problem in a tree-structured graph.

It is now easy to define a relaxation of (3) that will have the form of eq. (1) (see e.g. [44]). First, for each $t \in T$ let us define polytope $\mathcal{P}^t = \text{conv}(\{X f(X) \mid X \in \text{dom } f_t\}) \subseteq \mathbb{R}^{A_t} \times R$ where $\text{dom } f_t = \{X \in \{0, 1\}^{A_t} \mid f_t(X) < +\infty\}$ is the effective domain of f_t . Problem (3) can then be equivalently written as

$$\min_{X \in \{0,1\}^d, X^1 \in \mathcal{P}^1, \dots, X^m \in \mathcal{P}^m} \sum_{t \in T} X_{\circ}^t \quad \text{s.t.} \quad X_v = X_v^t \quad \forall t \in T, v \in A_t \quad (4)$$

where X_{\circ}^t denotes the last component of vector $X^t \in \mathbb{R}^{A_t} \times R$. By dropping constraint $X \in \{0, 1\}^d$, dualizing equalities $X_v^t - X_v = 0$ with a Lagrange multiplier Y_v^t , and eliminating variables X , we obtain the following saddle point problem which is a relaxation of (4) (its optimal value is a lower bound on (4)):

$$\min_{x=(X^1, \dots, X^m) \in \mathcal{P}^1 \times \dots \times \mathcal{P}^m} \max_{y=(Y_v^t)_{t \in T, v \in A_t} \in \mathcal{Y}} \sum_{t \in T} X_{\circ}^t + \sum_{t \in T, v \in A_t} X_v^t Y_v^t \quad (5)$$

where \mathcal{Y} is the set of vectors $y = (Y_v^t)_{t \in T, v \in A_t}$ satisfying constraints $\sum_{t: v \in A_t} Y_v^t = 0$ for all $v \in [d]$. This problem has the form of eq. (1) where function $f(x)$ is linear and $h^*(y)$ is the indicator function of a linear constraint on y .

As an example, the MAP-MRF inference problem on an undirected graph can be cast in the framework above by decomposing the graph into tree-structured subproblems. It is well-known that the Lagrangian relaxation is equivalent to a standard LP relaxation, aka the local polytope relaxation [19, 34]. MAP-MRF problems find numerous applications in machine learning and computer vision [7]. In more recent work, they also appears as the final inference layer in deep convolutional neural networks [18].

1.2 Related work

Saddle point problems in the form of (1) can be solved by a large number of proximal primal-dual algorithms (see for example the recent work [10] for a very comprehensive overview) as soon as the

proximal maps for both the primal and dual functions can be solved efficiently. On the other hand, Gidel *et al.* proposed in [13] an extension of the Frank-Wolfe algorithm to saddle-point problems $\min_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} \mathcal{L}(x, y)$ by assuming the existence of an efficient linear minimization oracle for the product space $\mathcal{X} \times \mathcal{Y}$ (which is assumed to be a compact set). In this paper, we are assuming the existence of an efficient linear minimization oracle just on the primal and an efficient proximal map on the dual. Therefore, our algorithms somewhat stand between the two aforementioned techniques.

Our proposed algorithms rely on the inexact accelerated proximal gradient algorithm of Aujol and Dossal [2] and the inexact primal-dual algorithm of Rasch and Chambolle [29, Section 3.1]. Note that the first method only generates a dual sequence $\{y_n\}$. We extend the method and the analysis to also generate a primal sequence $\{x_n\}$, which is needed to solve the saddle problem (1).

Several authors studied a special case of (1) given in (2), or equivalently the problem of minimizing function $f_{\mathcal{P}}(x) = f(x) + \delta_{\mathcal{P}}(x)$ subject to linear constraints $Ax = b$ [14, 23, 46]. (The last paper actually considered a more general class of saddle problems). Papers [23, 46] achieve an accuracy of $O(n^{-1/2})$ after n iterations on the primal and infeasibility gaps, where [46] uses one `lmo` call per iteration and [23] uses $O(k^2)$ `lmo` calls at k -th iteration assuming that a standard Frank-Wolfe method is employed. Note, the papers above do not give bounds on suboptimality gaps of the dual function H ; instead, [14, 23] bound residuals of the *augmented* Lagrangian, which is not directly related to the residuals of the Lagrangian in eq. (2). A similar but slightly more general class of composite optimization problems was also recently considered in [42]. In a setting similar to our paper (∇f is Lipschitz continuous) they show $O(n^{-1/3})$ accuracy on Lagrangian values after n calls to `lmo`.

Frank-Wolfe algorithms for saddle-point problems have also been used in [1, 22]. The former paper achieved a $O(n^{-1/2})$ convergence rate on a rather general class of constrained optimization problems. The paper [22] shares some high-level similarities with our approach (such as solving smoothed primal subproblem to a given accuracy with Frank-Wolfe), but uses a different smoothing strategy that requires both primal and dual domains to be compact. This assumption rules out many interesting applications, including the one considered in Section 1.1.

There is a large body of literature on the special case of problem (1) corresponding to Lagrangian relaxation of discrete optimization problems, see e.g. [43, 36, 16, 30, 17, 32, 38, 19, 25, 33, 24, 40, 39, 44]. Some of these methods apply only to MAP inference problems in pairwise (or low-order) graphical models, because they need to compute marginals in tree-structured subproblems [16, 17, 33] or because they explicitly exploit the fact that the relaxation can be described by polynomial many constraints [38, 25, 40, 39]. The papers [17, 32] obtained accuracy $O(1/n)$ on the dual objective after n iterations, by applying accelerated gradient methods [26].

The first method that we develop can be viewed as an extension of the technique in [44], which applied an inexact proximal point algorithm (PPA) to the dual objective. In contrast to [44], we apply an accelerated version of inexact PPA, specify to which accuracy the subproblems need to be solved, and analyze the convergence rate.

1.3 Notation for approximate solutions and organization of the paper

We introduce the following notation for a function ϕ and accuracy $\varepsilon \geq 0$:

$$\begin{aligned} z \approx_{\varepsilon} \arg \min_z \phi(z) &\Leftrightarrow \phi(z) \leq \min_z \phi(z) + \varepsilon \\ z \approx_{\varepsilon} \arg \max_z \phi(z) &\Leftrightarrow \phi(z) \geq \max_z \phi(z) - \varepsilon \\ z \approx_{\varepsilon} \mathbf{prox}_{\tau\phi}(\bar{z}) &\Leftrightarrow z \approx_{\varepsilon} \arg \min_z \left[\phi(z) + \frac{1}{2\tau} \|z - \bar{z}\|^2 \right] \end{aligned}$$

The rest of the paper is organized as follows. The next section describes Frank-Wolfe algorithms for minimizing a smooth function over a convex polytope. Then in Section 3 we will present our first

approach, which is based on an inexact accelerated proximal point algorithm on the dual problem. In Section 4 we present our second approach, which is based on an inexact proximal primal-dual algorithm and directly solves the saddle point problem. Preliminary numerical results are given in Section 5. To make the paper more self-contained, Appendix A discusses some important results of convex optimization. Additionally, for a better readability, the more technical proofs are all moved to Appendices B-E.

2 Frank-Wolfe algorithms

Frank-Wolfe style algorithms is a class of algorithms for minimizing functions $g_{\mathcal{P}} : \mathcal{X} \rightarrow \mathbb{R}$ of the form $g_{\mathcal{P}}(x) = g(x) + \delta_{\mathcal{P}}(x)$ where g a convex continuously differentiable function with a Lipschitz continuous gradient and \mathcal{P} is a convex polytope. They are typically iterative techniques that work by applying a certain procedure $\text{FWstep}(x; g_{\mathcal{P}}) \mapsto x'$ where $g_{\mathcal{P}}$ is the objective function, and x and x' are respectively the old and the new iterates with $g_{\mathcal{P}}(x') \leq g_{\mathcal{P}}(x)$. We will apply such steps to functions $g_{\mathcal{P}}$ that change from time to time, which is why $g_{\mathcal{P}}$ is made a part of the notation.

It will be convenient to denote $g_{\mathcal{P}}^{\downarrow}(x) = g_{\mathcal{P}}(x) - \min_{x \in \mathcal{X}} g_{\mathcal{P}}(x)$ to be a shifted version of $g_{\mathcal{P}}$ with $\min_{x \in \mathcal{X}} g_{\mathcal{P}}^{\downarrow}(x) = 0$. The following fact is known (see [20]). For completeness, in Appendix D.1 we give a proof of the second inequality, expanding some derivations that were omitted in the proof of [20, Theorem 2].

Lemma 1 ([20]). *For a point $\hat{x} \in \mathcal{P}$ denote $\text{gap}^{\text{FW}}(\hat{x}; g_{\mathcal{P}}) = \langle \nabla g(\hat{x}), \hat{x} - s \rangle$ where $s = \text{mo}_{\mathcal{P}}(\nabla g(\hat{x}))$. Then*

$$g_{\mathcal{P}}^{\downarrow}(\hat{x}) \leq \text{gap}^{\text{FW}}(\hat{x}; g_{\mathcal{P}}) \leq \begin{cases} g_{\mathcal{P}}^{\downarrow}(\hat{x}) + LD^2/2 & \text{if } g_{\mathcal{P}}^{\downarrow}(\hat{x}) > LD^2/2 \\ D\sqrt{2L \cdot g_{\mathcal{P}}^{\downarrow}(\hat{x})} & \text{if } g_{\mathcal{P}}^{\downarrow}(\hat{x}) \leq LD^2/2 \end{cases}$$

where D is the diameter of \mathcal{P} and $L = L_g$ is the Lipschitz constant of ∇g .

While the original FW algorithm has a sublinear convergence rate, there are several variants that achieve a linear convergence rate under some assumptions on g . Examples include *Frank-Wolfe with away steps* (AFW) [20], *Decomposition-invariant Conditional Gradient* (DiCG) [12, 3], and *Blended Conditional Gradient* (BCG) [8]. Each step in these methods is classified as either *good* or *bad*. Good steps are guaranteed to decrease $g_{\mathcal{P}}^{\downarrow}(x)$ by a constant factor. Bad steps do not have such guarantee (because they hit the boundary of the polytope), but they make x “sparser” in a certain sense and thus cannot happen too often.

More formally, consider a class of functions \mathfrak{F} where each function $g_{\mathcal{P}} \in \mathfrak{F}$ is associated with a parameter vector $\Theta_{g_{\mathcal{P}}} \in \mathbb{R}^p$, and $\mathcal{P} = \text{dom } g_{\mathcal{P}}$ is the same for all $g_{\mathcal{P}} \in \mathfrak{F}$. We say that procedure FWstep has a *linear convergence rate on \mathfrak{F}* if there exist continuous function $\theta : \mathbb{R}^p \rightarrow (0, 1)$ and integers $R_0, R_1 \geq 0$ with the following properties: (i) if the step $\text{FWstep}(x; g_{\mathcal{P}}) \mapsto x'$ for $g_{\mathcal{P}} \in \mathfrak{F}$ is good then $g_{\mathcal{P}}^{\downarrow}(x') \leq \theta(\Theta_{g_{\mathcal{P}}}) \cdot g_{\mathcal{P}}^{\downarrow}(x)$; (ii) when applying $\text{FWstep}(x; g_{\mathcal{P}})$ iteratively to some initial vector x_0 (possibly for different functions $g_{\mathcal{P}} \in \mathfrak{F}$), at any point we have $N_{\text{bad}} \leq R_0 + R_1 N_{\text{good}}$ where N_{good} and N_{bad} are respectively the numbers of good and bad steps.

We will consider two classes of functions:

- $\mathfrak{F}_{\text{strong}} = \{g_{\mathcal{P}}(x) = g(x) + \delta_{\mathcal{P}}(x) : g \text{ is a strongly convex differentiable function with a Lipschitz-continuous gradient, with } \Theta_{g_{\mathcal{P}}} = (\mu_g, L_g)\}$.
- $\mathfrak{F}_{\text{weak}} = \{g_{\mathcal{P}}(x) = g(Ex) + \langle b, x \rangle + \delta_{\mathcal{P}}(x) : g \text{ is a strongly convex differentiable function with a Lipschitz-continuous gradient, and } E, b \text{ are matrix and vector of appropriate dimensions, with } \Theta_{g_{\mathcal{P}}} = (\mu_g, L_g, E, b)\}$.

Note that class $\mathfrak{F}_{\text{strong}}$ is implicitly parameterized by the dimension of vector x , and class $\mathfrak{F}_{\text{weak}}$ is implicitly parameterized by the dimensions of vector x and matrix E .

The AFW method is known to have linear convergence on $\mathfrak{F}_{\text{strong}}$ [20] and also on $\mathfrak{F}_{\text{weak}}$ [4, 20]. From the result of [4, 20] it is easy to deduce that the DiCG method with away steps also has linear convergence on $\mathfrak{F}_{\text{weak}}$, using [3, Property 1]. The BCG method [8] has been shown to have linear convergence on class $\mathfrak{F}_{\text{strong}}$.

Remark 1. *Some of the techniques above maintain some additional information about current iterate x . In particular, AFW and BCG represent x as a convex combination of “atoms” (vertices of \mathcal{P}): $x = \sum_i \alpha_i a_i$ where $\alpha_i \geq 0$, $\sum_i \alpha_i = 1$ and a_i are atoms. Coefficients α_i are updated together with x . For brevity, we omitted this from the notation.*

Remark 2. *The claims about the number of bad steps are proven in [20, 12, 3, 8] assuming that the function $g_{\mathcal{P}}$ is fixed. However, the proofs only use “structural” properties of current iterate x ; they are easily extended to the case when $g_{\mathcal{P}}$ is changing, as long as \mathcal{P} is fixed.*

Iterative application of FWstep Procedure FWstep can be used in a natural way to solve problems $x \approx_{\varepsilon} \arg \min_x g_{\mathcal{P}}(x)$ up to desired accuracy ε .

Algorithm 1: Algorithm $\text{FW}_{\varepsilon}(x; g_{\mathcal{P}})$.

Output: vector $x' \approx_{\varepsilon} \arg \min_x g_{\mathcal{P}}(x)$.

```

1 while true do
2   update  $x \leftarrow \text{FWstep}(x; g_{\mathcal{P}})$ 
3   if  $\text{gap}^{\text{FW}}(x; g_{\mathcal{P}}) \leq \varepsilon$  then return  $x$ 

```

Proposition 2. *Suppose that procedure FWstep has a linear convergence rate on class \mathfrak{F} that contains $g_{\mathcal{P}}$. Then,*

(a) *The number of good steps made during $\text{FW}_{\varepsilon}(x_0; g_{\mathcal{P}})$ satisfies*

$$N_{\text{good}} \leq \log_{1/\theta(\Theta_{g_{\mathcal{P}}})} \frac{g_{\mathcal{P}}^{\downarrow}(x_0)}{\min\{\frac{1}{2}LD^2, \frac{1}{2L}(\frac{\varepsilon}{D})^2\}} \quad (6)$$

where D is the diameter of \mathcal{P} and $L > 0$ is any constant satisfying $L \geq L_g$.

(b) *Suppose that $g_{\mathcal{P}} \in \tilde{\mathfrak{F}} \subseteq \mathfrak{F}$ where $\sup_{g_{\mathcal{P}} \in \tilde{\mathfrak{F}}, x \in \mathcal{P}} g_{\mathcal{P}}^{\downarrow}(x) < \infty$, $\sup_{g_{\mathcal{P}} \in \tilde{\mathfrak{F}}} L_g < \infty$, and $\{\Theta_{g_{\mathcal{P}}} \mid g_{\mathcal{P}} \in \tilde{\mathfrak{F}}\}$ is a compact subset of \mathbb{R}^p . Then $N_{\text{good}} = O(\log \frac{1}{\varepsilon})$ where the constant in the $O(\cdot)$ notation depends on $\tilde{\mathfrak{F}}$.*

Proof. (a) By the definition of linear convergence, after the given number of good steps we obtain vector x satisfying $g_{\mathcal{P}}^{\downarrow}(x) \leq \min\{\frac{1}{2}LD^2, \frac{1}{2L}(\frac{\varepsilon}{D})^2\}$. By Lemma 1, such x satisfies $\text{gap}^{\text{FW}}(x; g_{\mathcal{P}}) \leq \varepsilon$, and therefore the algorithm will immediately terminate.

(b) Since set $\{\Theta_{g_{\mathcal{P}}} \mid g_{\mathcal{P}} \in \tilde{\mathfrak{F}}\} \subseteq \mathbb{R}^p$ is compact and function $\theta : \mathbb{R}^p \rightarrow (0, 1)$ is continuous, there exists $\theta^* \in (0, 1)$ such that $\theta(\Theta_{g_{\mathcal{P}}}) \leq \theta^*$ for all $g_{\mathcal{P}} \in \tilde{\mathfrak{F}}$. Thus, all quantities present in (6) (except for ε) are bounded by constants for all $g_{\mathcal{P}} \in \tilde{\mathfrak{F}}$. The claim follows. \square

3 First approach: dual proximal point algorithm

The first approach that we consider is a proximal point algorithm (PPA) applied to the dual problem:

$$\max_{y \in \mathcal{Y}} \left\{ H(y) := \min_{x \in \mathcal{X}} \mathcal{L}(x, y) \right\}.$$

For a point $\bar{y} \in \mathcal{Y}$ and a smoothing parameter $\gamma > 0$, we let

$$\mathcal{L}_{\gamma, \bar{y}}(x, y) = \mathcal{L}(x, y) - \frac{1}{2\gamma} \|y - \bar{y}\|^2,$$

which can be seen as the original saddle-point problem, but with an additional proximal regularization on the dual variable. In each iteration, the PPA solves a maximization problem of the following form:

$$\hat{y} = \arg \max_{y \in \mathcal{Y}} \left\{ H_{\gamma, \bar{y}}(y) := \min_{x \in \mathcal{X}} \mathcal{L}_{\gamma, \bar{y}}(x, y) = H(y) - \frac{1}{2\gamma} \|y - \bar{y}\|^2 \right\}.$$

Based on our structure, it will be beneficial to first solve for \hat{x} and then to solve for \hat{y} via its proximal map, that is

$$\hat{x} = \arg \min_{x \in \mathcal{X}} \left\{ F_{\gamma, \bar{y}}(x) := \max_{y \in \mathcal{Y}} \mathcal{L}_{\gamma, \bar{y}}(x, y) \right\}, \quad \hat{y} = \arg \max_{y \in \mathcal{Y}} \mathcal{L}_{\gamma, \bar{y}}(\hat{x}, y) = \mathbf{prox}_{\gamma h^*}(\bar{y} + \gamma K \hat{x}).$$

Note that also for the smoothed saddle-point problem, strong duality holds,

$$\min_{x \in \mathcal{X}} F_{\gamma, \bar{y}}(x) = \min_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} \mathcal{L}_{\gamma, \bar{y}}(x, y) = \max_{y \in \mathcal{Y}} H_{\gamma, \bar{y}}(y),$$

and hence each step of the PPA can be equivalently written as minimizing the primal-dual gap

$$(\hat{x}, \hat{y}) = \arg \min_{(x, y) \in \mathcal{X} \times \mathcal{Y}} F_{\gamma, \bar{y}}(x) - H_{\gamma, \bar{y}}(y) \quad (7)$$

It is a well-known fact that the basic proximal point algorithm can be accelerated to achieve a $O(1/n^2)$ convergence rate [15, 31], which follows from the fact that the PPA can be seen as a steepest descent on the Moreau envelope (see Appendix A), which has a Lipschitz continuous gradient and hence can be accelerated using the technique of Nesterov [26].

However, based on our general assumptions on the problem (1), we will not be able to solve the proximal subproblems (7) exactly but only up to a certain error $\varepsilon > 0$ that is

$$(\hat{x}, \hat{y}) \approx_{\varepsilon} \arg \min_{(x, y) \in \mathcal{X} \times \mathcal{Y}} F_{\gamma, \bar{y}}(x) - H_{\gamma, \bar{y}}(y),$$

which clearly implies that $\hat{x} \approx_{\varepsilon} \arg \min_{x \in \mathcal{X}} F_{\gamma, \bar{y}}(x)$ as well as $\hat{y} \approx_{\varepsilon} \arg \max_{y \in \mathcal{Y}} H_{\gamma, \bar{y}}(y)$. However, we can still apply the recently proposed inexact accelerated proximal gradient algorithm of Aujol and Dossal [2], that can handle such approximation while still achieving an optimal $O(1/n^2)$ convergence rate on the dual objective. Note that the original method given in [2] only generates the dual sequence $\{y_n\}$ but in Algorithm 2 below we also keep the primal sequence $\{x_n\}$ which is needed to obtain a solution of the original saddle-point problem (1). Therefore, the algorithm below can also be seen as a generalization for solving saddlepoint problems.

Algorithm 2: Approximate accelerated proximal gradient algorithm.

1 choose nonnegative sequences $\{t_n\}$, $\{\varepsilon_n\}$ so that $t_1 = 1$ and $\rho_n \stackrel{\text{def}}{=} t_{n-1}^2 - t_n^2 + t_n > 0$ for all $n \geq 2$

2 choose initial point $y_0 \in \mathcal{Y}$, set $\bar{y}_0 = y_0$

3 **for** $n = 1, 2, \dots$ **do**

4

$$(x_n, y_n) \approx_{\varepsilon_n} \arg \min_{(x, y) \in \mathcal{X} \times \mathcal{Y}} F_{\gamma, \bar{y}_{n-1}}(x) - H_{\gamma, \bar{y}_{n-1}}(y) \quad (8)$$

$$\bar{y}_n = y_n + \frac{t_n - 1}{t_{n+1}} (y_n - y_{n-1}) \quad (9)$$

In order to analyze this algorithm, let us introduce the following quantities:

$$u_0 = y_0, \quad u_n = y_{n-1} + t_n(y_n - y_{n-1}) \quad \forall n \geq 1 \quad (10)$$

$$A_n = \sum_{k=1}^n t_k \sqrt{2\gamma\varepsilon_k} \quad (11)$$

$$B_n = \sum_{k=1}^n \gamma t_k^2 \varepsilon_k \quad (12)$$

$$W_n = t_n^2 [H(y^*) - H(y_n)] + \sum_{k=2}^n \rho_k [H(y^*) - H(y_{k-1})] \quad (13)$$

$$T_n = t_n^2 + \sum_{k=2}^n \rho_k = \sum_{k=1}^n t_k \quad (14)$$

First, we recall the following result from [2].

Theorem 3 ([2]). *For any $n \geq 1$ there holds*

$$W_n + \frac{1}{2\gamma} \|u_n - y^*\|^2 \leq \frac{C_n^*}{2\gamma} \quad (15)$$

where

$$C_n^* = \|y_0 - y^*\|^2 + 2A_n \left(\|y_0 - y^*\| + 2A_n + \sqrt{2B_n} \right) + 2B_n \quad (16)$$

$$\leq \left(\|y_0 - y^*\| + 2A_n + \sqrt{2B_n} \right)^2 \quad (17)$$

This theorem immediately implies the following results. (Note, some of the statements below are slightly modified versions of statements from [2], but follow exactly the same proofs).

Corollary 4 ([2]). *Suppose that sequences $\{A_n\}$ and $\{B_n\}$ are bounded. Then*

(a) $H(y^*) - H(y_n) = O(1/t_n^2)$.

(b) $H(y^*) - H(y_n^e) = O(1/T_n)$ where $y_n^e = (t_n^2 y_n + \sum_{k=2}^n \rho_k y_{k-1})/T_n$.

(c) If function $-H(y)$ is coercive then sequence $\{y_n\}$ is bounded, and $\|y_n - y_{n-1}\| = O(1/t_n)$.

Note that the rate of convergence in Corollary 4 depends on the choice of sequence $\{t_n\}$. These are some of the choices that have appeared in the literature:

- PPA: $t_n = 1$ for all $n \geq 1$. Then $T_n = \Theta(n)$.
- Nesterov [26, 6]: $t_{n+1} = (1 + \sqrt{1 + 4t_n^2})/2$ for $n \geq 1$. Then $t_n = \Theta(n)$ and $T_n = \Theta(n^2)$.
- Aujol-Dossal [2]: $t_n = \left(\frac{n+a-1}{a}\right)^d$ with $d \in (0, 1]$ and $a > \max\{1, (2d)^{1/d}\}$. Then $t_n = \Theta(n^d)$ and $T_n = \Theta(n^{d+1})$.

When stating complexities, we will implicitly assume below that either the second case or the third case with $d = 1$ is used, meaning that $t_n = \Theta(n)$ and $T_n = \Theta(n^2)$.

We now generalize Theorem 3 to the situation in this section. This generalization is somewhat analogous to the generalization obtained by Tseng [45] (for a different Nesterov-type algorithm and with a different proof).

Theorem 5. *Denote $x_n^e = \sum_{k=1}^n t_k x_k / T_n$. For any $y \in \mathcal{Y}$ and any $n \geq 1$ there holds*

$$T_n [\mathcal{L}(x_n^e, y) - H(y^*)] + W_n + \frac{1}{2\gamma} \|u_n - y\|^2 \leq \frac{C_n(y)}{2\gamma} \quad (18)$$

where

$$C_n(y) = \|y_0 - y\|^2 + 2A_n \left(\|y - y^*\| + \|y_0 - y^*\| + 2A_n + \sqrt{2B_n} \right) + 2B_n \quad (19)$$

Note that setting $y = y^*$ in Theorem 5 recovers Theorem 3, since in this case we have $\mathcal{L}(x_n^e, y^*) \geq H(y^*)$ and $C_n(y^*) = C_n^*$.

Corollary 6. *Suppose that sequences $\{A_n\}, \{B_n\}$ are bounded and $\text{dom } h^* \subseteq \mathcal{Y}$ is a compact set. Then $F(x_n^e) - F(x^*) = O(1/T_n) = O(1/n^2)$.*

Proof. By the assumption of the corollary, quantity $C_n(y)$ is bounded for any $y \in \text{dom } h^*$ and $n \geq 1$. We also have $F(x_n^e) = \max_{y \in \text{dom } h^*} \mathcal{L}(x_n^e, y)$ and $F(x^*) = H(y^*)$. The claim now follows directly from Theorem 5. \square

Next, we analyze the special case of problem (1) corresponding to constrained optimization problem $\min_{x \in \mathbb{R}^d} \{f_{\mathcal{P}}(x) \mid Ax = b\}$.

Theorem 7. *Suppose we are in the case of the saddle problem in eq. (2). (a) There holds*

$$f_{\mathcal{P}}(x_n^e) - f_{\mathcal{P}}(x^*) \leq \frac{C_n(0)}{2\gamma T_n} \quad (20)$$

$$\|Ax_n^e - b\| \leq \sqrt{\frac{2 \max\{f_{\mathcal{P}}(x^*) - f_{\mathcal{P}}(x_n^e), 0\}}{\gamma T_n}} + \frac{\hat{C}_n}{\gamma T_n} \quad (21)$$

where

$$\hat{C}_n = \|y_0\| + A_n + \sqrt{\|y_0\|^2 + 2A_n \left(\|y^*\| + \|y_0 - y^*\| + 2A_n + \sqrt{2B_n} \right) + 2B_n}$$

(b) *There exists constant $\beta \geq 0$ such that for any $x \in \mathcal{P}$ we have $f_{\mathcal{P}}(x^*) - f_{\mathcal{P}}(x) \leq \beta \|Ax - b\|$.*

(c) *If sequences $\{A_n\}$ and $\{B_n\}$ are bounded then $f_{\mathcal{P}}(x_n^e) - f_{\mathcal{P}}(x^*) = O(1/n^2)$ and $\|Ax_n^e - b\| = O(1/n^2)$.*

Note that part (a) is derived directly from Theorem 5. In part (b) we crucially exploit the facts that \mathcal{P} is a polytope, the feasible set $\{x \in \mathcal{P} : Ax - b = 0\}$ is non-empty, and function f has a bounded gradient on \mathcal{P} . Part (c) is an easy consequence of parts (a) and (b).

3.1 Overall algorithm

In this section we fix n , and denote $\bar{y} = \bar{y}_{n-1}$ and $\varepsilon = \varepsilon_n$. In order to implement Algorithm 2 for solving the saddle-point problem (1), we need to specify how to solve subproblem (8) for vector \bar{y} up to accuracy ε :

$$(x_n, y_n) \approx_{\varepsilon} \arg \min_{(x,y) \in \mathcal{X} \times \mathcal{Y}} F_{\gamma, \bar{y}}(x) - H_{\gamma, \bar{y}}(y) \quad (22)$$

We will first compute $x_n \approx_{\varepsilon} \arg \min_{x \in \mathcal{X}} F_{\gamma, \bar{y}}(x)$ by invoking Algorithm 1 for function $F_{\gamma, \bar{y}}$, and then solve for y_n via its proximal map:

$$x_n = \text{FW}_{\varepsilon}(x_{n-1}; F_{\gamma, \bar{y}}) \quad y_n = \arg \max_{y \in \mathcal{Y}} \mathcal{L}_{\gamma, \bar{y}}(x_n, y) = \text{prox}_{\gamma h^*}(\bar{y} + \gamma K x_n)$$

(As we will see later, function $F_{\gamma, \bar{y}}$ has the form $F_{\gamma, \bar{y}}(x) = g(x) + \delta_{\mathcal{P}}(x)$ for some differentiable convex function g with a L_g -Lipschitz continuous gradient, and so Algorithm 1 is indeed applicable). By construction, vector x_n satisfies $\text{gap}^{\text{FW}}(x_n; F_{\gamma, \bar{y}}) \leq \varepsilon$. The following lemma thus implies that the pair (x_n, y_n) indeed solves problem (22).

Lemma 8. *Suppose that $\hat{x} \in \mathcal{P}$ and $\hat{y} = \arg \max_y \mathcal{L}_{\gamma, \bar{y}}(\hat{x}, y)$. Then $H_{\gamma, \bar{y}}(\hat{y}) \geq \mathcal{L}_{\gamma, \bar{y}}(\hat{x}, \hat{y}) - \varepsilon = F_{\gamma, \bar{y}}(\hat{x}) - \varepsilon$ where $\varepsilon = \text{gap}^{\text{FW}}(\hat{x}; F_{\gamma, \bar{y}})$.*

Next, we derive an explicit expression for function $F_{\gamma, \bar{y}}$ (which is needed for implementing the call $x_n = \text{FW}_\varepsilon(x_{n-1}; F_{\gamma, \bar{y}})$), and formulate sufficient conditions on \mathcal{L} that will guarantee that $F_{\gamma, \bar{y}} \in \mathfrak{F}_{\text{weak}}$ (this would yield a good bound on the complexity of Algorithm 2).

Recall that the function $F_{\gamma, \bar{y}}(x)$ is given by

$$F_{\gamma, \bar{y}}(x) = \max_{y \in \mathcal{Y}} \mathcal{L}_{\gamma, \bar{y}}(x, y) = \max_{y \in \mathcal{Y}} \langle Kx, y \rangle + f_{\mathcal{P}}(x) - h^*(y) - \frac{1}{2\gamma} \|y - \bar{y}\|^2.$$

We now show that it can be written as

$$F_{\gamma, \bar{y}}(x) = f_{\mathcal{P}}(x) + h_{\gamma, \bar{y}}(Kx)$$

with $h_{\gamma, \bar{y}}(Kx)$ a differentiable function with Lipschitz continuous gradient.

Lemma 9. *Let the function $h_{\gamma, \bar{y}}(Kx)$ be defined as*

$$h_{\gamma, \bar{y}}(Kx) = \max_{y \in \mathcal{Y}} \langle Kx, y \rangle - h^*(y) - \frac{1}{2\gamma} \|y - \bar{y}\|^2.$$

We have the following two representations:

$$h_{\gamma, \bar{y}}(Kx) = \frac{\gamma}{2} \|Kx\|^2 + \langle Kx, \bar{y} \rangle - m_{h^*}^\gamma(\bar{y} + \gamma Kx), \quad (23)$$

$$= m_h^{\gamma^{-1}}(\gamma^{-1} \bar{y} + Kx) - \frac{1}{2\gamma} \|\bar{y}\|^2. \quad (24)$$

where $m_{h^*}^\gamma$ is the Moreau envelope of h^* with smoothing parameter γ and $m_h^{\gamma^{-1}}$ is the Moreau envelope of h with smoothing parameter γ^{-1} .

Moreover, the function $h_{\gamma, \bar{y}}(Kx)$ is convex, continuously differentiable in x with a γL_K^2 -Lipschitz continuous gradient given by

$$\nabla_x h_{\gamma, \bar{y}}(Kx) = K^* \text{prox}_{\gamma h^*}(\bar{y} + \gamma Kx), \quad (25)$$

$$= \gamma K^* (\gamma^{-1} \bar{y} + Kx - \text{prox}_{\gamma^{-1} h}(\gamma^{-1} \bar{y} + Kx)). \quad (26)$$

In practical applications, we will mostly be interested in the situation where h^* is a linear constraint.

Lemma 10. *Let $\mathcal{X} = \mathcal{Y} = \mathbb{R}^\ell$, and let $h^*(y) = \delta_S(y)$ be a linear constraint of the form $S = \{y \in \mathbb{R}^\ell : Cy = d\}$ for some matrix C with full row rank and vector d . Then, the function $h_{\gamma, \bar{y}}(Kx)$ is a quadratic function of the form*

$$h_{\gamma, \bar{y}}(Kx) = \frac{\gamma}{2} \|Kx\|^2 + \langle Kx, \bar{y} \rangle - \frac{1}{2\gamma} \|C^*(CC^*)^{-1}(d - C(\bar{y} + \gamma Kx))\|^2.$$

Moreover, its gradient is a linear map given by

$$\nabla_x h_{\gamma, \bar{y}}(Kx) = K^* (\bar{y} + \gamma Kx + C^*(CC^*)^{-1}(d - C(\bar{y} + \gamma Kx))).$$

We therefore obtain the following sufficient condition where the functions $F_{\gamma, \bar{y}}(x)$ for any $\bar{y} \in \mathcal{Y}$ and any $\gamma > 0$ fall into the class $\mathfrak{F}_{\text{weak}}$.

Lemma 11. *Let $f(x)$ be a quadratic function of the form $f(x) = \frac{1}{2} \langle Qx, x \rangle + \langle q, x \rangle$, with a symmetric positive semidefinite matrix Q and vector q . Furthermore, let h^* satisfy the condition of Lemma 10. Then $\{F_{\gamma, \bar{y}} | \bar{y} \in \mathcal{Y}\} \subseteq \mathfrak{F}_{\text{weak}}$.*

Proof. First note that both $f(x)$ and $h_{\gamma, \bar{y}}(Kx)$ are quadratic functions. By completing the squares and ignoring constant terms, it follows that $F_{\gamma, \bar{y}}(x)$ can be written as $F_{\gamma, \bar{y}}(x) = \frac{1}{2}\|Ex\|^2 + \langle b, x \rangle + \delta_{\mathcal{P}}(x)$ for some matrix E and vector b , where matrix E may depend on γ but not on \bar{y} . \square

It remains to specify how to set the sequence $\{\varepsilon_n\}$. We want sequences $\{A_n\}$ and $\{B_n\}$ to be bounded; this can be achieved by setting $\varepsilon_n = \Theta(n^{-4-\delta})$ for some $\delta > 0$. With these choices, we obtain the main result of this section:

Theorem 12. *Suppose that function \mathcal{L} satisfies the precondition of Lemma 11, function $-H(y)$ is coercive, and procedure **FWstep** has a linear convergence rate on $\mathfrak{F}_{\text{weak}}$ (e.g. it is one step of **AFW** or **DiCG**). Then Algorithm 2 makes $O(n \log n)$ calls to **FWstep** during the first n iterations, and obtains iterates x_n^e and y_n^e satisfying $H(y^*) - H(y_n^e) = O(1/n^2)$. Furthermore, in the case of the problem in eq. (2) the iterates satisfy $f_{\mathcal{P}}(x_n^e) - f_{\mathcal{P}}(x^*) = O(1/n^2)$ and $\|Ax_n^e - b\| = O(1/n^2)$.*

Proof. By Lemma 11, all functions $F_{\gamma, \bar{y}_{n-1}}$ encountered during the algorithm belong to $\mathfrak{F}_{\text{weak}}$. Furthermore, vectors \bar{y}_{n-1} for $n \geq 1$ belong to a compact set, since the sequence $\{y_n\}$ (and thus the sequence $\{\bar{y}_n\}$) is bounded by Corollary 4(c). By Proposition 2(b) the number of good FW steps during n -th iteration is $O(\log \frac{1}{\varepsilon_n}) = O(\log n)$, and during the first n iterations is $\sum_{k=1}^n O(\log k) = O(n \log n)$. The number of bad FW steps is thus also $O(n \log n)$ by the definition of linear convergence and by the fact that the call $x_n = \text{FW}_{\varepsilon_n}(x_{n-1}; F_{\gamma, \bar{y}_{n-1}})$ is initialized with vector x_{n-1} . The remaining claims follow from Corollaries 4 and 6 and Theorem 7. \square

4 Second approach: primal-dual proximal algorithm

In this section we consider solving (1) without the restriction that h^* is the indicator function of a linear constraint. Therefore we make use of proximal primal-dual algorithms such as [9] which in each step of the algorithm need to compute proximal maps with respect to $f_{\mathcal{P}}$ and h^* . By our problem assumptions, the proximal map with respect to h^* is tractable but the proximal map with respect to $f_{\mathcal{P}}$ requires to solve for any $\bar{x} \in \mathcal{X}$ and $\tau > 0$ an optimization problem of the form

$$\text{prox}_{\tau f_{\mathcal{P}}}(\bar{x}) = \arg \min_{x \in \mathcal{P}} f(x) + \frac{1}{2\tau} \|x - \bar{x}\|^2.$$

We note the obvious fact that each proximal subproblem is a τ^{-1} -strongly convex function with $L_f + \tau^{-1}$ -Lipschitz continuous gradient over a convex polytope. Hence, it falls into the class $\mathfrak{F}_{\text{strong}}$, on which Frank-Wolfe algorithms achieve a linear rate of convergence. Similar to the previous section, we will not be able to solve the subproblems exactly but up to a certain accuracy $\varepsilon > 0$. We therefore need to resort to the recently proposed inexact primal-dual algorithm by Rasch and Chambolle [29, Section 3.1] which can handle such inaccuracy. The algorithm adapted to our situation is given below as Algorithm 3.

Algorithm 3: Inexact primal-dual algorithm.

- 1 choose $\tau, \sigma > 0$ such that $\sigma\tau L_K^2 < 1$
- 2 choose initial points $x_0 \in \mathcal{X}$ and $y_0 \in \mathcal{Y}$, set $x_{-1} = x_0$
- 3 **for** $n = 0, 1, \dots$ **do**

4

$$y_{n+1} = \text{prox}_{\sigma h^*}(y_n + \sigma K(2x_n - x_{n-1})) \quad (27)$$

$$x_{n+1} \approx_{\varepsilon_{n+1}} \text{prox}_{\tau f_{\mathcal{P}}}(x_n - \tau K^* y_{n+1}) \quad (28)$$

The following result has been shown in [29].¹

¹Part (a) is not formulated explicitly as a theorem in [29], but can be found on page 396 before Theorem 2. Part (b) appears as Theorem 2 in [29].

Theorem 13 ([29]).

(a) Define $x_n^e = \frac{1}{n} \sum_{k=1}^n x_k$ and $y_n^e = \frac{1}{n} \sum_{k=1}^n y_k$. Then for any $n \geq 1$ and $(x, y) \in \mathcal{X} \times \mathcal{Y}$

$$\mathcal{L}(x_n^e, y) - \mathcal{L}(x, y_n^e) \leq \frac{1}{n} \left(\frac{\|x - x_0\|^2}{2\tau} + \frac{\|y - y_0\|^2}{2\sigma} + \frac{\text{diam}(\mathcal{P})}{\tau} A_n + \frac{1}{\tau} B_n \right)$$

where

$$A_n = \sum_{k=1}^n \sqrt{2\tau\varepsilon_k} \qquad B_n = \sum_{k=1}^n \tau\varepsilon_k \qquad (29)$$

(b) If sequences $\{A_n\}$ and $\{B_n\}$ are bounded then there exists saddle point (x^*, y^*) of problem (1) such that $x_n \rightarrow x^*$ and $y_n \rightarrow y^*$.

To solve the subproblem in eq. (28), we call Algorithm 1 via

$$x_{n+1} \leftarrow \text{FW}_{\varepsilon_{n+1}}(x_n; g_n), \quad g_n(x) = f_{\mathcal{P}}(x) + \frac{1}{2\tau} \|x - (x_n - \tau K^* y_{n+1})\|^2$$

To make sequences $\{A_n\}$ and $\{B_n\}$ bounded, we can set $\varepsilon_n = \Theta(n^{-2-\delta})$ for some $\delta > 0$. With these choices, we obtain

Theorem 14. Suppose that procedure **FWstep** has a linear convergence rate on $\mathfrak{F}_{\text{strong}}$ (e.g. it is one step of AFW, DiCG or BCG). Then Algorithm 3 makes $O(n \log n)$ calls to **FWstep** during the first n iterations, and obtains iterates x_n^e and y_n^e satisfying $F(x_n^e) - F(x^*) = O(1/n)$ (if $\text{dom } h^*$ is a compact set) and $H(y^*) - H(y_n^e) = O(1/n)$.

Proof. Clearly, all proximal subproblems g_n encountered during the algorithm belong to $\mathfrak{F}_{\text{strong}}$. Furthermore, vectors $\bar{x}_{n+1} \stackrel{\text{def}}{=} x_n - \tau K^* y_{n+1}$ for $n \geq 0$ belong to a compact set (by Theorem 13(b)). By Proposition 2(b) the number of good FW steps during n -th iteration is $O(\log \frac{1}{\varepsilon_{n+1}}) = O(\log n)$, and during the first n iterations is $\sum_{k=1}^n O(\log k) = O(n \log n)$. The number of bad FW steps is thus also $O(n \log n)$ by the definition of linear convergence.

Now assume that $\text{dom } h^*$ is a compact set, and define $x \in \arg \min_{x \in \mathcal{P}} \mathcal{L}(x, y_n^e)$, $y \in \arg \max_{y \in \text{dom } h^*} \mathcal{L}(x_n^e, y)$.

We have $F(x_n^e) - F(x^*) = \mathcal{L}(x_n^e, y) - F(x^*) \leq \mathcal{L}(x_n^e, y) - H(y_n^e) = \mathcal{L}(x_n^e, y) - \mathcal{L}(x, y_n^e)$, so applying Theorem 13(a) for (x, y) gives $F(x_n^e) - F(x^*) = O(1/n)$.

It remains to prove the last claim. Since the sets $\{x_n\}$ and $\{y_n\}$ are bounded, we can assume w.l.o.g. that $\text{dom } h^*$ is a compact set (we can add indicator function $\delta_S(y)$ to $h^*(y)$ for a suitable set $S \subseteq \mathcal{Y}$ so that the algorithm and values $H(y^*)$, $H(y_n^e)$ are not affected). Using the same argument as before, we obtain $H(y^*) - H(y_n^e) = O(1/n)$. \square

5 Numerical Results

In this section, we show preliminary results for solving MRFs arising from computer vision. The goal is to solve the following discrete minimization problem:

$$\min_{X \in D^{\mathcal{V}}} E(X) := \sum_{i \in \mathcal{V}} \theta_i(X_i) + \sum_{ij \in \mathcal{E}} \theta_{ij}(X_i, X_j)$$

where $(\mathcal{V}, \mathcal{E})$ is a 4-connected 2D grid graph, D is a finite set of labels, and $\theta_i(\cdot), \theta_{ij}(\cdot, \cdot)$ are given unary and pairwise costs, respectively. We decompose the problem into horizontal and vertical chains, and convert it to the saddle point problem (1) as described in Section 1.1.

We compare two versions of Algorithm 2: accelerated proximal point algorithm (denoted as A-PPA) with the aggressive choice $t_n = (n + 1)/2$ (which corresponds to the Aujol-Dossal scheme

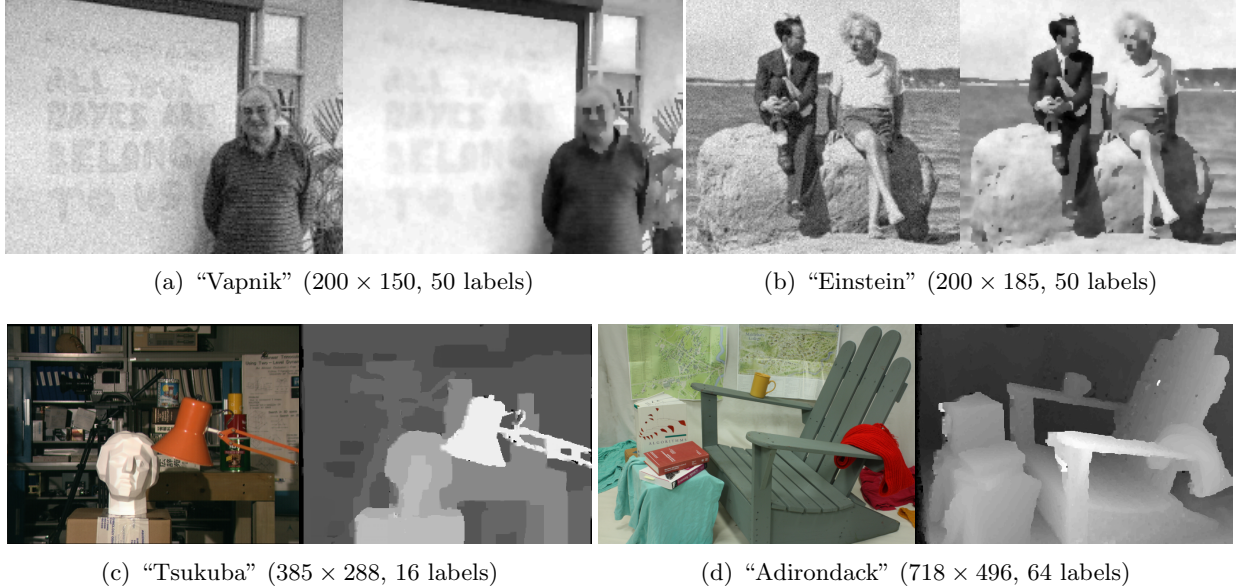


Figure 1: Results of the proposed algorithm for MRF-based image denoising (first row) and disparity estimation (second row)

with $d = 1, a = 2$), and the standard proximal point algorithm (denoted as PPA) which is realized by setting $t_n = 1$. Their convergence rates after n iterations are $O(1/n^2)$ and $O(1/n)$ respectively, assuming that sequences $\{A_n\}$ and $\{B_n\}$ are bounded. We invoke Algorithm 1 to minimize the functions $F_{\gamma, \bar{y}}$ up to accuracy

$$\varepsilon_n = \mathbf{gap}_0 \cdot n^{-\alpha}, \quad (30)$$

for a constant $\alpha > 0$, where \mathbf{gap}_0 denotes the initial gap of the function $F_{\gamma, \bar{y}}$. Note that for $O(1/n^2)$ convergence rate for A-PPA we would have to set $\alpha > 4$, but in practice smaller settings of α give larger speedups.

The smoothing parameter γ is chosen by hand, but as a general rule of thumb, it should be adapted to the scale of the primal respectively dual cost functions. Note that the selection of an optimal value of γ is beneficial mainly from a practical point of view as it does not influence the asymptotic convergence rate. We leave the selection of an optimal value of γ for future work.

In the case of PPA we tested two different versions. The first version is the method proposed in [44]. It uses a fixed number of Frank-Wolfe steps and is hence denoted as PPA-fw. The second version is motivated from our analysis for the case $t_n = 1$. It uses (30) for determining the accuracy of the subproblems and hence we denote it by PPA- α . Note that in order to guarantee an $O(1/n)$ convergence of PPA, one needs $\alpha > 2$ but as before, smaller values yield larger speedups.

Procedure `FWstep` was implemented as follows. We explicitly maintain the current iterate x as a convex combination of atoms. In the beginning of `FWstep` we first run a standard Frank-Wolfe step, and then re-optimize the objective over the current set of atoms. For that one needs to solve a low-dimensional strongly convex, quadratic subproblem over the unit simplex; we used the linearly converging accelerated proximal gradient method described in [27]. The resulting method can be viewed as a version of the BCG method [8], which is known to be linearly convergent on $\mathfrak{F}_{\text{strong}}$.

Remark 3. *Note that there is an extensive literature on FW variants. Potential alternatives to the method above include BCFW [21] and its variants [41, 28], DiCG [12, 3], and Frank-Wolfe with in-face directions [11]. However, testing different Frank-Wolfe methods is outside the scope of this paper; instead, we study what is the best way to use a given FW method.*

Next, we describe the two vision applications that we used.

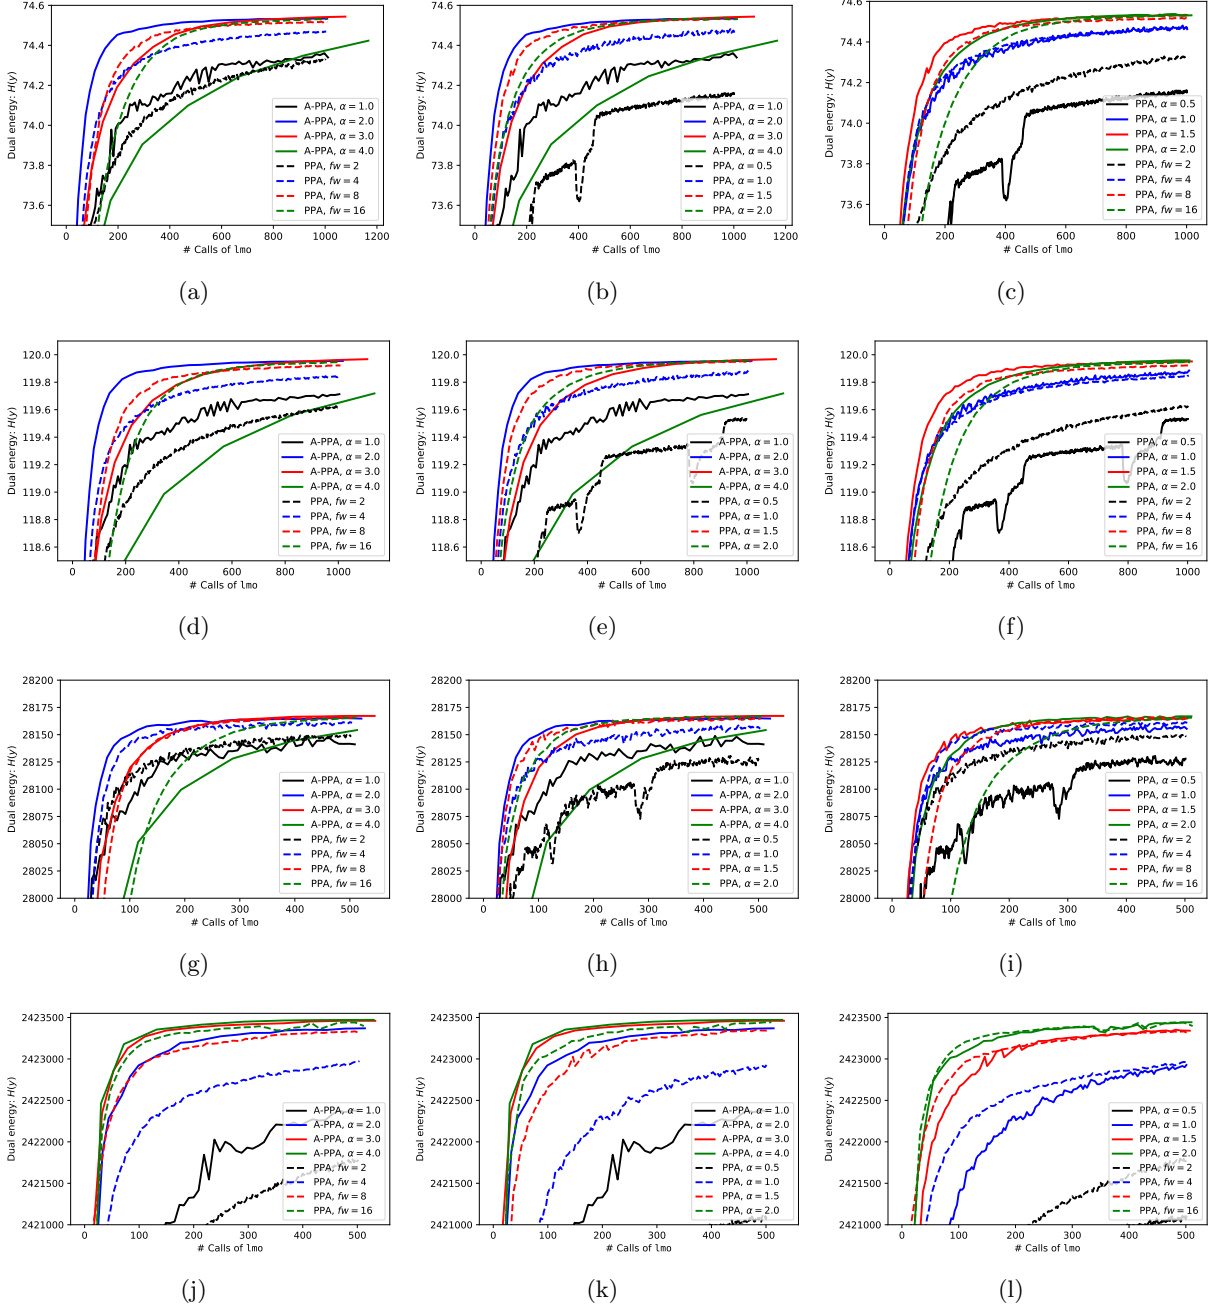


Figure 2: Convergence plots of the dual energy for “Vapnik” (first row), ”Einstein” (second row), “Tsukuba” (third row) and “Adirondack” (fourth row). See the text for more details.

Image denoising The first application is a simple Gaussian image denoising problem. The unary terms are given by a quadratic potential function and the pairwise terms are given by a truncated quadratic potential function. Hence, this model resembles a discrete version of the celebrated Mumford-Shah functional. Figure 1 (a) and (b) show the noisy input images together with their denoised versions. The value of γ was set to $\gamma = 1e - 03$ in both cases.

Stereo The second application is a classical disparity estimation problem from rectified stereo image pairs [35]. Figure 1 (c) and (d) show the left input images together with the estimated disparity images. For the “Tsukuba” image, the unaries are computed based on the absolute color differences between the left and right input image. The unaries of the “Adirondack” image are

computed using a CNN-based correlation network [18]. In both cases the pairwise costs are given by a truncated absolute potential function. The value of γ was set to $\gamma = 1e - 01$ in the case of “Tsukuba”, and $\gamma = 1.0$ in the case of “Adirondack”.

Results In Figure 2 we provide convergence plots of the three different algorithms: A-PPA, PPA-fw and PPA- α on all 4 examples. We plot the convergence of the dual energy $H(y)$ over the total number of calls of `lmo`.

The plots show that A-PPA systematically outperforms both the baseline method PPA-fw (first column) as well as its variant PPA- α (second column). The results also indicate that the choice of α significantly influences the global convergence behavior. For example in Figure 2 (a) $\alpha = 2$ yields the largest speedup in the first iterations but is caught up by $\alpha = 3$ in later iterations of the algorithm. This suggests, that a more flexible sequence ε_n has the potential for an additional speedup but such a development is left for future research.

From the plots in the third column one can also observe that the new variant PPA- α has a small speedup over the existing baseline method PPA-fw. This can be seen as a byproduct of our novel analysis in the case $t_n = 1$. Here one can also observe that values of $\alpha < 2$ yield the fastest convergence in the first iterations of the algorithm.

From a practical point of view, the plots show that α should be chosen according to the desired accuracy of the solution. This is of particular interest if one is only interested in a fast approximate solution to the problem, for example if the MRF is used as the last inference layer in a CNN.

6 Conclusion

In this work, we have proposed new primal-dual algorithms based on a mixture of proximal and Frank-Wolfe algorithms to solve a class of convex-concave saddle point problems arising in Lagrangian relaxations of discrete optimization problems. As our main result, we have shown after $O(n \log n)$ calls to `lmo` a $O(1/n)$ convergence rate in the most general case (Alg. 3) and a $O(1/n^2)$ convergence rate with certain regularity assumptions on the dual objective (Alg. 2). To the best of our knowledge, this improves on the known rates from the literature. Our preliminary numerical results also show an improved practical performance of Alg. 2 on MAP inference problems in computer vision. Note, we have not implemented yet Alg. 3 since its rate is worse on the application that we consider; at the moment the primary purpose of Alg. 3 is to show which rates are achievable for different classes of saddle point problems.

A Some results from convex optimization

In order to make the paper more self-contained, we recall here some well-known definitions and results from convex optimization, which will be useful later. The results with their proofs can for example be found in the recent monograph by Amir Beck [5].

The convex conjugate f^* of an extended real valued function f is defined as

$$f^*(y) = \sup_x \langle x, y \rangle - f(x), \quad (31)$$

which, since being a pointwise maximum over linear affine functions, is a convex lower semicontinuous function.

The infimal convolution between two proper functions f and g (written as $f \square g$) is defined by

$$(f \square g)(x) = \inf_u f(u) + g(x - u) = \inf_{x=u+v} f(u) + g(v). \quad (32)$$

which itself is a convex function, if f is proper and convex, and g is convex and real valued.

The infimal convolution and the convex conjugate for two proper functions f, g are related via

$$(f \square g)^* = f^* + g^*. \quad (33)$$

Moreover, assuming that f is proper and convex and g is a real-valued convex function, one also has that

$$(f + g)^* = f^* \square g^*. \quad (34)$$

The Moreau envelope of a convex function f with smoothness parameter $\mu > 0$ is the infimal convolution of f with μ^{-1} times a quadratic function.

$$m_f^\mu(x) = \inf_{x=u+v} f(u) + \frac{1}{2\mu} \|v\|^2 = \left(f \square \frac{1}{2\mu} \|\cdot\|^2 \right) (x). \quad (35)$$

Its unique minimizing argument is precisely given by the proximal map

$$\arg \min_u f(u) + \frac{1}{2\mu} \|x - u\|^2 = \mathbf{prox}_{\mu f}(x), \quad (36)$$

so that the Moreau envelope can explicitly be written as

$$m_f^\mu(x) = f(\mathbf{prox}_{\mu f}(x)) + \frac{1}{2\mu} \|x - \mathbf{prox}_{\mu f}(x)\|^2. \quad (37)$$

It is a well-known fact (and it easily follows from the monotonicity of the subdifferential), that the proximal map is a firmly nonexpansive operator, that is for any proper, convex lower semicontinuous function f it holds that

$$\|\mathbf{prox}_f(x) - \mathbf{prox}_f(y)\|^2 \leq \langle \mathbf{prox}_f(x) - \mathbf{prox}_f(y), x - y \rangle, \quad \forall x, y. \quad (38)$$

Furthermore, from the Cauchy-Schwartz inequality it also follows that the proximal map is a nonexpansive operator

$$\|\mathbf{prox}_f(x) - \mathbf{prox}_f(y)\| \leq \|x - y\|, \quad \forall x, y. \quad (39)$$

The Moreau envelope is continuously differentiable with an explicit representation of the gradient in terms of its proximal map

$$\nabla m_f^\mu(x) = \frac{1}{\mu} (x - \mathbf{prox}_{\mu f}(x)). \quad (40)$$

Furthermore, the gradient of the Moreau envelope is Lipschitz continuous with parameter μ^{-1} .

The next computation (which makes use of some results above) shows the relation between the Moreau envelopes of a convex function f and its convex conjugate f^* .

$$\begin{aligned} m_f^\mu(x) &= \inf_u f(u) + \frac{1}{2\mu} \|x\|^2 - \frac{1}{\mu} \langle x, u \rangle + \frac{1}{2\mu} \|u\|^2 \\ &= \frac{1}{2\mu} \|x\|^2 - \frac{1}{\mu} \sup_u \langle x, u \rangle - \mu f(u) - \frac{1}{2} \|u\|^2 \\ &= \frac{1}{2\mu} \|x\|^2 - \frac{1}{\mu} \left(\mu f(u) + \frac{1}{2} \|\cdot\|^2 \right)^* (x) \\ &= \frac{1}{2\mu} \|x\|^2 - \left(f^* \square \frac{\mu}{2} \|\cdot\|^2 \right) \left(\frac{x}{\mu} \right), \end{aligned}$$

from which it follows the beautiful relation

$$\frac{1}{2\mu} \|x\|^2 = m_f^\mu(x) + m_{f^*}^{\mu^{-1}} \left(\frac{x}{\mu} \right). \quad (41)$$

Now, taking the gradient with respect to x on both sides and multiplying by μ , recovers the celebrated Moreau identity

$$x = \mathbf{prox}_{\mu f}(x) + \mu \mathbf{prox}_{\frac{1}{\mu} f^*} \left(\frac{x}{\mu} \right). \quad (42)$$

B Analysis of Algorithm 2: proof of Theorem 5

B.1 Proof preliminaries

We first state a few technical results that will be used in the proof of Theorem 5.

Lemma 15. *Suppose that $(\hat{x}, \hat{y}) \approx_\varepsilon \arg \min_{x,y} [F_{\gamma, \bar{y}}(x) - H_{\gamma, \bar{y}}(y)]$. Then $\hat{y} \approx_{\varepsilon - \delta} \arg \max_y \mathcal{L}_{\gamma, \bar{y}}(\hat{x}, y)$ where $\delta = \mathcal{L}(\hat{x}, \hat{y}) - H(\hat{y}) \geq 0$.*

Proof. The claim follows from the following inequalities:

$$\max_y \mathcal{L}_{\gamma, \bar{y}}(\hat{x}, y) - \mathcal{L}_{\gamma, \bar{y}}(\hat{x}, \hat{y}) = F_{\gamma, \bar{y}}(\hat{x}) - \mathcal{L}_{\gamma, \bar{y}}(\hat{x}, \hat{y}) = F_{\gamma, \bar{y}}(\hat{x}) - H_{\gamma, \bar{y}}(\hat{y}) - \delta \leq \varepsilon - \delta$$

□

Lemma 16 ([37, Lemma 2], [2, Lemma 2.5]). *Suppose that $\hat{y} \approx_\varepsilon \text{prox}_{\gamma\phi}(\bar{y})$ for a proper convex l.s.c. function ϕ . Then there exists r with $\|r\| \leq \sqrt{2\gamma\varepsilon}$ such that*

$$\phi(z) \geq \phi(\hat{y}) + \left\langle \frac{\bar{y} - \hat{y} - r}{\gamma}, z - \hat{y} \right\rangle - \varepsilon \quad \forall z \in \mathcal{Y}$$

Lemma 17. *Suppose that $(\hat{x}, \hat{y}) \approx_\varepsilon \arg \min_{x,y} [F_{\gamma, \bar{y}}(x) - H_{\gamma, \bar{y}}(y)]$. Then*

$$H(\hat{y}) \geq \mathcal{L}(\hat{x}, z) + \frac{\|\hat{y} - z\|^2 - \|z - \bar{y}\|^2 + \|\hat{y} - \bar{y}\|^2}{2\gamma} - \varepsilon - \sqrt{\frac{2\varepsilon}{\gamma}} \cdot \|\hat{y} - z\| \quad \forall z \in \mathcal{Y}$$

Proof. Denote $\phi(z) = -\mathcal{L}(\hat{x}, z)$. By Lemma 15, we have $\hat{y} \approx_{\varepsilon - \delta} \text{prox}_{\gamma\phi}(\bar{y})$ where $\delta = \mathcal{L}(\hat{x}, \hat{y}) - H(\hat{y}) \geq 0$. By Lemma 16, there exists r with $\|r\| \leq \sqrt{2\gamma(\varepsilon - \delta)} \leq \sqrt{2\gamma\varepsilon}$ such that

$$\begin{aligned} -H(\hat{y}) &= \delta + \phi(\hat{y}) \leq \delta + \phi(z) - \left\langle \frac{\bar{y} - \hat{y} - r}{\gamma}, z - \hat{y} \right\rangle + \varepsilon - \delta \\ &= -\mathcal{L}(\hat{x}, z) - \frac{1}{\gamma} \langle \bar{y} - \hat{y}, z - \hat{y} \rangle + \frac{1}{\gamma} \langle r, z - \hat{y} \rangle + \varepsilon \\ &= -\mathcal{L}(\hat{x}, z) + \frac{1}{2\gamma} (\|\bar{y} - z\|^2 - \|\hat{y} - z\|^2 - \|\hat{y} - \bar{y}\|^2) + \frac{1}{\gamma} \langle r, z - \hat{y} \rangle + \varepsilon \end{aligned}$$

It remains to observe that $\langle r, z - \hat{y} \rangle \leq \|r\| \cdot \|\hat{y} - z\| \leq \sqrt{2\gamma\varepsilon} \cdot \|\hat{y} - z\|$. □

The following identity will also be useful (it follows from eq. (9) and (10)):

$$\bar{y}_n = \left(1 - \frac{1}{t_{n+1}}\right) y_n + \frac{1}{t_{n+1}} u_n \quad (43)$$

B.2 Proof of Theorem 5

Proof. Let us fix $y \in \mathcal{Y}$ (to be determined later). For brevity, denote $a_k = a_k(y) = \|u_k - y\|$ for $k \geq 0$. Now fix $k \geq 1$, and apply Lemma 17 to $\bar{y} = \bar{y}_{k-1}$, $\hat{y} = y_k$ and $z = (1 - \frac{1}{t_k})y_{k-1} + \frac{1}{t_k}y$. We get

$$H(y_k) \geq \mathcal{L}(x_k, z) + \frac{\|y_k - z\|^2 - \|z - \bar{y}_{k-1}\|^2}{2\gamma} - \varepsilon_k - \sqrt{\frac{2\varepsilon_k}{\gamma}} \cdot \|y - y_k\| \quad (44)$$

Using (43), it can be checked that $y_k - z = \frac{1}{t_k}(u_k - y)$ and $z - \bar{y}_{k-1} = \frac{1}{t_k}(y - u_{k-1})$. Furthermore, by concavity of function $\mathcal{L}(x_k, \cdot)$ we get

$$\begin{aligned} \mathcal{L}(x_k, y) &\geq \left(1 - \frac{1}{t_k}\right) \mathcal{L}(x_k, y_{k-1}) + \frac{1}{t_k} \mathcal{L}(x_k, y) \\ &\geq \left(1 - \frac{1}{t_k}\right) H(y_{k-1}) + \frac{1}{t_k} \mathcal{L}(x_k, y) \end{aligned}$$

Plugging these equations into (44) and multiplying by t_k^2 gives

$$t_k^2 H(y_k) \geq (t_k^2 - t_k) H(y_{k-1}) + t_k \mathcal{L}(x_k, y) + \frac{a_k^2 - a_{k-1}^2}{2\gamma} - t_k^2 \varepsilon_k - t_k \sqrt{\frac{2\varepsilon_k}{\gamma}} \cdot a_k$$

Let us sum this inequality over $k = 1, \dots, n$ and move terms with $H(\cdot)$ to the LHS. We obtain

$$t_n^2 H(y_n) + \sum_{k=1}^{n-1} \rho_{k+1} H(y_k) \geq \sum_{k=1}^n t_k \mathcal{L}(x_k, y) + \frac{a_n^2 - a_0^2}{2\gamma} - \sum_{k=1}^n t_k^2 \varepsilon_k - \sum_{k=1}^n t_k \sqrt{\frac{2\varepsilon_k}{\gamma}} \cdot a_k$$

Note that the LHS equals $T_n H(y^*) - W_n$. By convexity of function $\mathcal{L}(\cdot, y)$, we have $\sum_{k=1}^n t_k \mathcal{L}(x_k, y) \geq T_n \mathcal{L}(x_n^e, y)$. Therefore,

$$T_n H(y^*) - W_n \geq T_n \mathcal{L}(x_n^e, y) + \frac{\|u_n - y\|^2}{2\gamma} - \frac{\tilde{C}_n(y)}{2\gamma} \quad (45)$$

where we denoted

$$\tilde{C}_n(y) = \|y_0 - y\|^2 + 2 \sum_{k=1}^n \gamma t_k^2 \varepsilon_k + 2 \sum_{k=1}^n t_k \sqrt{2\gamma \varepsilon_k} \cdot \|u_k - y\|$$

Equivalently,

$$T_n [\mathcal{L}(x_n^e, y) - H(y^*)] + W_n + \frac{1}{2\gamma} \|u_n - y\|^2 \leq \frac{\tilde{C}_n(y)}{2\gamma} \quad (46)$$

□

Our next goal is to upper bound quantities $\|u_k - y^*\|$. We use the same argument as in [37, 2]. It relies on the following fact.

Lemma 18. [37, Lemma 1] *Assume that a nonnegative sequence $\{a_n\}$ satisfies $a_n^2 \leq S_n + \sum_{k=1}^n \lambda_k a_k$ where $\{S_k\}$ is a nondecreasing sequence, $S_0 \geq a_0^2$ and $\lambda_k \geq 0$ for all k . Then*

$$a_n \leq A_n + \sqrt{S_n + A_n^2}, \quad A_n = \frac{1}{2} \sum_{k=1}^n \lambda_k$$

Plugging $y = y^*$ into (46) and observing that $\mathcal{L}(x_n^e, y^*) - H(y^*) \geq 0$ gives

$$\|u_n - y^*\|^2 \leq \tilde{C}_n(y^*)$$

Therefore, the sequence $\{a_n\}$ defined via $a_n = \|u_n - y^*\|$ satisfies the precondition of Lemma 18 with

$$B_n = \sum_{k=1}^n \gamma t_k^2 \varepsilon_k \quad S_n = \|y_0 - y^*\|^2 + 2B_n \quad \lambda_k = 2t_k \sqrt{2\gamma \varepsilon_k} \quad A_n = \sum_{k=1}^n t_k \sqrt{2\gamma \varepsilon_k}$$

We can thus conclude

$$\|u_n - y^*\| \leq A_n + \sqrt{\|y_0 - y^*\|^2 + 2B_n + A_n^2} \leq \|y_0 - y^*\| + 2A_n + \sqrt{2B_n}$$

By observing that $\|u_k - y\| \leq \|u_k - y^*\| + \|y - y^*\|$ we obtain

$$\begin{aligned} \|u_k - y\| &\leq \|y - y^*\| + \|y_0 - y^*\| + 2A_k + \sqrt{2B_k} \\ &\leq \|y - y^*\| + \|y_0 - y^*\| + 2A_n + \sqrt{2B_n} \quad \forall 1 \leq k \leq n \end{aligned}$$

Plugging this bound into (45) gives the desired result:

$$\tilde{C}_n(y) \leq \|y_0 - y\|^2 + 2B_n + 2A_n \cdot \left(\|y - y^*\| + \|y_0 - y^*\| + 2A_n + \sqrt{2B_n} \right) = C_n(y)$$

C Algorithm 2 for solving $\min_x \{f_{\mathcal{P}}(x) \mid Ax = b\}$: proof of Theorem 7

C.1 Proof of Theorem 7(a)

(a) Let us fix $n \geq 1$. Using Theorem 5 and the strong duality assumption ($H(y^*) = f_{\mathcal{P}}(x^*)$), we obtain

$$f_{\mathcal{P}}(x_n^e) - f_{\mathcal{P}}(x^*) + \langle y, Ax_n^e - b \rangle \leq \frac{C_n(y)}{2\gamma T_n} \quad (47)$$

Plugging $y = 0$ gives eq. (20). Next, from (19) we obtain

$$\begin{aligned} C_n(y) &\leq \|y_0\|^2 + 2\|y_0\| \cdot \|y\| + \|y\|^2 + 2A_n \left(\|y\| + \|y^*\| + \|y_0 - y^*\| + 2A_n + \sqrt{2B_n} \right) + 2B_n \\ &= \|y\|^2 + P \cdot \|y\| + Q \end{aligned}$$

where we denoted

$$P = 2(\|y_0\| + A_n) \quad Q = \|y_0\|^2 + 2A_n \left(\|y^*\| + \|y_0 - y^*\| + 2A_n + \sqrt{2B_n} \right) + 2B_n$$

Now denote $r = Ax_n^e - b$, $\Delta = \max\{f_{\mathcal{P}}(x^*) - f_{\mathcal{P}}(x_n^e), 0\}$, and for scalar $z > 0$ define vector $y = z \cdot \frac{r}{\|r\|}$ (assuming that $\|r\| \neq 0$). From eq. (47) we get

$$z \cdot \|r\| = \langle y, r \rangle \leq \Delta + \frac{\|y\|^2 + P \cdot \|y\| + Q}{2\gamma T_n} = \Delta + \frac{z^2 + P \cdot z + Q}{2\gamma T_n}$$

and therefore

$$\|r\| \leq \frac{z}{2\gamma T_n} + \left(\Delta + \frac{Q}{2\gamma T_n} \right) \frac{1}{z} + \frac{P}{2\gamma T_n}$$

Since $\min_{z>0} \left[\frac{z}{\alpha} + \frac{\beta}{z} \right] = 2\sqrt{\frac{\beta}{\alpha}}$ for $\alpha, \beta > 0$, we obtain the claim in eq. (21):

$$\|r\| \leq 2\sqrt{\frac{\Delta + \frac{Q}{2\gamma T_n}}{2\gamma T_n}} + \frac{P}{2\gamma T_n} \leq \sqrt{\frac{2\Delta}{\gamma T_n}} + \frac{\sqrt{Q}}{\gamma T_n} + \frac{P}{2\gamma T_n} = \sqrt{\frac{2\Delta}{\gamma T_n}} + \frac{P/2 + \sqrt{Q}}{\gamma T_n}$$

C.2 Proof of Theorem 7(b)

In this section for a point $x \in \mathbb{R}^d$ and convex set $C \subseteq \mathbb{R}^d$ we denote x_C to be the projection of x to C , and $\text{dist}(x, C) = \|x - x_C\|$. We will prove the following key result.

Theorem 19. *Consider polytope $\mathcal{P} \subseteq \mathbb{R}^d$ and linear subspace $\mathcal{H} \subseteq \mathbb{R}^d$ with $\mathcal{Q} = \mathcal{P} \cap \mathcal{H} \neq \emptyset$. There exists constant β such that*

$$\text{dist}(x, \mathcal{Q}) \leq \beta \cdot \text{dist}(x, \mathcal{H}) \quad \forall x \in \mathcal{P} \quad (48)$$

This would easily imply Theorem 7(b). Indeed, let $\mathcal{H} = \{x \in \mathbb{R}^d : Ax - b = 0\}$, then for any $x \in \mathcal{P}$ we have $\|x - x_{\mathcal{H}}\| \leq \frac{\|Ax - b\|}{\sigma_{\min}(A)}$, where $\sigma_{\min}(A)$ denotes the smallest singular value of A . This easily follows from the projection formula $x_{\mathcal{H}} = x - A^*(AA^*)^{-1}(Ax - b)$ and the fact that any A can be factorized as $A = U\Sigma V^*$ so that $A^*(AA^*)^{-1} = V\Sigma^{-1}U^*$ and $\|V\Sigma^{-1}U^*\| = \sigma_{\min}^{-1}(A)$. It follows from Theorem 19 that $\|\tilde{x} - x\| \leq \frac{\beta}{\sigma_{\min}(A)} \|Ax - b\|$ where $\tilde{x} = x_{\mathcal{P} \cap \mathcal{H}}$. Since $\tilde{x} \in \mathcal{P} \cap \mathcal{H}$, we have $f_{\mathcal{P}}(\tilde{x}) \geq f_{\mathcal{P}}(x^*)$. Convexity of $f_{\mathcal{P}}$ gives

$$f_{\mathcal{P}}(x) \geq f_{\mathcal{P}}(\tilde{x}) + \langle \nabla f(\tilde{x}), x - \tilde{x} \rangle \geq f_{\mathcal{P}}(x^*) - \|\nabla f(\tilde{x})\| \cdot \|\tilde{x} - x\| \geq f_{\mathcal{P}}(x^*) - \|\nabla f(\tilde{x})\| \cdot \frac{\beta}{\sigma_{\min}(A)} \|Ax - b\|$$

which yields Theorem 7(b), since $\|\nabla f(\tilde{x})\|$ is bounded on \mathcal{P} .

In the remainder of this section we prove Theorem 19. We need to show that $\sup_{x \in \mathcal{P} - \mathcal{H}} \beta(x) < +\infty$ where we defined $\beta(x) = \frac{\text{dist}(x, \mathcal{Q})}{\text{dist}(x, \mathcal{H})}$. (Note that $\text{dist}(x, \mathcal{H}) > 0$ for $x \in \mathcal{P} - \mathcal{H}$, and $\text{dist}(x, \mathcal{H}) = \text{dist}(x, \mathcal{Q}) = 0$ for $x \in \mathcal{P} \cap \mathcal{H} = \mathcal{Q}$).

Clearly, set $\mathcal{Q} = \mathcal{P} \cap \mathcal{H}$ is a polytope. Let $\mathbb{F}_{\mathcal{P}}, \mathbb{F}_{\mathcal{Q}}$ be the sets of faces of \mathcal{P}, \mathcal{Q} respectively. For a pair of faces $(\mathcal{A}, \mathcal{B}) \in \mathbb{F}_{\mathcal{P}} \times \mathbb{F}_{\mathcal{Q}}$ define

$$\mathcal{B}_\circ^+ = \{x \in \mathcal{P} \mid \mathcal{B} \text{ is the minimal face of } \mathcal{Q} \text{ containing } x_\mathcal{Q}\} \quad (49)$$

$$\mathcal{B}^+ = \{x \in \mathcal{P} \mid x_\mathcal{Q} = x_\mathcal{B} = x_{\text{affine}(\mathcal{B})} \in \mathcal{B}\} \quad (50)$$

$$\beta(\mathcal{A}, \mathcal{B}) = \sup_{x \in (\mathcal{A} \cap \mathcal{B}^+) - \mathcal{H}} \beta(x) \quad (51)$$

where $\text{affine}(\mathcal{B})$ is the affine hull of \mathcal{B} , and we assume that $\sup \emptyset = -\infty$. Denote

$$\mathbb{F}_{\mathcal{P}\mathcal{Q}} = \{(\mathcal{A}, \mathcal{B}) \in \mathbb{F}_{\mathcal{P}} \times \mathbb{F}_{\mathcal{Q}} \mid (\mathcal{A} \cap \mathcal{B}^+) - \mathcal{H} \neq \emptyset\}$$

We also define $\dim(\mathcal{A}, \mathcal{B}) = \dim \mathcal{A} + \dim \mathcal{B}$.

Lemma 20. *For each $\mathcal{B} \in \mathbb{F}_{\mathcal{Q}}$ set \mathcal{B}^+ are compact, and $\mathcal{B}_\circ^+ \subseteq \mathcal{B}^+$.*

Proof. The first claim follows from continuity of projections and compactness of faces. To show the second claim, consider $x \in \mathcal{B}_\circ^+$. Clearly, $x_\mathcal{B} = x_\mathcal{Q}$. Also, $x_\mathcal{B}$ belongs to the interior of \mathcal{B} (and thus $x_{\text{affine}(\mathcal{B})} = x_\mathcal{B}$), otherwise there would exist a face $\mathcal{B}' \subsetneq \mathcal{B}$ of \mathcal{Q} containing $x_\mathcal{B}$, contradicting minimality of \mathcal{B} . \square

From this lemma we can conclude that $\mathcal{P} - \mathcal{H} = \bigcup_{(\mathcal{A}, \mathcal{B}) \in \mathbb{F}_{\mathcal{P}\mathcal{Q}}} ((\mathcal{A} \cap \mathcal{B}^+) - \mathcal{H})$. (Indeed, for each $x \in \mathcal{P} - \mathcal{H}$ pick arbitrary face $\mathcal{A} \in \mathbb{F}_{\mathcal{P}}$ containing x , and minimal face $\mathcal{B} \in \mathbb{F}_{\mathcal{Q}}$ containing $x_\mathcal{Q}$. Then $x \in \mathcal{A}$, $x \in \mathcal{B}_\circ^+ \subseteq \mathcal{B}^+$ and so $x \in (\mathcal{A} \cap \mathcal{B}^+) - \mathcal{H}$, implying that $(\mathcal{A}, \mathcal{B}) \in \mathbb{F}_{\mathcal{P}\mathcal{Q}}$). It now suffices to show that $\beta(\mathcal{A}, \mathcal{B}) < +\infty$ for all $(\mathcal{A}, \mathcal{B}) \in \mathbb{F}_{\mathcal{P}\mathcal{Q}}$. Pair $(\mathcal{A}, \mathcal{B}) \in \mathbb{F}_{\mathcal{P}\mathcal{Q}}$ will be called *\mathcal{Q} -nonadjacent* if $\mathcal{A} \cap \mathcal{B}^+ \cap \mathcal{B} \cap \mathcal{H}$ is empty, and *\mathcal{Q} -adjacent* otherwise. It is easy to show the boundedness of $\beta(\mathcal{A}, \mathcal{B})$ for pairs of the former type.

Lemma 21. *If $(\mathcal{A}, \mathcal{B})$ is \mathcal{Q} -nonadjacent then $\beta(\mathcal{A}, \mathcal{B}) < +\infty$.*

Proof. We will prove the contrapositive statement. Suppose that $\beta(\mathcal{A}, \mathcal{B}) = +\infty$, then set $\mathcal{A} \cap \mathcal{B}^+$ contains a sequence x^1, x^2, \dots with $\lim_{i \rightarrow \infty} \beta(x^i) = +\infty$. Since values $\text{dist}(x^i, \mathcal{Q})$ are upper-bounded, we must have $\lim_{i \rightarrow \infty} \text{dist}(x^i, \mathcal{H}) = 0$. Let y be a limit point of sequence $\{x^i\}$, then $\text{dist}(y, \mathcal{H}) = 0$ and $y \in \mathcal{A} \cap \mathcal{B}^+$ by compactness of \mathcal{A} and \mathcal{B}^+ .

We have $y \in \mathcal{Q} = \mathcal{P} \cap \mathcal{H}$, and so $y_\mathcal{Q} = y$. Since $y \in \mathcal{B}^+$, we also have $y_\mathcal{Q} \in \mathcal{B}$, and so $y \in \mathcal{B}$. We showed that $y \in \mathcal{A} \cap \mathcal{B}^+ \cap \mathcal{B} \cap \mathcal{H}$, and so $(\mathcal{A}, \mathcal{B})$ is \mathcal{Q} -adjacent. \square

Next, we analyze \mathcal{Q} -adjacent pairs.

Lemma 22. *Consider a \mathcal{Q} -adjacent pair $(\mathcal{A}, \mathcal{B})$ and points $x \in (\mathcal{A} \cap \mathcal{B}^+) - \mathcal{H}$ and $y \in \mathcal{A} \cap \mathcal{B}^+ \cap \mathcal{B} \cap \mathcal{H}$. For a value $\lambda \geq 0$ define $x(\lambda) = \lambda x + (1 - \lambda)y$.*

(a) *For every linear subspace \mathcal{I} containing y there exists constant $\tau > 0$ (that depends on x, y, \mathcal{I}) such that $\text{dist}(x(\lambda), \mathcal{I}) = \tau \lambda$ for all $\lambda \geq 0$.*

(b) *If $\lambda > 0$ and $x(\lambda) \in \mathcal{A} \cap \mathcal{B}^+$ then $x(\lambda) \notin \mathcal{H}$ and $\beta(x(\lambda)) = \beta(x)$.*

Proof. (a) By applying translations and orthogonal rotations we can assume w.l.o.g. that $\mathcal{I} = \{z \in \mathbb{R}^d \mid z_1 = z_2 = \dots = z_m = 0\}$ for some $m \leq d$, and $y = (0, 0, \dots, 0)$. (Such transformations do not affect the claim of the lemma). We have $x(\lambda) = \lambda \cdot x$ and $\text{dist}(x(\lambda), \mathcal{I}) = \lambda \cdot (x_1^2 + x_2^2 + \dots + x_m^2)^{1/2}$ for any $\lambda \geq 0$. This implies the claim.

(b) Let τ_1 and τ_2 be the values from part (a) applied respectively to linear subspaces $\mathcal{I} = \text{affine}(\mathcal{B})$ and $\mathcal{I} = \mathcal{H}$. For any $\lambda > 0$ with $x(\lambda) \in \mathcal{A} \cap \mathcal{B}^+$ (and in particular for $\lambda = 1$) we have $\text{dist}(x(\lambda), \mathcal{Q}) = \text{dist}(x(\lambda), \text{affine}(\mathcal{B})) = \tau_1 \lambda$ and $\text{dist}(x(\lambda), \mathcal{H}) = \tau_2 \lambda$. The claim follows. \square

Lemma 23. *For every \mathcal{Q} -adjacent pair $(\mathcal{A}, \mathcal{B})$ there exists pair $(\mathcal{A}', \mathcal{B}') \in \mathbb{F}_{\mathcal{P}\mathcal{Q}}$ with $\beta(\mathcal{A}', \mathcal{B}') \geq \beta(\mathcal{A}, \mathcal{B})$ and $\dim(\mathcal{A}', \mathcal{B}') < \dim(\mathcal{A}, \mathcal{B})$.*

Proof. Consider $x \in (\mathcal{A} \cap \mathcal{B}^+) - \mathcal{H}$. Using Lemma 22(b), we conclude that there exists $\hat{x} \in (\mathcal{A} \cap \mathcal{B}^+) - \mathcal{H}$ such that $\beta(\hat{x}) = \beta(x)$ and \hat{x} lies on the boundary of $\mathcal{A} \cap \mathcal{B}^+$. (To obtain such \hat{x} , we can take maximum value of $\hat{\lambda}$ with $x(\hat{\lambda}) \in \mathcal{A} \cap \mathcal{B}^+$, and set $\hat{x} = x(\hat{\lambda})$; the maximum exists since set $\mathcal{A} \cap \mathcal{B}^+$ is compact and $y \neq x$.)

It now suffices to show that $\hat{x} \in \mathcal{A}' \cap (\mathcal{B}')^+$ for some $(\mathcal{A}', \mathcal{B}') \in \mathbb{F}_{\mathcal{P}} \times \mathbb{F}_{\mathcal{Q}}$ with $\dim(\mathcal{A}', \mathcal{B}') < \dim(\mathcal{A}, \mathcal{B})$. (This will imply that $(\mathcal{A}', \mathcal{B}') \in \mathbb{F}_{\mathcal{P}\mathcal{Q}}$ since $\hat{x} \notin \mathcal{H}$). Two cases are possible:

- \hat{x} lies on the boundary of \mathcal{A} . Then there exists face $\mathcal{A}' \subsetneq \mathcal{A}$ of \mathcal{P} containing \hat{x} is the minimal face of \mathcal{P} containing \hat{x} . We then have $\hat{x} \in \mathcal{A}' \cap \mathcal{B}^+$, as desired.
- \hat{x} lies on the boundary of \mathcal{B}^+ . This means that $\hat{x}_{\mathcal{Q}} = \hat{x}_{\mathcal{B}} \in \mathcal{B}$ lies on boundary of \mathcal{B} . Let \mathcal{B}' be the minimal face of \mathcal{Q} containing $\hat{x}_{\mathcal{Q}}$; by the previous fact we must have $\mathcal{B}' \subsetneq \mathcal{B}$. By definition, $\hat{x} \in (\mathcal{B}')^+$, and so $\hat{x} \in (\mathcal{B}')^+$ by Lemma 20. We then have $\hat{x} \in \mathcal{A} \cap (\mathcal{B}')^+$, as desired.

□

By applying Lemma 23 repeatedly (with an induction on $\dim(\mathcal{A}, \mathcal{B})$) we conclude that for every \mathcal{Q} -adjacent pair $(\mathcal{A}, \mathcal{B})$ there exists \mathcal{Q} -nonadjacent pair $(\mathcal{A}', \mathcal{B}') \in \mathbb{F}_{\mathcal{P}\mathcal{Q}}$ with $\beta(\mathcal{A}', \mathcal{B}') \geq \beta(\mathcal{A}, \mathcal{B})$. Combining this fact with Lemma 21 concludes the proof.

C.3 Proof of Theorem 7(c)

Denote $r_n = \|Ax_n^e - b\|$ and $\Delta_n = \max\{f_{\mathcal{P}}(x^*) - f_{\mathcal{P}}(x_n^e), 0\}$. Since sequences $\{A_n\}$ and $\{B_n\}$ are bounded and $T_n = \Theta(n^2)$, by Theorem 7(a) there exist positive constants P, Q such that $r_n \leq P\frac{\sqrt{\Delta_n}}{n} + \frac{Q}{n^2}$ for sufficiently large n 's, or equivalently $n^2 r_n \leq Pn\sqrt{\Delta_n} + Q$. Also, we have $\Delta_n \leq \beta r_n$ by Theorem 7(b), and so $n^2 r_n \leq Pn\sqrt{\beta r_n} + Q$. Denote $z_n = n\sqrt{r_n}$, then $z_n^2 - P\sqrt{\beta}z_n - Q \leq 0$. This means that $z_n \leq z^*$ where z^* is the positive solution of quadratic equation $z^2 - P\sqrt{\beta}z - Q = 0$. We showed that $n\sqrt{r_n} = O(1)$, or equivalently $r_n = O(1/n^2)$.

D Properties of the Frank-Wolfe gap: proofs of Lemmas 1 and 8

D.1 Proof of Lemma 1

Let $x^* \in \arg \min_{x \in \mathcal{P}} g(x)$. Since the gradient of g is L -Lipschitz, for any d with $\hat{x} + d \in \mathcal{P}$ we have

$$g(x^*) \leq g(\hat{x} + d) \leq g(\hat{x}) + \langle \nabla g(\hat{x}), d \rangle + \frac{L}{2} \|d\|^2$$

Let us plug $d = \gamma(s - \hat{x})$ where $\gamma \in [0, 1]$ and $s = \text{1mo}_{\mathcal{P}}(\nabla g(\hat{x}))$. We obtain

$$-c = g(x^*) - g(\hat{x}) \leq \gamma \langle \nabla g(\hat{x}), s - \hat{x} \rangle + \frac{\gamma^2}{2} L \|s - \hat{x}\|^2 \leq -b\gamma + \frac{a\gamma^2}{2}$$

where we denoted $c = g_{\mathcal{P}}^{\downarrow}(\hat{x})$, $b = \langle \nabla g(\hat{x}), \hat{x} - s \rangle = \text{gap}^{\text{FW}}(\hat{x}; g_{\mathcal{P}})$ and $a = LD^2 \geq L\|s - \hat{x}\|^2$. Optimizing the RHS in the last inequality over γ gives

$$c \geq \max_{\gamma \in [0, 1]} \left(-\frac{a\gamma^2}{2} + b\gamma \right) = \begin{cases} b - \frac{a}{2} & \text{if } b \geq \frac{a}{2} & (\gamma = 1) \\ \frac{b^2}{2a} & \text{if } b < \frac{a}{2} & (\gamma = \frac{b}{a}) \end{cases} \quad (52)$$

To prove the second inequality in Lemma 1, we need to show two claims:

- (i) If $c > a/2$ then $b \leq c + a/2$. Suppose this is false: $c > a/2$ and $c < b - a/2$. From (52) we get $b < a$, and so $c < b - a/2 < a - a/2 = a/2$ - a contradiction.
- (ii) If $c \leq a/2$ then $b \leq \sqrt{2ac}$. Suppose this is false: $c \leq a/2$ and $c < b^2/2a$. From (52) we get $b > a$ and thus $c \geq b - a/2$, implying $c > a - a/2 = a/2$ - a contradiction.

D.2 Proof of Lemma 8

Let $\tilde{\mathcal{L}}_{\gamma, \bar{y}}(x, y) = \langle Kx, y \rangle + f(x) - h^*(y) - \frac{1}{2\gamma} \|y - \bar{y}\|^2$ be the function that includes all terms of $\mathcal{L}_{\gamma, \bar{y}}$ except for $\delta_{\mathcal{P}}(x)$, so that $\mathcal{L}(x, y) = \tilde{\mathcal{L}}(x, y) + \delta_{\mathcal{P}}(x)$. Denote $g(x) = \max_y \tilde{\mathcal{L}}_{\gamma, \bar{y}}(x, y)$, $\hat{g}(x) = \tilde{\mathcal{L}}_{\gamma, \bar{y}}(x, \hat{y})$ and $\hat{g}_{\mathcal{P}}(x) = g(x) + \delta_{\mathcal{P}}(x) = \mathcal{L}_{\gamma, \bar{y}}(x, \hat{y})$.

By construction, we have $F_{\gamma, \bar{y}}(x) = g(x) + \delta_{\mathcal{P}}(x)$ and $\hat{g}(x) \leq g(x)$ for any x . Furthermore, functions g and \hat{g} are differentiable convex with a Lipschitz continuous gradient (since the same holds for function $f(x)$ and also for function $\max_y [\langle Kx, y \rangle - h^*(y) - \frac{1}{2\gamma} \|y - \bar{y}\|^2]$ by Lemma 9).

Since $\hat{y} = \arg \max_y \tilde{\mathcal{L}}_{\gamma, \bar{y}}(\hat{x}, y)$, we have $\hat{g}(\hat{x}) = \tilde{\mathcal{L}}_{\gamma, \bar{y}}(\hat{x}, \hat{y}) = \max_y \tilde{\mathcal{L}}_{\gamma, \bar{y}}(\hat{x}, y) = g(\hat{x})$. Since $\hat{g}(\hat{x}) = g(\hat{x})$ and $\hat{g}(x) \leq g(x)$ for all $x \in \mathcal{X}$, we obtain that $\nabla \hat{g}(\hat{x}) = \nabla g(\hat{x})$.

Conditions $\hat{g}(\hat{x}) = g(\hat{x})$ and $\nabla \hat{g}(\hat{x}) = \nabla g(\hat{x})$ mean that $\varepsilon = \mathbf{gap}^{\text{FW}}(\hat{x}; F_{\gamma, \bar{y}}) = \mathbf{gap}^{\text{FW}}(\hat{x}; \hat{g}_{\mathcal{P}})$. Let us apply Lemma 1 to function $\hat{g}_{\mathcal{P}}(x)$. We obtain $\hat{g}_{\mathcal{P}}^{\downarrow}(\hat{x}) \leq \varepsilon$, or equivalently $\mathcal{L}_{\gamma, \bar{y}}(\hat{x}, \hat{y}) \leq \varepsilon + \min_x \mathcal{L}_{\gamma, \bar{y}}(x, \hat{y}) = \varepsilon + H_{\gamma, \bar{y}}(\hat{y})$.

E Properties of function $h_{\gamma, \bar{y}}(Kx)$: proofs of Lemma 9 and 10

E.1 Proof of Lemma 9

We first show the effect on Moreau envelope of a function f when adding a linear function. Let f be a convex functions $f(x)$ and define $g(x) = f(x) + \langle x, b \rangle$ for some $b \in \mathcal{X}$. From the definition of the Moreau envelope and by completing the squares it directly follows that

$$\begin{aligned} m_g^\mu(x) &= \inf_u f(u) + \langle u, b \rangle + \frac{1}{2\mu} \|x - u\|^2 \\ &= \inf_u f(u) + \frac{1}{2\mu} \|(x - \mu b) - u\|^2 + \langle x, b \rangle - \frac{\mu}{2} \|b\|^2 \\ &= m_f^\mu(x - \mu b) + \langle x, b \rangle - \frac{\mu}{2} \|b\|^2. \end{aligned} \tag{53}$$

Next, we rewrite $h_{\gamma, \bar{y}}(Kx)$ as

$$h_{\gamma, \bar{y}}(Kx) = - \inf_{y \in \mathcal{Y}} -\langle Kx, y \rangle + h^*(y) + \frac{1}{2\gamma} \|y - \bar{y}\|^2,$$

which is now in the form of (minus) an infimal convolution of the function $h^*(y)$ minus the linear function $\langle Kx, y \rangle$. The application of (53) immediately yields the first representation in terms of the Moreau envelope of h^* , that is

$$h_{\gamma, \bar{y}}(Kx) = \frac{\gamma}{2} \|Kx\|^2 + \langle Kx, \bar{y} \rangle - m_{h^*}^\gamma(\bar{y} + \gamma Kx).$$

The second representation in terms of the Moreau envelope of h is obtained by applying the formula (41).

The gradient formula in terms of the proximal map of h^* can be deduced from the gradient formula (40) applied to $m_{h^*}^\gamma(\bar{y} + \gamma Kx)$, which yields

$$\nabla_x h_{\gamma, \bar{y}}(x) = \gamma K^* Kx + K^* \bar{y} - \gamma K^* \frac{1}{\gamma} (\bar{y} + \gamma Kx - \text{prox}_{\gamma h^*}(\bar{y} + \gamma Kx)) = K^* \text{prox}_{\gamma h^*}(\bar{y} + \gamma Kx),$$

and Moreau’s identity (42) easily gives the second formula in terms of the proximal map of h .

It remains to compute the Lipschitz constant, which by invoking the the nonexpansiveness (39) of the proximal map shows that

$$\begin{aligned} \|K^* \mathbf{prox}_{\gamma h^*}(\bar{y} + \gamma Kx) - K^* \mathbf{prox}_{\gamma h^*}(\bar{y} + \gamma Ky)\| &\leq \|K\| \|\mathbf{prox}_{\gamma h^*}(\bar{y} + \gamma Kx) - \mathbf{prox}_{\gamma h^*}(\bar{y} + \gamma Ky)\| \\ &\leq \|K\| \|(\bar{y} + \gamma Kx) - (\bar{y} + \gamma Ky)\| \\ &\leq \gamma \|K\|^2 = \gamma L_K^2. \end{aligned}$$

E.2 Proof of Lemma 10

The proximal map to the indicator function of the linear constraint $S = \{y \in \mathcal{Y} : Cy = d\}$ amounts to computing for a given point $z \in \mathcal{Y}$ the projection

$$\mathbf{prox}_{\gamma h^*}(z) = \arg \min_{y \in \mathcal{Y}} \frac{1}{2} \|y - z\|^2, \quad \text{s.t. } Cy = d.$$

It is a well-known fact that such projection can be computed by

$$\mathbf{prox}_{\gamma h^*}(z) = z + (CC^*)^{-1}(d - Cz),$$

where we have assumed that the matrix C has full row rank, that is there are no redundant or conflicting constraints.

With this result in mind and by combining (23) with (37), we can easily obtain a closed form representation of the function $h_{\gamma, \bar{y}}(Kx)$, that is

$$\begin{aligned} h_{\gamma, \bar{y}}(Kx) &= \frac{\gamma}{2} \|Kx\|^2 + \langle Kx, \bar{y} \rangle - m_{h^*}^\gamma(\bar{y} + \gamma Kx) \\ &= \frac{\gamma}{2} \|Kx\|^2 + \langle Kx, \bar{y} \rangle - \delta_S(\mathbf{prox}_{\gamma h^*}(\bar{y} + \gamma Kx)) - \frac{1}{2\gamma} \|\bar{y} + \gamma Kx - \mathbf{prox}_{\gamma h^*}(\bar{y} + \gamma Kx)\|^2 \\ &= \frac{\gamma}{2} \|Kx\|^2 + \langle Kx, \bar{y} \rangle - \frac{1}{2\gamma} \|\bar{y} + \gamma Kx - \mathbf{prox}_{\gamma h^*}(\bar{y} + \gamma Kx)\|^2, \end{aligned}$$

where in the last line we have used the fact that $\delta_S(\mathbf{prox}_{\gamma h^*}(z)) = 0$, for any $z \in \mathcal{Y}$. Obviously, $h_{\gamma, \bar{y}}(Kx)$ is a quadratic function and its (linear) gradient with respect to x follows from the gradient formula (25), that is

$$\nabla_x h_{\gamma, \bar{y}}(Kx) = K^* \mathbf{prox}_{\gamma h^*}(\bar{y} + \gamma Kx) = K^* (\bar{y} + \gamma Kx + (CC^*)^{-1}(d - C(\bar{y} + \gamma Kx))).$$

References

- [1] A. Argyriou, M. Signoretto, and J. A. K. Suykens. Hybrid conditional gradient-smoothing algorithms with applications to sparse and low rank regularization. In *Regularization, Optimization, Kernels, and Support Vector Machines*, pages 53–82. Chapman and Hall/CRC, 2014.
- [2] J.-F. Aujol and C. Dossal. Stability of over-relaxations for the forward-backward algorithm, application to FISTA. *SIAM Journal on Optimization*, 25(4):2408–2433, 2015.
- [3] M. A. Bashiri and X. Zhang. Decomposition-invariant conditional gradient for general polytopes with line search. In *Conference on Neural Information Processing Systems (NIPS)*, 2017.
- [4] A. Beck and S. Shtern. Linearly convergent away-step conditional gradient for non-strongly convex functions. *Mathematical Programming*, 164:1–27, 2017.

- [5] Amir Beck. *First-order methods in optimization*. SIAM, 2017.
- [6] Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sci.*, 2(1):183–202, 2009.
- [7] Andrew Blake, Pushmeet Kohli, and Carsten Rother. *Markov random fields for vision and image processing*. MIT press, 2011.
- [8] G. Braun, S. Pokutta, D. Tu, and S. Wright. Blended conditional gradients: the unconditioning of conditional gradients. In *International Conference on Machine Learning (ICML)*, 2019.
- [9] Antonin Chambolle and Thomas Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *J. Math. Imaging Vision*, 40(1):120–145, 2011.
- [10] Laurent Condat, Daichi Kitahara, Andrés Contreras, and Akira Hirabayashi. Proximal splitting algorithms: A tour of recent advances, with new twists. *arXiv preprint arXiv:1912.00137*, 2019.
- [11] Robert M. Freund, Paul Grigas, and Rahul Mazumder. An extended Frank-Wolfe method with “in-face” directions, and its application to low-rank matrix completion. *SIAM J. Optimization*, 27(1):319–346, 2017.
- [12] D. Garber and O. Meshi. Linear-memory and decomposition-invariant linearly convergent conditional gradient algorithm for structured polytopes. In *Conference on Neural Information Processing Systems (NIPS)*, 2016.
- [13] Gauthier Gidel, Tony Jebara, and Simon Lacoste-Julien. Frank-Wolfe algorithms for saddle point problems. In *Artificial Intelligence and Statistics*, pages 362–371. PMLR, 2017.
- [14] Gauthier Gidel, Fabian Pedregosa, and Simon Lacoste-Julien. Frank-Wolfe splitting via augmented lagrangian method. In *AISTATS*, 2018.
- [15] Osman Güler. New proximal point algorithms for convex minimization. *SIAM J. Optim.*, 2(4):649–664, 1992.
- [16] J. Johnson, D. M. Malioutov, and A. S. Willsky. Lagrangian relaxation for map estimation in graphical models. In *45th Annual Allerton Conference on Communication, Control and Computing*, 2007.
- [17] V. Jojic, S. Gould, and D. Koller. Accelerated dual decomposition for MAP inference. In *International Conference on Machine Learning (ICML)*, 2010.
- [18] Patrick Knöbelreiter, Christian Sormann, Alexander Shekhovtsov, Friedrich Fraundorfer, and Thomas Pock. Belief propagation reloaded: Learning bp-layers for labeling problems. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7900–7909, 2020.
- [19] N. Komodakis, N. Paragios, and G. Tziritas. MRF energy minimization and beyond via dual decomposition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(3):531–552, March 2011.
- [20] S. Lacoste-Julien and M. Jaggi. On the global linear convergence of Frank-Wolfe optimization variants. In *Conference on Neural Information Processing Systems (NIPS)*, 2015.
- [21] S. Lacoste-Julien, M. Jaggi, M. Schmidt, and P. Pletscher. Block-coordinate Frank-Wolfe optimization for structural SVMs. In *International Conference on Machine Learning (ICML)*, 2013.

- [22] Guanghui Lan and Yi Zhou. Conditional gradient sliding for convex optimization. *SIAM Journal on Optimization*, 26(2):1379–1409, 2016.
- [23] Ya-Feng Liu, Xin Liu, and Shiqian Ma. On the non-ergodic convergence rate of an inexact augmented lagrangian framework for composite convex programming. *Mathematics of Operations Research*, 44:2, 2019.
- [24] D. V. N. Luong, P. Parpas, D. Rueckert, and B. Rustem. Solving MRF minimization by mirror descent. In *Advances in Visual Computing - 8th International Symposium, ISVC 2012, Rethymnon, Crete, Greece, July 16-18, 2012, Revised Selected Papers, Part I*, pages 587–598, 2012.
- [25] A. F. T. Martins, M. A. T. Figueiredo, P. M. Q. Aguiar, N. A. Smith, and E. P. Xing. An augmented Lagrangian approach to constrained MAP inference. In *International Conference on Machine Learning (ICML)*, 2011.
- [26] Yurii Nesterov. A method for solving the convex programming problem with convergence rate $O(1/k^2)$. *Dokl. Akad. Nauk SSSR*, 269(3):543–547, 1983.
- [27] Yurii Nesterov. *Introductory lectures on convex optimization*, volume 87 of *Applied Optimization*. Kluwer Academic Publishers, Boston, MA, 2004. A basic course.
- [28] A. Osokin, J.-B. Alayrac, I. Lukasevitz, P. K. Dokania, and S. Lacoste-Julien. Minding the gaps for block Frank-Wolfe optimization of structured SVMs. In *International Conference on Machine Learning (ICML)*, 2016.
- [29] Julian Rasch and Antonin Chambolle. Inexact first-order primal–dual algorithms. *Computational Optimization and Applications*, 76:381–430, 2020.
- [30] P. Ravikumar, A. Agarwal, and M. J. Wainwright. Message-passing for graph-structured linear programs: Proximal methods and rounding schemes. *Journal of Machine Learning Research*, 11:1043–1080, 2010.
- [31] Saverio Salzo and Silvia Villa. Inexact and accelerated proximal point algorithms. *J. Convex Anal.*, 19(4):1167–1192, 2012.
- [32] B. Savchynskyy, J. Kappes, S. Schmidt, and C. Schnörr. A study of nesterov’s scheme for lagrangian decomposition and map labeling. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 1817–1823. IEEE, 2011.
- [33] B. Savchynskyy, S. Schmidt, J. H. Kappes, and C. Schnörr. Efficient MRF energy minimization via adaptive diminishing smoothing. In *Uncertainty in Artificial Intelligence (UAI)*, 2012.
- [34] Bogdan Savchynskyy. Discrete graphical models – an optimization perspective. *Foundations and Trends® in Computer Graphics and Vision*, 11(3-4):160–429, 2019.
- [35] D. Scharstein, H. Hirschmüller, Y. Kitajima, G. Krathwohl, N. Nestic, X. Wang, and P. Westling. High-resolution stereo datasets with subpixel-accurate ground truth. In *German Pattern Recognition Conference (GCPR)*, pages 31–42, 2014.
- [36] M. I. Schlesinger and V. V. Giginyak. Solution to structural recognition (MAX,+)-problems by their equivalent transformations. (2):3–18, 2007.
- [37] M. Schmidt, N. Le Roux, and F. Bach. Convergence rates of inexact proximal-gradient methods for convex optimization. *CoRR*, arXiv:1109.2415, 2011.

- [38] S. Schmidt, B. Savchynskyy, J. H. Kappes, and C. Schnörr. Evaluation of a first-order primal-dual algorithm for mrf energy minimization. In *International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition*, pages 89–103. Springer, 2011.
- [39] Alexander Schwing, Tamir Hazan, Marc Pollefeys, and Raquel Urtasun. Globally convergent parallel MAP LP relaxation solver using the Frank-Wolfe algorithm. In *International Conference on Machine Learning (ICML)*, 2014.
- [40] Alexander G. Schwing, Tamir Hazan, Marc Pollefeys, and Raquel Urtasun. Globally convergent dual MAP LP relaxation solvers using Fenchel-Young margins. In *NIPS*, 2012.
- [41] N. Shah, V. Kolmogorov, and C. H. Lampert. A multi-plane block-coordinate Frank-Wolfe algorithm for training structural SVMs with a costly max-oracle. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [42] Antonio Silveti-Falls, Cesare Molinari, and Jalal Fadili. Generalized conditional gradient with augmented lagrangian for composite minimization. *SIAM Journal on Optimization*, 30(4):2687–2725, 2020.
- [43] G. Storvik and G. Dahl. Lagrangian-based methods for finding map. *IEEE Trans. on Image Processing*, 9(3):469–479, march 2000.
- [44] Paul Swoboda and Vladimir Kolmogorov. MAP inference via block-coordinate Frank-Wolfe algorithm. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [45] P. Tseng. On accelerated proximal gradient methods for convex-concave optimization. Technical report, University of Washington, Seattle, 2008.
- [46] Alp Yurtsever, Olivier Fercoq, Francesco Locatello, and Volkan Cevher. A conditional gradient framework for composite convex minimization with applications to semidefinite programming. In *ICML*, 2018.