

BEYOND SHORT STEPS IN FRANK-WOLFE ALGORITHMS

Anonymous authors

Paper under double-blind review

ABSTRACT

We introduce novel techniques to enhance Frank-Wolfe algorithms by leveraging function smoothness beyond traditional short steps. Our study focuses on Frank-Wolfe algorithms with step sizes that incorporate primal-dual guarantees, offering practical stopping criteria. We present a new Frank-Wolfe algorithm utilizing an optimistic framework and provide a primal-dual convergence proof. Additionally, we propose a generalized short-step strategy aimed at optimizing a computable primal-dual gap. Interestingly, this new generalized short-step strategy is also applicable to gradient descent algorithms beyond Frank-Wolfe methods. Empirical results demonstrate that our optimistic algorithm outperforms existing methods, highlighting its practical advantages.

1 INTRODUCTION

We are interested in solving the following optimization problem:

$$\min_{x \in \mathcal{X}} f(x),$$

where \mathcal{X} is a compact convex set and f is a convex and L -smooth function. The Frank-Wolfe (FW) algorithm (Frank & Wolfe, 1956), also known as the conditional gradient algorithm (Levitin & Polyak, 1966), is a key algorithm for this problem class, particularly for problems where projection onto a constraint set is computationally expensive, as it does not require projections. It leverages a linear minimization oracle (LMO) for \mathcal{X} , which upon presentation with a linear function c returns $v \leftarrow \arg \min_{v \in \mathcal{X}} \langle c, v \rangle$, and a gradient oracle for f , which given a point $x \in \text{dom}(f)$ returns $\nabla f(x)$. Given its low cost per iteration it is often highly effective in various machine learning applications. These include optimal transport (Luise et al., 2019), neural network pruning (Ye et al., 2020), adversarial attacks (Chen et al., 2020), non-negative matrix factorization (Nguyen et al., 2022), particle filtering (Lacoste-Julien et al., 2015), and distributed learning (Bellet et al., 2015), among others.

A more general version of the Frank-Wolfe algorithm can handle the optimization of $f(x) + \psi(x)$ for a convex function ψ . In this case however, we require access to a gradient oracle for f and an oracle that can solve $\arg \min_v \{\langle w, v \rangle + \psi(v)\}$ for any $w \in \mathbb{R}^d$, as discussed in (Nesterov, 2018). We recover the classical FW algorithm by choosing $\psi(x)$ as the indicator function of the set \mathcal{X} .

In this work, we will focus on primal-dual analyses for FW algorithms. In contrast to classical analyses where a step-size strategy and corresponding convergence rate need to be heuristically estimated and then proven by induction, in the primal-dual setup they emerge as a natural consequence of the analysis. This approach does not only provide tighter dual gaps and hence stopping criteria but also enhances the convergence properties of the algorithm. In particular, rather than considering primal progress without clear indication how good the solution is, primal-dual progress refers to the reduction in the primal-dual gap, which is a measure of how close the current solution is to optimality; this will be our measure here. A well-known strategy for FW algorithms is the so-called short step, which essentially chooses the step size to maximize primal guaranteed progress from the descent lemma, see (Frank & Wolfe, 1956); and also (Braun et al., 2022) for an in-depth discussion.

We introduce the concept of primal-dual short steps, that maximize primal-dual guaranteed progress based on a model obtained from the primal-dual analysis. We also show that this new step-size strategy is also applicable to gradient descent algorithms. Moreover, recent advancements have

introduced adaptive step-size strategies, such as, e.g., (Pedregosa et al., 2020) (see also (Pokutta, 2024) for a numerically improved version), which further refine the short-step. These strategies aim to decrease a model of the function adaptively, enhancing the algorithm’s performance and stability and in particular allow for leveraging local curvature information. However, they come at the extra cost of a line search like procedure and are subject to the same lower bounds as the classical short steps (Guélat & Marcotte, 1986).

MAIN CONTRIBUTIONS

Optimistic Frank-Wolfe Algorithm We introduce a novel Frank-Wolfe method that leverages the concept of optimism (Rakhlin & Sridharan, 2013; Steinhardt & Liang, 2014). This algorithm is designed to adapt effectively to varying conditions and provides a robust analysis of the convergence rate associated with a primal-dual gap.

Primal-Dual Short Steps We propose a new class of step-size rules for existing Frank-Wolfe algorithms, termed primal-dual short steps. These steps are based on the sequential minimization of a primal-dual gap defined in the algorithm’s analysis. This approach is flexible and extends to gradient descent algorithms, allowing for line search over the primal-dual gap in both algorithm classes.

Numerical Experiments We conduct numerical experiments that showcase the practical advantages of the proposed optimistic Frank-Wolfe algorithm. In our tests, the results consistently show the effectiveness of this new variant in minimizing the primal-dual gap more efficiently than traditional methods. In particular, in most experiments the empirical order of convergence is better, as demonstrated by the steeper slope in a log plot.

Additionally, we have obtained primal-dual analyses for algorithms and results where so far only classical primal convergence analysis where available, often improving constants and providing effective stopping criteria for these algorithms. These findings are detailed in Appendix E.

RELATED WORK

The Frank-Wolfe algorithm, introduced by Frank & Wolfe (1956) where also short steps were introduced, is a foundational method in projection-free optimization, particularly for problems involving convex constraints. Levitin & Polyak (1966) further extended the method, which they referred to as *conditional gradient algorithm*. A significant advancement for the Frank-Wolfe algorithm was made by Jaggi (2013), who introduced the FW gap and provided convergence guarantees on this measure, marking the first instance of FW guarantees for a primal-dual gap.

The development of primal-dual methods for conditional gradient algorithms has been explored in works such as Nesterov (2018) and Diakonikolas & Orecchia (2019), which have contributed to the broader understanding of these algorithms. Additionally, Abernethy & Wang (2017) described a Frank-Wolfe variant that incorporates cumulative gradients, albeit with equal weights, resulting in an additional logarithmic factor in complexity.

Several methods apply the LMO at a convex combination of previously obtained directions. In particular, variants of the Frank-Wolfe algorithm have been developed by Mokhtari et al. (2020), Lu & Freund (2021), and Négier et al. (2020) using stochastic gradients. Notably, the analysis of Lu & Freund (2021) reduces to the optimal rate in the deterministic case, up to constant factors. Wang et al. (2024) also presented a cumulative-gradient FW algorithm, with additional logarithmic factors. Further contributions include the work of Li et al. (2021), who provided a primal-dual analysis of the deterministic algorithm, heavy-ball Frank-Wolfe (HB-FW), using a convex combination of gradients.

In the context of online learning, the concept of optimism, which involves using hints such as past losses to predict future losses, originated in (Azoury & Warmuth, 2001; Chiang et al., 2012; Rakhlin & Sridharan, 2013). This idea had already been developed in saddle point optimization, as seen in the work of Popov (1980). A duality was found and further explored between mirror descent and Frank-Wolfe algorithms by Bach (2015) and Peña (2019). Wang et al. (2024) and Gutman & Peña (2023) have also shown that FW and mirror descent, among other methods, can be unified under a common framework. In this work, we generalize an anytime online-to-batch conversion (Nesterov & Shikhman, 2015; Cutkosky, 2019) for connecting FW algorithms with online learning ones.

Works that have studied modifications of the gradients to obtain faster convergence of FW algorithms and linking them back to gradient descent methods and gradient mappings include [Combettes & Pokutta \(2020\)](#); [Mortagy et al. \(2020\)](#) and in [Diakonikolas et al. \(2020\)](#) local acceleration in the Nesterov sense has been demonstrated for FW algorithms for strongly convex functions while whether local acceleration is possible for non-strongly convex but smooth functions remains an open question.

2 PRELIMINARIES AND GROUNDWORK

In this section, we introduce the necessary notation and an overview of some FW algorithms. In the following let \mathcal{X} be a compact convex set with diameter D . Further, let $f : \mathcal{X} \rightarrow \mathbb{R}$ be a differentiable function defined on an open set containing \mathcal{X} . We assume that f is convex and L -smooth on \mathcal{X} with respect to a norm $\|\cdot\|$. This means that for all $x, y \in \mathcal{X}$, the following inequality holds:

$$0 \leq f(y) - f(x) - \langle \nabla f(x), y - x \rangle \leq \frac{L}{2} \|x - y\|^2.$$

These two conditions imply f has an L -Lipschitz gradient $\|\nabla f(x) - \nabla f(y)\|_* \leq L\|x - y\|$, where $\|\cdot\|_*$ denotes the dual norm.

If not stated otherwise let $x^* \in \arg \min_{x \in \mathcal{X}} f(x)$ be a minimizer of f over \mathcal{X} . Further, let $I_{\mathcal{X}}(x)$ be the indicator function of the set \mathcal{X} that is 0 if $x \in \mathcal{X}$ and $+\infty$ otherwise, and let 1_B be the event indicator function that is 1 if the event B holds and 0 otherwise. We denote $[T] \stackrel{\text{def}}{=} \{1, 2, \dots, T\}$. We denote by $\partial\psi(x)$ the subdifferential of ψ at x so that if $g \in \partial\psi(x)$, we have $\psi(y) \geq \psi(x) + \langle g, y - x \rangle$.

We provide now a sketch of the proof structure for some primal-dual analysis of FW as well as HB-FW and other methods, and in [Appendix B](#) we provide full details of this overview. With respect to ([Li et al., 2021](#)), we provide a slightly more general analysis of HB-FW by allowing for a composite term and we provide an analysis of FW when decreasing regularization is used. We start by defining a lower bound L_t on the optimal value that we have access to at time t , and for iterates $x_t, t \geq 1$, we define a primal-dual gap as

$$G_t \stackrel{\text{def}}{=} f(x_{t+1}) - L_t \quad (1)$$

where $f(x_{t+1})$ is one step ahead of the lower bound, which is a subtle shift that differs slightly from known analysis and allows us to get very simple primal-dual analyses of and convergence results for the new and old algorithms we present in this work. In the sequel, if not stated otherwise, let $a_t > 0$ to be determined later and define $A_t = A_{t-1} + a_t = \sum_{\ell=0}^t a_\ell$. For instance, for FW and HB-FW we use the following lower bounds, or rather, the choice of lower bound defines the iterate v_t of the algorithm:

$$\begin{aligned} A_t f(x^*) &\stackrel{\textcircled{1}}{\geq} \sum_{\ell=0}^t a_\ell f(x_\ell) + \sum_{\ell=0}^t a_\ell \langle \nabla f(x_\ell), x^* - x_\ell \rangle \\ &\stackrel{\textcircled{2}}{\geq} \sum_{\ell=0}^t a_\ell f(x_\ell) + \min_{v \in \mathcal{X}} \left\{ \sum_{\ell=0}^t a_\ell \langle \nabla f(x_\ell), v - x_\ell \rangle \right\} \stackrel{\text{def}}{=} A_t L_t^{\text{HB}} \\ &\stackrel{\textcircled{3}}{\geq} \sum_{\ell=0}^t a_\ell f(x_\ell) + \sum_{\ell=0}^t a_\ell \min_{v \in \mathcal{X}} \langle \nabla f(x_\ell), v - x_\ell \rangle \stackrel{\text{def}}{=} A_t L_t^{\text{FW}} = \sum_{\ell=0}^t a_\ell f(x_\ell) - \sum_{\ell=0}^t a_\ell g(x_\ell). \end{aligned} \quad (2)$$

where the term above $g(x_\ell) \stackrel{\text{def}}{=} \max_{v \in \mathcal{X}} \langle \nabla f(x_\ell), x_\ell - v \rangle$ ($\geq f(x_\ell) - f(x^*)$) is the so-called FW gap, which is also a primal dual gap, the best of which for $\ell \in [t]$ was shown to converge by [Jaggi \(2013\)](#) at a $O(\frac{LD^2}{t})$ rate for FW. Above $\textcircled{1}$ holds by convexity, $\textcircled{2}$ takes a min. This expression is A_t times the lower bound L_t^{HB} in HB-FW is and its arg min corresponds to the iterate v_t of HB-FW, cf. [Algorithm 1](#), whereas after $\textcircled{3}$ we have A_t times the lower bound L_t^{FW} used in the FW algorithm and the arg min of the ℓ -th summand is the iterate v_t of FW, cf. [Algorithm 1](#).

Algorithm 1 Frank-Wolfe and Heavy-Ball FW algorithms**Input:** Function f , feasible set \mathcal{X} , initial point x_0 . Weights a_t, γ_t .

```

1: for  $t = 0$  to  $T$  do
2:    $\diamond$  Choose either 3 (HB-FW) or 4 (FW) for the entire run:
3:    $v_t \leftarrow \arg \min_{v \in \mathcal{X}} \sum_{i=1}^{t-1} a_i \langle \nabla f(x_i), v \rangle$ 
4:    $v_t \leftarrow \arg \min_{v \in \mathcal{X}} \langle \nabla f(x_t), v \rangle$ 
5:    $x_{t+1} \leftarrow (1 - \gamma_t)x_t + \gamma_t v_t$ 
6: end for

```

If we now show a bound

$$A_t G_t - A_{t-1} G_{t-1} \leq E_t \text{ for all } t \geq 0, \quad (3)$$

then we have $f(x_{t+1}) - f(x^*) \leq G_t \leq \frac{1}{A_t} \sum_{i=0}^t E_i$. Our aim is thus to have small E_t and large A_t .

Note $A_{-1} = 0$ by definition. If f is L -smooth with respect to $\|\cdot\|$ and $D \stackrel{\text{def}}{=} \max_{x,y \in \mathcal{X}} \|x - y\|$, then choosing $a_t = 2t + 2$, $\gamma_t = a_t/A_t$ one can show that both FW and HB-FW in Algorithm 1 satisfy (3) with $E_t = \frac{LD^2 a_t^2}{2A_t}$ and consequently $G_t \leq \frac{2LD^2}{t+2}$, cf. Appendix B. The term $f(x_{t+1}) - f(x^*)$ cannot be computed in general since we usually do not have access to the value $f(x^*)$. However, the primal-dual bound G_t is computable and thus it can be used as a stopping criterion.

Remark 2.1 (Alternative step-size strategies). *The analysis uses the guaranteed primal progress, that is, the descent that is guaranteed on $f(x_{t+1}) - f(x_t)$, in particular with the choice $\gamma_t = \frac{a_t}{A_t} = \frac{2}{t+2}$ by using the upper bound that the smoothness inequality provides. Any step size strategy whose primal progress is greater than this guaranteed progress yields a better convergence rate. For instance, the so-called short step-size strategy, maximizes this guaranteed progress along the segment joining x_t and v_t . Indeed the best lower bound on $f(x_t) - f(x_{t+1})$ given by smoothness is*

$$\max_{\gamma \in [0,1]} \left\{ \gamma \langle \nabla f(x_t), x_t - v_t \rangle + \gamma^2 \frac{L \|x_t - v_t\|^2}{2} \right\}.$$

And the arg max is $\gamma = \min\{1, \frac{\langle \nabla f(x_t), x_t - v_t \rangle}{L \|x_t - v_t\|^2}\}$, which is the short step-size. Another alternative is performing a line search over the segment joining x_t and v_t . Even if the line search is performed with an error at iteration t , this error only contributes additively to the corresponding E_t , and so its global impact can be easily quantified.

Although short steps induce monotonic primal progress by definition, they do not necessarily induce monotonic primal-dual progress, which is the important measure to look at when we require a stopping criterion. To remedy this, in Section 4 we introduce a new primal-dual short step strategy that induces monotonic primal-dual progress.

On the other hand all of previous FW approaches exploit smoothness of f in the same way: a quadratic upper bound on the function is computed and the yielded guaranteed progress is used to compensate from some other errors in the analysis. In Section 3, we introduce a framework that goes beyond this by making use of optimism. The algorithm has the potential of better adapt to the environment, as we show in the experiments section, where Optimisitic Frank-Wolfe performs better than other approaches.

3 AN OPTIMISTIC FRANK-WOLFE ALGORITHM

In this section, we propose a new Frank-Wolfe method that at each iteration, uses a prediction of the next gradient in order to minimize a suitable regularized lower model of the objective. The more accurate this prediction is, the better the resulting convergence rate. For functions with Lipschitz continuous gradients, the previous gradient serves as a sufficiently accurate predictor for the next one. We sketch the main ideas of our new algorithm and provide all details in Appendix C.

The algorithm and analysis is based on optimistic versions of the Online Mirror Descent (OMD) algorithm and of the Follow the Regularized Leader (FTRL) algorithm. We provide a slight generalization over these online learning algorithms in order to allow for non-differentiable regularizers, in order to cover typical cases where Frank-Wolfe is applied. To that effect, we make use of the following

Bregman divergence definition, that specifies a subgradient of ϕ of the regularizer for its definition: $D_\psi(x, y, \phi) \stackrel{\text{def}}{=} \psi(x) - \psi(y) - \langle \phi, x - y \rangle$, where $\phi \in \partial\psi(y)$. Note that in the pseudocode of [Algorithm 2](#) we write $D_\psi(v, v_{t-1}, \phi_t \in \partial\psi(v_{t-1}))$ to mean that the algorithm can use $D_\psi(v, v_{t-1}, \phi_t)$ for any $\phi_t \in \partial\psi(v_{t-1})$.

Online learning algorithms typically use regularizers ψ that are strongly convex or enjoy any other curvature property such as uniform convexity, in order to obtain a low enough regret. Even though we do not assume strong convexity or any other curvature property of the regularizer ψ (in fact, ψ can be 0 in the feasible set), we show that we can apply an optimistic approach that leads to the optimal convergence rate in this setting, and we show in [Section 5](#) that our algorithm *empirically outperforms other approaches*.

The starting point of the method, as in [\(2\)](#), is defining a lower bound on the optimal value, which we do by taking inspiration on the anytime online-to-batch conversion of [Cutkosky \(2019\)](#), that connects the regret of online learning algorithms with convergence guarantees of optimization methods. This naturally leads to the definition $x_t = \frac{1}{A_t} \sum_{i=1}^t a_i v_i = \frac{a_t}{A_t} v_t + \frac{A_{t-1}}{A_t} x_{t-1}$ for all $t \geq 1$ in [Algorithm 2](#). We are also able to show that at each iteration, the theory allows for making use of a point y_{t-1} with lower function value than the previous x_{t-1} , in place of using x_{t-1} , and we show that this does not hurt the convergence guarantee, so it allows for heuristics like performing line search over the segment x_{t-1} and v_t , or other future heuristics that may be built on top of this algorithm. The definition of v_t comes from the optimistic online learning algorithmic schemes.

In a similar fashion to the one in [\(1\)](#), using said lower bound we define a primal-dual gap that we denote G_t^{OP} for both variants of the algorithm, cf. [Appendix C.2](#). We provide guarantees on this primal-dual gap by the regret of the optimistic procedure, gradient Lipschitzness of the function, along with the loss weighting and compactness of the domain.

We use an optimistic FTRL algorithm ([Algorithm 2](#)) or an optimistic MD algorithm ([Algorithm 2](#)) with *constant step size*, designed to work for subdifferentiable losses. Interesting, while it is well-known that for constant step size and unconstrained problems, FTRL and OMD have the same updates, we note the more general property that in the constrained setting, FTRL is an instance of OMD with subdifferentiable regularizers for a precise choice of subgradients, cf. [Remark C.5](#).

Algorithm 2 Optimistic Frank-Wolfe algorithms

Input: Convex subdifferentiable regularizer ψ such that $\arg \min_x \{\langle w, x \rangle + \psi(x)\}$ exists for all $w \in \mathbb{R}^d$. A convex function f , differentiable and L -smooth with respect to a norm $\|\cdot\|$ in $\text{dom}(\psi)$. Initial point $x_0 = v_0 \in \text{dom}(\psi)$.

- 1: $A_0 \leftarrow 0; a_0 \leftarrow 0; g_0 = 0; g_1 = \nabla f(x_0)$
- 2: **for** $t \leftarrow 1$ **to** T **do**
- 3: $a_t \leftarrow 2t; A_t \leftarrow A_{t-1} + a_t = \sum_{i=1}^t a_i = t(t+1)$
- 4: \diamond Choose either [5](#) or [6](#) for the entire run, for $\phi_t \in \partial\psi(v_{t-1})$:
- 5: $v_t \leftarrow \text{point in } \arg \min_{v \in \mathbb{R}^d} \{\sum_{i=1}^{t-1} a_i \langle \nabla f(x_i), v \rangle + a_t \langle g_t, v \rangle + \psi(v)\}$
- 6: $v_t \leftarrow \text{point in } \arg \min_{v \in \mathbb{R}^d} \{a_{t-1} \langle \nabla f(x_{t-1}) - g_{t-1}, v \rangle + a_t \langle g_t, v \rangle + D_\psi(v, v_{t-1}; \phi_t)\}$
- 7: $y_{t-1} \leftarrow \text{point such that } f(y_{t-1}) \leq f(x_{t-1})$
- 8: $x_t \leftarrow \frac{A_{t-1}}{A_t} y_{t-1} + \frac{a_t}{A_t} v_t$ \diamond that is, $x_t \leftarrow \frac{t-1}{t+1} y_{t-1} + \frac{2}{t+1} v_t$
- 9: $g_{t+1} \leftarrow \nabla f(x_t)$
- 10: **end for**
- 11: **return** x_T .

One subtlety of G_t^{OP} is that it is not directly computable as is, since the direction we apply the LMO depends on the hint that we choose, that is, a prediction for the next gradient. We can compute a simple close bound of it at every iteration, or if we want to compute it after certain number of iterations, we can do it by performing an extra LMO, cf. [Remark C.6](#). The guarantee we obtain on the algorithm is the following.

Theorem 3.1. $\llbracket \downarrow \rrbracket$ Let \mathcal{X} be compact and convex, and let $\psi : \mathcal{X} \rightarrow \mathbb{R}$ be a closed convex function, subdifferentiable in \mathcal{X} . Let f be convex and L -smooth in the set \mathcal{X} of diameter D with respect to a norm $\|\cdot\|$. The iterates x_t of [Algorithm 2](#) satisfy:

$$f(x_t) - f(x^*) \leq G_t^{\text{OP}} \leq \frac{\psi(x^*) - \psi(x_1)}{t(t+1)} + \frac{4LD^2}{t+1},$$

for the variant in Line 5. For the variant in Line 6 we obtain the same except that $\psi(x^*) - \psi(x_1)$ is substituted by $D_\psi(x^*, x_0; \phi_0)$, where $\phi_0 \in \partial\psi(x_0)$.

We note that the term involving ψ is just a consequence of our choice of step sizes, which we used for simplicity. It is possible to keep the $O(\frac{LD^2}{t+1})$ rate and have an arbitrarily fast polynomial-on- t decay on the term involving ψ by making a different choice of step sizes, see Appendix C.2.

4 PRIMAL-DUAL SHORT STEPS

A key step of the analysis in most Frank-Wolfe algorithms for smooth problems, such as in Section 2, consists of using the smoothness inequality in order to guarantee some descent that compensates other per-iteration errors that appear. Several step-size rules have been devised, that may change depending on the specific setting. However, three families of step sizes stand out in almost every setting, which correspond to the ones we discussed in Section 2: (A) open-loop step-sizes, that only depend on the iteration count, (B) short steps, that minimize the upper bound given by the last computed gradient $\nabla f(x_t)$ and the smoothness inequality with respect to some norm, along the segment in between the current point x_t and the computed Frank-Wolfe vertex $v_t \in \arg \max_{v \in \mathcal{X}} \{\langle \nabla f(x_t), v \rangle\}$, and (C) line search in the aforementioned segment to maximize primal progress.

In the sequel, we devise a new class of step sizes, which are a generalization of (B). The idea of (B) is to greedily maximize the *guaranteed* primal progress along the segment in between x_t and v_t , where we know we are feasible. The key idea of our new step-size rule consists of taking the primal-dual gap (1) that is defined for the analyses with the structure of (3), and choosing the step size in order to maximize the *guaranteed* progress in terms of this primal-dual gap.

We also show that our primal-dual gap bound at iteration t is convex with respect to the step size, which implies that one can efficiently do a line search to maximize the primal-dual progress. In order to show the flexibility of this paradigm, we also generalize these ideas to the gradient descent algorithm.

4.1 PRIMAL-DUAL SHORT STEPS FOR FW ALGORITHMS

Let us consider first the case of FW and HB-FW. In this section, we denote G_t , and x_t and v_t the respective primal-dual gaps, and iterates. From the analyses in Appendix B, cf. (8) or (11), we have the following

$$A_t G_t - A_{t-1} G_{t-1} \leq \frac{L a_t^2}{2 A_t} \|v_t - x_t\|^2,$$

and defining $\gamma_t \stackrel{\text{def}}{=} a_t / A_t$, dividing the above by A_t on both sides, and rearranging gives

$$G_t \leq (1 - \gamma_t) G_{t-1} + \gamma_t^2 \frac{L}{2} \|v_t - x_t\|^2. \quad (4)$$

The right-hand side is minimized for the choice

$$\gamma_t = \min \left\{ 1, \frac{G_{t-1}}{L \|v_t - x_t\|^2} \right\}, \quad (5)$$

which is what we refer to as the *primal-dual short step* for these two algorithms, and focuses on maximizing guaranteed progress of the primal-dual gap, as discussed above. We show that this step-size is sound in the sense that we still keep the optimal convergence guarantees of other approaches. We also show that we can perform a line search over the primal-dual gap bound obtained before using the smoothness inequality in the analyses.

Proposition 4.1. [↓] *Let f be convex and differentiable. The FW and HB-FW algorithms satisfy*

$$G_t \leq (1 - \gamma) G_{t-1} - \gamma \langle \nabla f(x_t), v_t - x_t \rangle - f(x_t) + f((1 - \gamma)x_t + \gamma v_t), \quad (6)$$

for all $\gamma \in [0, 1]$, $t > 1$ and the RHS is convex on γ . If further f is L -smooth w.r.t. a norm $\|\cdot\|$, and $D \stackrel{\text{def}}{=} \max_{x, y \in \mathcal{X}} \|x - y\|$, then using (5) or line search on the RHS of (6), we obtain: $G_t \leq \frac{4LD^2}{t+2}$.

The convergence above for the line search is derived from the one for the primal-dual short step, since the right hand side of (5) upper bounds the one of (6).

We note that, naturally, if we define the gap as $G_t \stackrel{\text{def}}{=} f(x_{t+1}) - f(x^*)$, which is the primal gap, then analyzing the algorithm with the strategy in (3), yields that the primal-dual short steps become regular short steps. In that case, even though we do not know the value $f(x^*)$ in the gap definition, the step can still be defined without this knowledge, cf. Remark D.1.

4.2 PRIMAL-DUAL SHORT STEPS FOR GRADIENT DESCENT

We now extend our primal-dual short steps to gradient descent (GD) whose updates are given by $x_{t+1} \leftarrow x_t - a_t \nabla f(x_t)$. We use the Euclidean norm in this section for the problem $\min_{x \in \mathbb{R}^d} f(x)$, where f is convex and L -smooth. One possible analysis of GD relies on defining a primal-dual gap similarly to the one in Appendix B.2, and then show $A_t G_t \leq A_{t-1} G_{t-1}$ for $a_t = 1/L$ and $A_1 G_1 = \frac{1}{2} \|x_1 - x^*\|_2^2$, for an initial point x_1 and a minimizer x^* . We provide a sketch of the steps that we take and leave the details to Appendix D.1. The lower bound on $f(x^*)$ that we use is

$$A_t f(x^*) \geq \sum_{i=1}^t a_i f(x_i) + \sum_{i=1}^t a_i \langle \nabla f(x_i), x_{t+1} - x_i \rangle + \frac{1}{2} \|x_{t+1} - x_1\|_2^2 - D^2 \stackrel{\text{def}}{=} A_t L_t,$$

where D is an upper bound on the initial distance to a minimizer $\|x^* - x_1\|_2$. Defining the gap as $G_t = f(x_{t+1}) - L_t$, we arrive to

$$G_t \leq \frac{A_{t-1}}{A_t} G_{t-1} + \|\nabla f(x_t)\|_2^2 \left(-\frac{a_t A_{t-1}}{A_t} + \frac{a_t^2 L}{2} - \frac{a_t^2}{2A_t} \right),$$

for $t > 1$ and a similar expression for $t = 1$. Recall that $A_t \stackrel{\text{def}}{=} \sum_{i=1}^t a_i$. So at this stage, one can optimize a_t in order to minimize the right hand side, by solving a simple cubic equation, which is what we term the primal-dual short steps for GD. We show that the optimal value satisfies $a_t \geq \frac{1}{2L}$, which ultimately yields to a convergence rate no slower than $G_t \leq \frac{LD^2}{t}$, as we formalize in the following. We also show, as above, that we can perform a line search for minimizing a better bound on the primal-dual gap, yielding no worse convergence rates. Recall that $f(x_{t+1}) - f(x^*) \leq G_t$.

Proposition 4.2. *[↓] Let f be convex and differentiable. The primal-dual gap of GD is bounded by (26) which is convex on the step a_t . If further f is L -smooth with respect to $\|\cdot\|_2$, GD using the primal-dual short step-size or line search on the aforementioned bound then the step size satisfies $a_t \geq \frac{1}{2L}$ and we have $G_t \leq \frac{LD^2}{t}$.*

A notable difference between the line search on these primal-dual bounds in GD and FW algorithms is that convexity was shown in Proposition 4.1 to hold with respect to the parameter γ which corresponds to $\frac{a_t}{A_t}$ whereas for GD in Proposition 4.2, the convexity is shown to hold for the parameter a_t .

5 EXPERIMENTS

Here we provide experiments demonstrating the good performance of the optimistic variant. All experiments were performed in Julia based on the `FrankWolfe.jl` package run on a MacBook Pro with an Apple M1 chip with Julia 1.11.1. The code will be made publicly available upon publication. In our experiments we compare the optimistic FW variant (with OFTRL) with the heavy ball variant, and the vanilla FW algorithm. For the FW variant, apart from the agnostic step size rule $\gamma_t \stackrel{\text{def}}{=} \frac{2}{t+2}$ we also consider the adaptive line search (indicated by `adaptive` in the plots) from (Pokutta, 2024), a numerically improved variant of (Pedregosa et al., 2020), which is the default line search in the `FrankWolfe.jl` package. As shown in (Guélat & Marcotte, 1986), line search is subject to a lower bound that does not hold for, e.g., open-loop step-sizes and in particular the heavy ball variant and the optimistic variants. In fact, there are cases where the line search variant is slower than the vanilla variant as asymptotically shown in (Bach, 2021) and later in (Wirth et al., 2023; 2024) for the non-asymptotic case; we observe a similar behavior in our experiments as they satisfy the conditions of (Wirth et al., 2023; 2024); see also (Kerdreux et al., 2021).

Our numerical experiments are over various polytopes P , so that in particular $\psi(x) \stackrel{\text{def}}{=} I_P(x)$ is the indicator function of the respective polytope. In the following, we will use the dual gap the definition

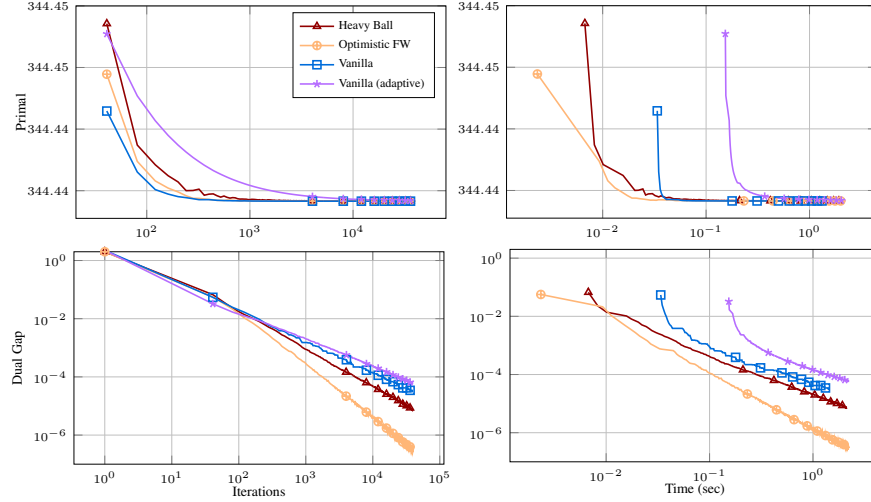


Figure 1: Comparison over the probability simplex with objective $f(x) = \|x - x_0\|_2^2$, where x_0 is a random point outside the probability simplex. The optimistic method converges faster in iterations and time.

of $G_t \stackrel{\text{def}}{=} f(x_t) - L_t^{\text{FW}}$, as in (2). In the case of the heavy ball variant from Appendix B.1 we set $v_i = v_t$, for all $i \leq t$, and $v_t \in \arg \min_{v \in P} \sum_{i=0}^t \frac{a_i}{A_t} f(x_i) + \sum_{i=0}^t \frac{a_i}{A_t} \langle \nabla f(x_i), v - x_i \rangle$, and we refer to this lower bound as L_t^{HB} if not clear from the context. Clearly this is a lower bound and in fact this is precisely the heavy ball vertex that is computed in the heavy ball algorithm. It is important to note that this lower bound and the associated dual gap is valid for *any algorithm* as it is simply a convex combination of linear lower bound functions minimized over P . For comparability and due to its natural form and being valid for any algorithm, we use this definition of the gap also for reporting the results of the optimistic variant in our experiments.

In the case of the vanilla FW variant (either with open loop step sizes or with line search) the v_i in the expression above are chosen to be the FW vertices, i.e., $v_i \in \arg \min_{v \in P} f(x_i) + \langle \nabla f(x_i), v - x_i \rangle$, minimizing each summand separately. This bound is thus separable in contrast to the heavy ball one as we will discuss further below in Remark 5.1 and we can use $f(x^*) \geq \max_{0 \leq i \leq t} f(x_i) + \langle \nabla f(x_i), v_i - x_i \rangle$, as lower bound so that the gap becomes the running minimum of the FW gaps across the iterations, which we refer to as L_t^{FW} .

Remark 5.1 (Strength of lower bounds). Suppose that $g(x_t)$ denotes a generic gap function that bounds the primal gap at x_t , i.e., $f(x_t) - f(x^*) \leq g(x_t)$. We will choose the specific gap function later depending on the context. Observe that the gap function immediately gives a lower bound for $f(x^*)$ simply by rewriting as $f(x^*) \geq f(x_t) - g(x_t)$. Then in line with the above if we make the choice for the lower bound L_t as $A_t f(x^*) \geq \sum_{\ell=0}^t a_\ell f(x_\ell) - \sum_{\ell=0}^t a_\ell g(x_\ell) = A_t L_t$, then in this case, the lower bound L_t cannot be stronger than taking the best lower bound observed so far, since $f(x^*) \geq \max_{\ell=0, \dots, t} f(x_\ell) - g(x_\ell) \geq \frac{1}{A_t} \sum_{\ell=0}^t a_\ell (f(x_\ell) - g(x_\ell))$. This is the case because the lower bound is a convex combination of lower bound terms from individual iterations.

As mentioned, this is different, e.g., for the heavy ball lower bound function, which does not decompose in individual iterations as the Frank-Wolfe vertex is computed for the cumulative function across rounds. The reason why the primal-dual gap for HB-FW can be significantly better (as in: lower) than the Frank-Wolfe gap is illustrated in the following example: The FW algorithm, commonly used with polytopal constraints, sometimes suffers from the so-called zigzag problem (Wolfe, 1970; Guélat & Marcotte, 1986) (see also Braun et al. (2022) for an in-depth discussion), that usually is due to a minimizer being in the relative interior of a face while the points v_t , being the result of an LMO, are chosen as vertices of the polytope. In this scenario, it is possible that $\nabla f(x_t)$ is aligned with the direction $v_t - x_t$ for all t , while a convex combination of gradients is close to being perpendicular to the optimal face making the lower bound L_t^{HB} be closer to the optimal value.

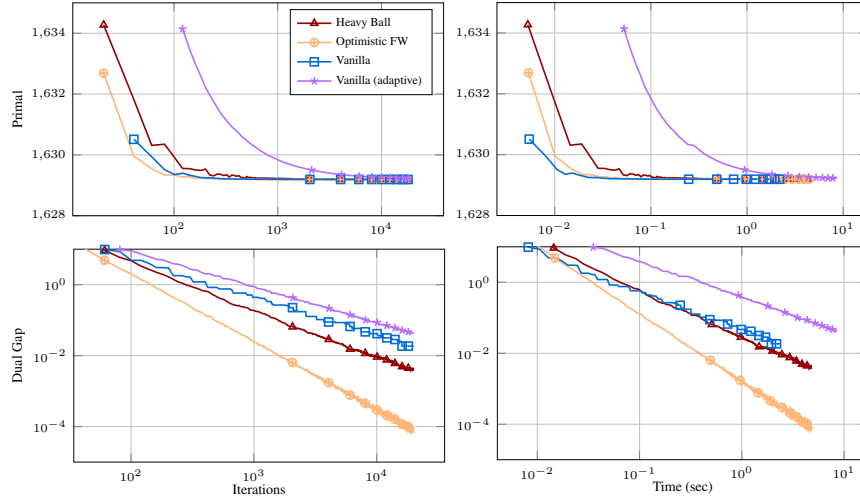


Figure 2: Comparison over k -sparse polytope with $k = 10$ with objective $f(x) = \|Ax - b\|_2^2$, where A and b are random. The optimistic method converges faster in iterations and time.

In the plots we show the primal value and the dual gap vs. the number of iterations and time. Points with excessively large values are not shown in the plots, leading to apparent different starting points. We use log-log plots so that the slope is equal to the polynomial order of convergence. As our variants compute one gradient per iteration, the number of iterations is equal to the number of gradient evaluations. Throughout our experiments we see that the optimistic variant significantly outperforms the heavyball variant, the vanilla FW, and often, although not always, also the FW variant with the adaptive line search strategy from (Pokutta, 2024) in the order of convergence.

We also run experiments for the primal-dual short step from Section 4. However, we found that it behaved similarly than the vanilla FW variant with standard short-steps (or line search), without outperforming this already good heuristic; see Figure 3 in Appendix F. The reason seems to be that because the primal-dual gap measure is better, the step size is actually often *smaller*, leading to no faster convergence; this is reminiscent of the lower bound in (Guélat & Marcotte, 1986).

Finally, we also tested whether optimism is really the main source of improvement or whether most of it is explained by using the heavy ball lower bound which shows to be stronger than the vanilla FW lower bound. In fact, we show in Figure 4 in Appendix F, that indeed optimism is what is making the algorithm be faster. In fact the heavy ball lower bound applied to the vanilla FW variant is weaker than the normal FW gap in our tests. This strongly suggests that it is really the trajectory of the optimistic variant that is better.

We tested various combinations of functions and feasible regions. For the functions, we used (a) quadratic optimization consisting essentially on (b) ill-conditioned quadratic optimization and (c) portfolio optimization problems, which are all widespread benchmarks for Frank-Wolfe algorithms (Braun et al., 2022), see Appendix F. In terms of the feasible regions, we considered (i) the probability simplex $P = \{x \in \mathbb{R}^d \mid \sum_i x_i = 1, x \geq 0\}$, (ii) and the more complex k -sparse polytope P which is the convex hull of all $0/1$ -vectors with k non-zero entries. We detail the instance parameters in the captions of the figures.

In our experiments we report results for regression problems. The reason for this is to allow us to build specific setups where the different convergence behaviors of FW variants are known to emerge. We also performed preliminary experiments over portfolio optimization instances (here the objective is the negative log-likelihood of the portfolio returns) and also Optimal Experiment Design instances (Hendrych et al., 2024) (here the objective arises via matrix means) and found that the optimistic variant also outperforms the other variants similarly.

The optimistic approach presents a good alternative for current Frank-Wolfe methods that is theoretically justified and enjoys faster convergence in practice.

REFERENCES

- Jacob D. Abernethy and Jun-Kun Wang. On frank-wolfe and equilibrium computation. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 6584–6593, 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/7371364b3d72ac9a3ed8638e6f0be2c9-Abstract.html>.
- Katy S. Azoury and Manfred K. Warmuth. Relative loss bounds for on-line density estimation with the exponential family of distributions. *Mach. Learn.*, 43(3):211–246, 2001. doi: 10.1023/A:1010896012157. URL <https://doi.org/10.1023/A:1010896012157>.
- Francis Bach. On the effectiveness of richardson extrapolation in data science. *SIAM Journal on Mathematics of Data Science*, 3(4):1251–1277, 2021.
- Francis R. Bach. Duality between subgradient and conditional gradient methods. *SIAM J. Optim.*, 25(1):115–129, 2015. doi: 10.1137/130941961. URL <https://doi.org/10.1137/130941961>.
- Aurélien Bellet, Yingyu Liang, Alireza Bagheri Garakani, Maria-Florina Balcan, and Fei Sha. A distributed frank-wolfe algorithm for communication-efficient sparse learning. In Suresh Venkatasubramanian and Jieping Ye (eds.), *Proceedings of the 2015 SIAM International Conference on Data Mining, Vancouver, BC, Canada, April 30 - May 2, 2015*, pp. 478–486. SIAM, 2015. doi: 10.1137/1.9781611974010.54. URL <https://doi.org/10.1137/1.9781611974010.54>.
- Gábor Braun, Alejandro Carderera, Cyrille W. Combettes, Hamed Hassani, Amin Karbasi, Aryan Mokthari, and Sebastian Pokutta. Conditional gradient methods. *preprint*, 11 2022.
- Jinghui Chen, Dongruo Zhou, Jinfeng Yi, and Quanquan Gu. A frank-wolfe framework for efficient and effective adversarial attacks. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pp. 3486–3494. AAAI Press, 2020. doi: 10.1609/AAAI.V34i04.5753. URL <https://doi.org/10.1609/aaai.v34i04.5753>.
- Chao-Kai Chiang, Tianbao Yang, Chia-Jung Lee, Mehrdad Mahdavi, Chi-Jen Lu, Rong Jin, and Shenghuo Zhu. Online optimization with gradual variations. In Shie Mannor, Nathan Srebro, and Robert C. Williamson (eds.), *COLT 2012 - The 25th Annual Conference on Learning Theory, June 25-27, 2012, Edinburgh, Scotland*, volume 23 of *JMLR Proceedings*, pp. 6.1–6.20. JMLR.org, 2012. URL <http://proceedings.mlr.press/v23/chiang12/chiang12.pdf>.
- Cyrille W Combettes and Sebastian Pokutta. Boosting Frank-Wolfe by Chasing Gradients. *Proceedings of ICML*, 3 2020.
- Ashok Cutkosky. Anytime online-to-batch, optimism and acceleration. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pp. 1446–1454. PMLR, 2019. URL <http://proceedings.mlr.press/v97/cutkosky19a.html>.
- Jelena Diakonikolas and Lorenzo Orecchia. The approximate duality gap technique: A unified theory of first-order methods. *SIAM J. Optim.*, 29(1):660–689, 2019. doi: 10.1137/18M1172314. URL <https://doi.org/10.1137/18M1172314>.
- Jelena Diakonikolas, Alejandro Carderera, and Sebastian Pokutta. Locally Accelerated Conditional Gradients. *Proceedings of AISTATS*, 2020.
- Marguerite Frank and Philip Wolfe. An algorithm for quadratic programming. *Naval research logistics quarterly*, 3(1-2):95–110, 1956.
- Dan Garber and Elad Hazan. Faster rates for the frank-wolfe method over strongly-convex sets. In *International Conference on Machine Learning*, pp. 541–549. PMLR, 2015.

- Jacques Guélat and Patrice Marcotte. Some comments on Wolfe’s ‘away step’. *Mathematical Programming*, 35(1):110–119, 1986.
- David H. Gutman and Javier F. Peña. Perturbed fenchel duality and first-order methods. *Math. Program.*, 198(1):443–469, 2023. doi: 10.1007/S10107-022-01779-7. URL <https://doi.org/10.1007/s10107-022-01779-7>.
- Deborah Hendrych, Mathieu Besançon, and Sebastian Pokutta. Solving the optimal experiment design problem with mixed-integer convex methods. In Leo Liberti (ed.), *22nd International Symposium on Experimental Algorithms (SEA 2024)*, volume 301 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pp. 16:1–16:22, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. ISBN 978-3-95977-325-6. doi: 10.4230/LIPIcs.SEA.2024.16. URL <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.SEA.2024.16>.
- Martin Jaggi. Revisiting frank-wolfe: Projection-free sparse convex optimization. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013*, volume 28 of *JMLR Workshop and Conference Proceedings*, pp. 427–435. JMLR.org, 2013. URL <http://proceedings.mlr.press/v28/jaggi13.html>.
- Thomas Kerdreux, Alexandre d’Aspremont, and Sebastian Pokutta. Local and Global Uniform Convexity Conditions. *preprint*, 2 2021.
- Simon Lacoste-Julien and Martin Jaggi. On the global linear convergence of frank-wolfe optimization variants. *Advances in neural information processing systems*, 28, 2015.
- Simon Lacoste-Julien, Fredrik Lindsten, and Francis R. Bach. Sequential kernel herding: Frank-wolfe optimization for particle filtering. In Guy Lebanon and S. V. N. Vishwanathan (eds.), *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2015, San Diego, California, USA, May 9-12, 2015*, volume 38 of *JMLR Workshop and Conference Proceedings*. JMLR.org, 2015. URL <http://proceedings.mlr.press/v38/lacoste-julien15.html>.
- Evgeny S Levitin and Boris T Polyak. Constrained minimization methods. *USSR Computational mathematics and mathematical physics*, 6(5):1–50, 1966.
- Bingcong Li, Alireza Sadeghi, and Georgios B. Giannakis. Heavy ball momentum for conditional gradient. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pp. 21244–21255, 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/b166b57d195370cd41f80dd29ed523d9-Abstract.html>.
- Haihao Lu and Robert M. Freund. Generalized stochastic frank-wolfe algorithm with stochastic “substitute” gradient for structured convex optimization. *Math. Program.*, 187(1):317–349, 2021. doi: 10.1007/S10107-020-01480-7. URL <https://doi.org/10.1007/s10107-020-01480-7>.
- Giulia Luise, Saverio Salzo, Massimiliano Pontil, and Carlo Ciliberto. Sinkhorn barycenters with free support via frank-wolfe algorithm. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 9318–9329, 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/9f96f36b7aae3b1ff847c26ac94c604e-Abstract.html>.
- Aryan Mokhtari, Hamed Hassani, and Amin Karbasi. Stochastic conditional gradient methods: From convex minimization to submodular maximization. *J. Mach. Learn. Res.*, 21:105:1–105:49, 2020. URL <https://jmlr.org/papers/v21/18-764.html>.
- Hassan Mortagy, Swati Gupta, and Sebastian Pokutta. Walking in the Shadow: A New Perspective on Descent Directions for Constrained Minimization. *Proceedings of NeurIPS*, 6 2020.
- Geoffrey Négier, Gideon Dresdner, Alicia Y. Tsai, Laurent El Ghaoui, Francesco Locatello, Robert Freund, and Fabian Pedregosa. Stochastic frank-wolfe for constrained finite-sum minimization. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 7253–7262. PMLR, 2020. URL <http://proceedings.mlr.press/v119/negiar20a.html>.

- Yu Nesterov and Vladimir Shikhman. Quasi-monotone subgradient methods for nonsmooth convex minimization. *Journal of Optimization Theory and Applications*, 165(3):917–940, 2015. URL <https://link.springer.com/article/10.1007/s10957-014-0677-5>.
- Yurii E. Nesterov. Complexity bounds for primal-dual methods minimizing the model of objective function. *Math. Program.*, 171(1-2):311–330, 2018. doi: 10.1007/S10107-017-1188-6. URL <https://doi.org/10.1007/s10107-017-1188-6>.
- Tri Nguyen, Xiao Fu, and Ruiyuan Wu. Memory-efficient convex optimization for self-dictionary separable nonnegative matrix factorization: A frank-wolfe approach. *IEEE Trans. Signal Process.*, 70:3221–3236, 2022. doi: 10.1109/TSP.2022.3177845. URL <https://doi.org/10.1109/TSP.2022.3177845>.
- Francesco Orabona. A modern introduction to online learning. *CoRR*, abs/1912.13213, 2019. URL <http://arxiv.org/abs/1912.13213>.
- Fabian Pedregosa, Geoffrey Negiar, Armin Askari, and Martin Jaggi. Linearly convergent frank-wolfe with backtracking line-search. In *International conference on artificial intelligence and statistics*, pp. 1–10. PMLR, 2020.
- Javier F Peña. Affine invariant convergence rates of the conditional gradient method. *SIAM Journal on Optimization*, 33(4):2654–2674, 2023.
- Javier F Peña and Daniel Rodríguez. Polytope conditioning and linear convergence of the Frank–Wolfe algorithm. *Mathematics of Operations Research*, 44(1):1–18, 2018. doi: 10.1287/moor.2017.0910.
- Javier Peña. Generalized conditional subgradient and generalized mirror descent: duality, convergence, and symmetry. *arXiv preprint arXiv:1903.00459*, 2019.
- S. Pokutta. The Frank-Wolfe algorithm: a short introduction. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 126:3–35, 1 2024.
- Leonid Denisovich Popov. A modification of the arrow-hurwitz method of search for saddle points. *Mat. Zametki*, 28(5):777–784, 1980.
- Alexander Rakhlin and Karthik Sridharan. Online learning with predictable sequences. In Shai Shalev-Shwartz and Ingo Steinwart (eds.), *COLT 2013 - The 26th Annual Conference on Learning Theory, June 12-14, 2013, Princeton University, NJ, USA*, volume 30 of *JMLR Workshop and Conference Proceedings*, pp. 993–1019. JMLR.org, 2013. URL <http://proceedings.mlr.press/v30/Rakhlin13.html>.
- Jacob Steinhardt and Percy Liang. Adaptivity and optimism: An improved exponentiated gradient algorithm. In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014*, volume 32 of *JMLR Workshop and Conference Proceedings*, pp. 1593–1601. JMLR.org, 2014. URL <http://proceedings.mlr.press/v32/steinhardt14.html>.
- Jun-Kun Wang, Jacob D. Abernethy, and Kfir Y. Levy. No-regret dynamics in the fenchel game: a unified framework for algorithmic convex optimization. *Math. Program.*, 205(1):203–268, 2024. doi: 10.1007/S10107-023-01976-Y. URL <https://doi.org/10.1007/s10107-023-01976-y>.
- Elias Wirth, Thomas Kerdreux, and Sebastian Pokutta. Acceleration of Frank-Wolfe algorithms with open loop step-sizes. *Proceedings of AISTATS*, 1 2023.
- Elias Wirth, Javier Peña, and Sebastian Pokutta. Accelerated affine-invariant convergence rates of the frank-wolfe algorithm with open-loop step-sizes. *Mathematical Programming A*, 12 2024.
- Philip Wolfe. Convergence theory in nonlinear programming. *Integer and nonlinear programming*, pp. 1–36, 1970.
- Mao Ye, Chengyue Gong, Lizhen Nie, Denny Zhou, Adam R. Klivans, and Qiang Liu. Good subnetworks provably exist: Pruning via greedy forward selection. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 10820–10830. PMLR, 2020. URL <http://proceedings.mlr.press/v119/ye20b.html>.

A USE OF GENERATIVE AI

Generative AI was used for polishing the writing of this paper. No other use of generative AI was made for this work.

B OVERVIEW AND SOME GENERALIZATIONS OF FRANK-WOLFE’S PRIMAL-DUAL ANALYSES

As discussed in the introduction, the analysis of Frank-Wolfe algorithms on compact convex sets can be generalized to the problem $f + \psi$ where we assume access to a gradient oracle for f and an that solves the global optimization of the function $\arg \min_{x \in \mathbb{R}^d} \langle w, x \rangle + \psi(x)$ for every x , as originally proposed in (Nesterov, 2018). Here, we provide a short proof of this fact, with a slightly better constant, which in particular proves the claims in Section 2 when ψ .

Assume $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is convex, L -smooth, and differentiable in an open set containing $\mathcal{X} \stackrel{\text{def}}{=} \text{dom}(\psi)$ and $\psi : \mathbb{R}^n \rightarrow R$ is a convex, proper, lower semi-continuous function. We assume that for every $u \in \mathbb{R}^n$, there exists $\arg \min_x \{\langle u, x \rangle + \psi(x)\}$. Finally, the convergence after t steps will depend on $D_t \stackrel{\text{def}}{=} \max_{\ell=0, \dots, t} \{\|v_\ell - x_\ell\|\}$ for the points x_ℓ, v_ℓ with $\ell = 0, \dots, t$ as defined in Algorithm 1 up to iteration t . Similar to the original FW algorithm, we can instead substitute D_t by an upper bound of such quantity, e.g., $\text{diam}(\text{dom}(\psi))$, assuming its value is finite. Recall that we define positive weights $a_\ell > 0$ for all $\ell \geq 0$, to be determined later, and $A_t \stackrel{\text{def}}{=} \sum_{i=0}^t a_i$. Thus, $A_{-1} = 0$.

Algorithm 3 Generalized Frank-Wolfe and Heavy-Ball FW algorithms

Input: Functions f and ψ , initial point x_0 . Weights a_t, γ_t .

- 1: **for** $t = 0$ **to** T **do**
 - 2: $a_t \leftarrow 2t + 2, A_t \leftarrow \sum_{i=0}^t a_i = (t+1)(t+2)$
 - 3: \diamond Choose either 3 (generalized HB-FW) or 4 (generalized FW) for the entire run:
 - 4: $v_t \leftarrow \arg \min_{v \in \mathbb{R}^d} \left\{ \sum_{i=0}^t a_i \langle \nabla f(x_i), v \rangle + A_t \psi(v) \right\}$
 - 5: $v_t \leftarrow \arg \min_{v \in \mathbb{R}^d} \{ \langle \nabla f(x_t), v \rangle + \psi(v) \}$
 - 6: $x_{t+1} \leftarrow \frac{A_{t-1}}{A_t} x_t + \frac{a_t}{A_t} v_t$
 - 7: **end for**
-

We define the following lower bound on $f(x^*) + \psi(x^*)$, which is a generalization of (2), where x^* here is defined as a minimizer in $\arg \min_{x \in \mathbb{R}^n} \{f(x) + \psi(x)\}$:

$$\begin{aligned}
 A_t(f(x^*) + \psi(x^*)) &\stackrel{\textcircled{1}}{\geq} \sum_{\ell=0}^t a_\ell f(x_\ell) + \sum_{\ell=0}^t a_\ell \langle \nabla f(x_\ell), x^* - x_\ell \rangle + \sum_{\ell=0}^t a_\ell \psi(x^*) \\
 &\stackrel{\textcircled{2}}{\geq} \sum_{\ell=0}^t a_\ell f(x_\ell) + \sum_{\ell=0}^t a_\ell \langle \nabla f(x_\ell), v_\ell - x_\ell \rangle + \sum_{\ell=0}^t a_\ell \psi(v_\ell) \\
 &\stackrel{\text{def}}{=} A_t L_t,
 \end{aligned}$$

where we applied convexity in $\textcircled{1}$, and for $\textcircled{2}$, we applied the definition of v_ℓ in the algorithm, which is why we define this point. Define the primal-dual gap

$$G_t \stackrel{\text{def}}{=} f(x_{t+1}) - L_t \tag{7}$$

and note that upper bound $f(x_{t+1})$ is one step ahead, which helps the analysis. We obtain the following, for $t \geq 0$:

$$\begin{aligned}
A_t G_t - A_{t-1} G_{t-1} &= A_t f(x_{t+1}) - A_{t-1} f(x_t) + A_t \psi(x_{t+1}) - A_{t-1} \psi(x_t) \\
&- \left(a_t f(x_t) + \sum_{\ell=0}^{t-1} a_\ell f(x_\ell) + a_t \langle \nabla f(x_t), v_t - x_t \rangle + a_t \psi(v_t) + \sum_{\ell=0}^{t-1} a_\ell (\langle \nabla f(x_\ell), v_\ell - x_\ell \rangle + \psi(v_\ell)) \right) \\
&+ \left(\sum_{\ell=0}^{t-1} a_\ell f(x_\ell) + \sum_{\ell=0}^{t-1} a_\ell (\langle \nabla f(x_\ell), v_\ell - x_\ell \rangle + \psi(v_\ell)) \right) \\
&\stackrel{\textcircled{1}}{\leq} \langle \nabla f(x_t), A_t(x_{t+1} - x_t) - a_t(v_t - x_t) \rangle + \frac{L A_t}{2} \|x_{t+1} - x_t\|^2 \\
&\stackrel{\textcircled{2}}{=} \frac{L a_t^2}{2 A_t} \|v_t - x_t\|^2 \stackrel{\textcircled{3}}{\leq} \frac{L D_t^2 a_t^2}{2 A_t} \stackrel{\text{def}}{=} E_t.
\end{aligned} \tag{8}$$

In $\textcircled{1}$ we grouped terms to get $A_t(f(x_{t+1}) - f(x_t))$ and applied smoothness to this term. We also used the definition of x_{t+1} which along with the convexity of ψ , implies $A_t \psi(x_{t+1}) \leq A_{t-1} \psi(x_t) + a_t \psi(v_t)$. In $\textcircled{2}$ we used twice the definition of x_{t+1} which implies $A_t(x_{t+1} - x_t) = a_t(v_t - x_t)$. In $\textcircled{3}$, we bounded the distance of points by the quantity D_t .

Now with the choice $a_t = 2t + 2$ and $A_t = \sum_{i=0}^t a_i = (t+1)(t+2)$ that

$$f(x_{t+1}) + \psi(x_{t+1}) - (f(x^*) + \psi(x^*)) \leq G_t \leq \frac{1}{A_t} \sum_{i=0}^t \frac{L D_t^2 a_i^2}{2 A_i} = \frac{1}{(t+1)(t+2)} \sum_{i=0}^t \frac{2 L D_t^2 (i+1)}{i+2} < \frac{2 L D_t^2}{t+2}. \tag{9}$$

Recall that the classical Frank-Wolfe algorithm corresponds to the algorithm and analysis presented above but for ψ being the indicator function $I_{\mathcal{X}}$ of a compact convex set \mathcal{X} . In such a case, we have the following.

Remark B.1 (Alternative step-sizes in Generalized Frank-Wolfe). *We note that step-size strategies like the ones described in Remark 2.1 also work in the general case $f + \psi$, by finding*

$$\arg \min_{x \in \text{conv } \hat{x}_t, v_t} \left\{ f(\hat{x}_t) + \langle \nabla f(\hat{x}_t), x - \hat{x}_t \rangle + \frac{L}{2} \|x - \hat{x}_t\|^2 + \psi(x) \right\} \text{ or } \arg \min_{x \in \text{conv } \hat{x}_t, v_t} \{ f(x) + \psi(x) \},$$

respectively.

B.1 THE HEAVY BALL FRANK-WOLFE ALGORITHM

We will now consider a variant of the Frank-Wolfe algorithm that uses a heavy ball step, similar to the one from Li et al. (2021), however allowing for more flexibility in the choice of the step size parameters and an additional term $\psi(v)$. As before, let $a_i > 0$ to be determined later and define $A_t = \sum_{i=0}^t a_i$. Let $v_t \stackrel{\text{def}}{=} \arg \min_{v \in \mathcal{X}} \left\{ \langle \sum_{i=0}^t a_i \nabla f(x_i), v \rangle + A_t \psi(v) \right\}$ and let $x_{t+1} \stackrel{\text{def}}{=} \frac{A_{t-1}}{A_t} x_t + \frac{a_t}{A_t} v_t$ be defined as a convex combination of x_t and v_t (and since $A_0 = a_0$, we have $x_1 = v_0$). We define the following lower bound on $f(x^*)$, where x^* is defined as a minimizer in $\arg \min_{x \in \mathcal{X}} \{ f(x) + \psi(x) \}$:

$$\begin{aligned}
A_t(f(x^*) + \psi(x^*)) &\stackrel{\textcircled{1}}{\geq} \sum_{i=0}^t a_i f(x_i) + \sum_{i=0}^t a_i \langle \nabla f(x_i), x^* - x_i \rangle + A_t \psi(x^*) \\
&\stackrel{\textcircled{2}}{\geq} \sum_{i=0}^t a_i f(x_i) + \sum_{i=0}^t a_i \langle \nabla f(x_i), v_t - x_i \rangle + A_t \psi(v_t) \stackrel{\text{def}}{=} A_t L_t
\end{aligned}$$

where we applied convexity in $\textcircled{1}$, and the definition of v_t in $\textcircled{2}$. As before, we define the primal-dual gap as

$$G_t \stackrel{\text{def}}{=} f(x_{t+1}) - L_t. \tag{10}$$

We obtain the following, for $t \geq 0$:

$$\begin{aligned}
A_t G_t - A_{t-1} G_{t-1} &= A_t f(x_{t+1}) - A_{t-1} f(x_t) + A_t \psi(x_{t+1}) - A_{t-1} \psi(x_t) \\
&\quad - \left(a_t f(x_t) + \sum_{i=0}^{t-1} a_i f(x_i) + \sum_{i=0}^{t-1} a_i \langle \nabla f(x_i), v_t - x_i \rangle + A_{t-1} \psi(v_t) \right) - a_t \langle \nabla f(x_t), v_t - x_t \rangle - a_t \psi(v_t) \\
&\quad + \left(\sum_{i=0}^{t-1} a_i f(x_i) + \sum_{i=0}^{t-1} a_i \langle \nabla f(x_i), v_{t-1} - x_i \rangle + A_{t-1} \psi(v_{t-1}) \right) \\
&\stackrel{\textcircled{1}}{\leq} \langle \nabla f(x_t), A_t(x_{t+1} - x_t) - a_t(v_t - x_t) \rangle + \frac{L A_t}{2} \|x_{t+1} - x_t\|^2 \\
&\stackrel{\textcircled{2}}{\leq} \frac{L a_t^2}{2 A_t} \|v_t - x_t\|^2 \stackrel{\textcircled{3}}{\leq} \frac{L D_t^2 a_t^2}{2 A_t} \stackrel{\text{def}}{=} E_t.
\end{aligned} \tag{11}$$

In $\textcircled{1}$ we grouped terms to get $A_t(f(x_{t+1}) - f(x_t))$ and applied smoothness to this term. We also dropped the rest of the terms in the parentheses by optimality of v_{t-1} and dropped the three other terms depending on ψ by its convexity and $A_t x_{t+1} = A_{t-1} x_t + a_t v_t$. In $\textcircled{2}$ we used the definition of x_{t+1} , which is the point that minimizes the right hand side of the smoothness inequality that we applied and so substituting x_{t+1} by $\frac{A_{t-1}}{A_t} x_t + \frac{a_t}{A_t} v_t$ leads to something greater. We simplified some terms. In $\textcircled{3}$, we bounded the distance of points by the diameter of \mathcal{X} . Concluding is now the same as for FW.

B.2 GENERALIZED FRANK-WOLFE ALGORITHM REVISITED

In (8) we have obtained a convergence rate for $\min_x f(x) + \psi(x)$. Alternatively, if we are interested in optimizing f , we can use a regularizer ψ and conclude convergence with a very similar analysis. As a consequence, $\psi(x^*)$ appears in the rates. If ψ is an indicator function of a set this is exactly the heavy ball algorithm from above but for a general ψ , it is different.

The lower bound that we define on $f(x^*)$ with $x^* \in \arg \min_{x \in \mathcal{X}} f(x)$, and ψ as before, is as follows:

$$\begin{aligned}
A_t f(x^*) &\stackrel{\textcircled{1}}{\geq} \sum_{i=0}^t a_i f(x_i) + \sum_{i=0}^t a_i \langle \nabla f(x_i), x^* - x_i \rangle \pm \psi(x^*) \\
&\geq \sum_{i=0}^t a_i f(x_i) + \sum_{i=0}^t a_i \langle \nabla f(x_i), v_t - x_i \rangle + \psi(v_t) - \psi(x^*) \stackrel{\text{def}}{=} A_t L_t,
\end{aligned} \tag{12}$$

where $v_t \stackrel{\text{def}}{=} \arg \min_{v \in \mathbb{R}^n} \{ \sum_{i=0}^t a_i \langle \nabla f(x_i), v \rangle + \psi(v) \}$. Similarly to the previous section, we assume access to an oracle that returns one such v_t , which is assumed to exist. We used convexity in $\textcircled{1}$ and we also added and subtracted ψ in order to compute a lower bound that allows us to reduce

the gap without knowing x^* . Finally, defining the gap $G_t \stackrel{\text{def}}{=} f(x_{t+1}) - L_t$, we have, for all $t \geq 0$

$$\begin{aligned}
A_t G_t - A_{t-1} G_{t-1} - 1_{\{t=0\}} \psi(x^*) &= A_t f(x_{t+1}) - A_{t-1} f(x_t) \\
&\quad - \left(a_t f(x_t) + \sum_{i=0}^{t-1} a_i f(x_i) + a_t \langle \nabla f(x_t), v_t - x_t \rangle + \sum_{i=0}^{t-1} a_i \langle \nabla f(x_i), v_t - x_i \rangle + \psi(v_t) \right) \\
&\quad + \left(\sum_{i=0}^{t-1} a_i f(x_i) + \sum_{i=0}^{t-1} a_i \langle \nabla f(x_i), v_{t-1} - x_i \rangle + \psi(v_{t-1}) \right) \\
&\stackrel{\textcircled{1}}{\leq} \langle \nabla f(x_t), A_t(x_{t+1} - x_t) - a_t(v_t - x_t) \rangle + \frac{L A_t}{2} \|x_{t+1} - x_t\|^2 \\
&\stackrel{\textcircled{2}}{=} \frac{L a_t^2}{2 A_t} \|v_t - x_t\|^2 \stackrel{\textcircled{3}}{\leq} \frac{L D^2 a_t^2}{2 A_t} \stackrel{\text{def}}{=} E_t.
\end{aligned} \tag{13}$$

Step $\textcircled{1}$ uses the optimality of v_{t-1} to bound some terms by 0, Steps $\textcircled{2}$ and $\textcircled{3}$ are identical to the previous analysis. Now to conclude we choose $a_t = 2t + 2$ as before, and so we have $A_t = \sum_{i=0}^t a_i = (t+1)(t+2)$. But we also have to take into account that we have the extra term $\psi(x^*)$ when bounding $A_0 G_0$:

$$\begin{aligned}
f(x_{t+1}) - f(x^*) &\leq \frac{1}{A_t} \left(\psi(x^*) + \sum_{i=0}^t E_i \right) = \frac{\psi(x^*)}{A_t} + \frac{1}{A_t} \sum_{i=0}^t \frac{L D^2 a_i^2}{2 A_i} \\
&= \frac{\psi(x^*)}{A_t} + \frac{1}{(t+1)(t+2)} \sum_{i=0}^t \frac{2 L D^2 (i+1)}{(i+2)} < \frac{\psi(x^*)}{(t+1)(t+2)} + \frac{2 L D^2}{t+2}.
\end{aligned}$$

Remark B.2 (Arbitrary fast rate for $\psi(x^*)$). In fact, the part of the rate involving $\psi(x^*)$ was arbitrary and made for simplicity. Indeed, the intuition is that in the lower bound (12) we are adding a regularizer that is in fact $\psi(x^*)/A_t$ (since the whole inequality is multiplied by A_t in particular $f(x^*)$), so the larger we make A_t be, the smaller the regularizer we are adding is, and the faster that term will go to 0. In algebra, if we set $a_i = \Theta(i^k)$, then we have that $\sum_{i=0}^t a_i^2/A_i = \Theta(\sum_{i=0}^t i^{2k-(k+1)}) = \Theta(t^k)$ and $A_t = \Theta(t^{k+1})$ so in any case the part without regularizer is

$$L D^2 \frac{1}{A_t} \sum_{i=0}^t \frac{a_i^2}{A_i} = \Theta\left(\frac{L D^2}{t}\right),$$

but now A_t grows at any polynomial rate that we want which we can use to decrease $\psi(x^*)/A_t$ fast. If we are adding a ψ ourselves on top of an indicator function because we want some properties to hold, most of the time we would like to keep this term being of the same order as the current gap so its contribution to the total rate of convergence is a multiplicative constant.

This reasoning also applies when finishing the analysis of [Theorem 3.1](#).

C PROOFS FOR OPTIMISTIC FW ALGORITHM

We start by defining the lower bound that we use on $f(x^*)$ for a family of algorithms, using which we define a primal-dual gap that allows to show convergence of our optimistic Frank-Wolfe algorithm. This proof is inspired by the anytime online-to-batch reduction of an optimization problem into an online learning one ([Cutkosky, 2019](#)), which is a generalization and independent work with respect to ([Nesterov & Shikhman, 2015](#)). An interesting modification of the technique that we make is that we allow for selecting a point y_{t-1} such that $f(y_{t-1}) \leq f(x_{t-1})$ before computing the coupling x_t . This allows to incorporate heuristics without degrading the convergence rate, such as performing a line search between x_{t-1} and v_t . We used this modified reduction in order to provide a primal-dual gap, which is computable if the regret of the corresponding online learning problem is computable. If we instead have a computable upper bound $\hat{R}_t(x^*)$ on the regret $R_t(x^*)$ we can correspondingly

define and compute a primal-dual gap based on this bound. This will be the case when we instantiate the framework with our optimistic algorithm.

Let $a_t > 0$ for $t \geq 1$ and define $A_t \stackrel{\text{def}}{=} \sum_{i=1}^t a_i$. Given some points $\{v_i\}_{i \geq 1}$, and given x_0 , we define $x_t \stackrel{\text{def}}{=} \frac{A_{t-1}}{A_t} y_{t-1} + \frac{a_t}{A_t} v_t$, for $t \geq 1$, where y_t is any point such that $f(y_t) \leq f(x_t)$. Then, we have for all $t \geq 1$ and $k = 0, 1, \dots, t$:

$$\begin{aligned}
 A_t f(x^*) &\stackrel{\textcircled{1}}{\geq} \sum_{i=1}^t a_i f(x_i) + \sum_{i=1}^t a_i \langle \nabla f(x_i), x^* - x_i \rangle \\
 &\stackrel{\textcircled{2}}{=} \Phi_k \stackrel{\text{def}}{=} A_k f(x_k) + \sum_{i=1}^k a_i \langle \nabla f(x_i), x^* - v_i \rangle + \sum_{i=k+1}^t a_i f(x_i) + \sum_{i=k+1}^t a_i \langle \nabla f(x_i), x^* - x_i \rangle \\
 &\quad + \sum_{i=0}^{k-1} \left(A_i D_f(y_i, x_{i+1}) + A_i (f(x_i) - f(y_i)) \right) \\
 &\stackrel{\textcircled{3}}{=} A_t f(x_t) + \sum_{i=1}^t a_i \langle \nabla f(x_i), x^* - v_i \rangle + \sum_{i=0}^{t-1} \left(A_i D_f(y_i, x_{i+1}) + A_i (f(x_i) - f(y_i)) \right) \\
 &\stackrel{\textcircled{4}}{\geq} A_t f(x_t) - \widehat{R}_t(x^*) \stackrel{\text{def}}{=} A_t L_t
 \end{aligned} \tag{14}$$

where $\textcircled{1}$ uses convexity of f . We have that $\textcircled{2}$ and $\textcircled{3}$ are due to the terms on the sides being Φ_0 and Φ_t and to $\Phi_k = \Phi_{k+1}$, which holds, since canceling terms such equality is equivalent to $\textcircled{5}$ below:

$$\begin{aligned}
 A_k f(x_k) &\stackrel{\textcircled{5}}{=} a_{k+1} \langle \nabla f(x_{k+1}), x_{k+1} - v_{k+1} \rangle + A_k f(x_{k+1}) + A_k D_f(y_k, x_{k+1}) + A_k (f(x_k) - f(y_k)) \\
 &\stackrel{\textcircled{6}}{=} A_k \langle \nabla f(x_{k+1}), y_k - x_{k+1} \rangle + A_k f(x_{k+1}) + A_k D_f(y_k, x_{k+1}) + A_k (f(x_k) - f(y_k)),
 \end{aligned}$$

where $\textcircled{6}$ holds by definition of x_{k+1} . Thus, $\textcircled{5}$ clearly holds since the right hand side of $\textcircled{6}$ equals the left hand side of $\textcircled{5}$, by definition of the Bregman divergence. Note that the second summand of the right hand side of $\textcircled{3}$ equals to minus the regret $R_t(x^*)$ of the online learning game with linear losses $a_i \langle \nabla f(x_t), \cdot \rangle$ and played points v_i with respect to the comparator x^* . We defined $\widehat{R}_t(x^*)$ as any computable upper bound on $R_t(x^*)$, and thus we obtain $\textcircled{4}$ above by this bound, the assumption on y_{t-1} , and $D_f(y_i, x_{i+1}) \geq 0$.

Now, we define our primal-dual gap as

$$G_t \stackrel{\text{def}}{=} f(x_t) - L_t = \frac{\widehat{R}_t(x^*)}{A_t}, \tag{15}$$

which, by construction, is an upper bound on the primal gap $f(x_t) - f(x^*)$. Thus, an online learning algorithm whose regret, or a bound of it, is computable, provides us with a computable primal-dual gap, and convergence is linked to the value of the regret. We can now apply optimistic follow-the-regularized leader (OFTRL) or optimistic Mirror Descent (OMD) online learning algorithms and provide regret bounds, by using a not necessarily strongly convex or differentiable regularizer in order to obtain an optimistic generalized FW algorithm, where we assume that we can solve subproblems of the form $\arg \min_x \{ \langle w, x \rangle + \psi(x) \}$, for any $w \in \mathbb{R}^d$. Typically in FW algorithms, ψ is just the indicator function of the feasible set. We provide an overview of these algorithms, that we have generalized to deal with subdifferentiable regularizers.

C.1 OPTIMISTIC FTRL AND OPTIMISTIC MD WITH CONVEX SUBDIFFERENTIABLE REGULARIZERS

Given a function ψ that is subdifferentiable in a closed convex feasible set \mathcal{X} , two points $x, y \in \mathcal{X}$ and $\phi \in \partial\psi(y)$, define the non-differentiable Bregman divergence $D_\psi(x, y, \phi)$ as

$$D_\psi(x, y, \phi) \stackrel{\text{def}}{=} \psi(x) - \psi(y) - \langle \phi, x - y \rangle,$$

We start by presenting a regret bound for Optimistic FTRL, by using our possibly non-differentiable non-strongly convex regularizers.

Theorem C.1 (Optimistic FTRL). *Let \mathcal{X} be a closed convex set and let $\psi_t, \ell_t : \mathcal{X} \rightarrow \mathbb{R}$ be closed, proper, convex and subdifferentiable functions in \mathcal{X} , for $t \geq 1$. For $t \in [T]$ and given some hints $\tilde{g}_t \in \mathbb{R}^d$, let $z_t \stackrel{\text{def}}{=} \arg \min_{z \in \mathcal{X}} \{\sum_{i=1}^{t-1} \ell_i(z) + \langle \tilde{g}_t, z \rangle + \psi_t(z)\}$, which we assume to exist. Also, define $z_{T+1} = u$ be an arbitrary point $u \in \mathcal{X}$. Then, the regret $R_T(u)$ satisfies:*

$$\sum_{t=1}^T (\ell_t(z_t) - \ell_t(u)) \leq \psi_{T+1}(u) - \min_{z \in \mathcal{X}} \psi_1(z) + \sum_{t=1}^T \left(\langle g_t - \tilde{g}_t, z_t - z_{t+1} \rangle - D_{f_t}(z_{t+1}, z_t; g_t - \tilde{g}_t) + \psi_t(z_{t+1}) - \psi_{t+1}(z_{t+1}) \right),$$

for all subgradients $g_t \in \partial \ell_t(z_t)$, where $f_t(z) \stackrel{\text{def}}{=} \sum_{i=1}^t \ell_i(z) + \psi_t(z)$.

Proof. First, note that since $z_t = \arg \min_{z \in \mathcal{X}} \{f_t(z) - \ell_t(z) + \langle \tilde{g}_t, z \rangle\}$, we have $0 \in \partial(f_t - \ell_t + \langle \tilde{g}_t, \cdot \rangle)(z_t)$ and thus $g_t - \tilde{g}_t \in \partial f_t(z_t)$ for any $g_t \in \partial \ell_t(z_t)$, so the expression $D_{f_t}(z_{t+1}, z_t; g_t - \tilde{g}_t)$ above makes sense. We bound the regret as

$$\begin{aligned} \sum_{t=1}^T (\ell_t(z_t) - \ell_t(u)) &= \sum_{t=1}^T \left[\left(\ell_t(z_t) + \sum_{i=1}^{t-1} \ell_i(z_t) + \psi_t(z_t) \right) - \left(\sum_{i=1}^t \ell_i(z_{t+1}) + \psi_{t+1}(z_{t+1}) \right) \right] \\ &\quad + \psi_{T+1}(u) - \psi_1(z_1) \\ &= \sum_{t=1}^T \left[\left(\sum_{i=1}^t \ell_i(z_t) + \psi_t(z_t) \right) - \left(\sum_{i=1}^t \ell_i(z_{t+1}) + \psi_t(z_{t+1}) \right) + \psi_t(z_{t+1}) - \psi_{t+1}(z_{t+1}) \right] \\ &\quad + \psi_{T+1}(u) - \psi_1(z_1) \\ &\leq \psi_{T+1}(u) - \min_{z \in \mathcal{X}} \psi_1(z) + \sum_{t=1}^T \left(\langle g_t - \tilde{g}_t, z_t - z_{t+1} \rangle - D_{f_t}(z_{t+1}, z_t; g_t - \tilde{g}_t) + \psi_t(z_{t+1}) - \psi_{t+1}(z_{t+1}) \right). \end{aligned} \tag{16}$$

Above, we simply add and subtract terms, use the definition of z_{T+1} , and in the inequality we just bound $-\psi_1(z_1) \leq -\min_{z \in \mathcal{X}} \psi_1(z)$. \square

Corollary C.2. *Under the assumptions of Theorem C.1, for time-invariant $\psi_t = \psi$ we have*

$$\begin{aligned} \sum_{t=1}^T (\ell_t(z_t) - \ell_t(u)) &\leq \psi_{T+1}(u) - \min_{z \in \mathcal{X}} \psi_1(z) + \sum_{t=1}^T \langle g_t - \tilde{g}_t, z_t - z_{t+1} \rangle - D_{f_t}(z_{t+1}, z_t; g_t - \tilde{g}_t) \\ &\leq \psi_{T+1}(u) - \min_{z \in \mathcal{X}} \psi_1(z) + \sum_{t=1}^T \langle g_t - \tilde{g}_t, z_t - z_{t+1} \rangle, \end{aligned} \tag{17}$$

Proof. The result follows by noticing that the terms with ψ_t in the sum cancel and that the Bregman divergences of the convex functions f_t are non-negative. \square

Now we present an alternative algorithm to the above, the optimistic Mirror Descent algorithm. First, we prove a lemma about the generic update rule of Mirror Descent: $x_{t+1} \in \arg \min_{x \in \mathcal{X}} \{\eta_t \langle g, x \rangle + D_\psi(x, x_t, \phi_t)\}$, where an assumption is made for ψ that the arg min is always non empty. This holds for instance, for ψ being the indicator function of a convex compact set, or ψ being a Legendre function. The following corresponds to the classical mirror descent lemma, but for non-differentiable maps.

Lemma C.3 (Mirror Lemma with non-Differentiable Mirror Map). *Given a closed convex set \mathcal{X} , let ψ be proper, closed, convex and subdifferentiable in \mathcal{X} , and let $g \in \mathbb{R}^d$. If $x_{t+1} \in \arg \min_{x \in \mathcal{X}} \{\langle g, x \rangle + D_\psi(x, x_t; \phi_t)\}$ exists, then for all $u \in \mathcal{X}$ and all $\phi_{t+1} \in \partial \psi(x_{t+1})$:*

$$\langle g, x_{t+1} - u \rangle \leq D_\psi(u, x_t; \phi_t) - D_\psi(x_{t+1}, x_t; \phi_t) - D_\psi(u, x_{t+1}; \phi_{t+1}).$$

We note that by optimality of x_{t+1} , we have $0 \in g + \partial \psi(x_{t+1}) - \phi_t$ and so a possible choice of ϕ_{t+1} is $\phi_t - g$.

Proof. The point x_{t+1} is a minimizer of $F(x) \stackrel{\text{def}}{=} \langle g, x \rangle + D_\psi(x, x_t; \phi_t)$. If we substitute the definition of F into the following expression, implied by the first order optimality condition of x_{t+1} , $F(x_{t+1}) + D_F(u, x_{t+1}; \phi_{t+1}) \leq F(u)$, then we obtain the lemma above. \square

The following theorem about mirror descent is a slight generalization over the common one using not necessarily differentiable regularizers and using a hint for the first step. Compare with, for instance, (Orabona, 2019, Theorem 6.20) with constant step size.

Theorem C.4 (Optimistic MD). *Let \mathcal{X} be a closed convex set and let $\psi, \ell_t : \mathcal{X} \rightarrow \mathbb{R}$ be closed, proper, convex and subdifferentiable functions in \mathcal{X} , for $t \in [T]$. For $t \in [T]$, and given some hints $\tilde{g}_t \in \mathbb{R}^d$ and $g_0 = \tilde{g}_0 = 0$, $z_0 \in \mathcal{X}$, let $z_t \stackrel{\text{def}}{=} \arg \min_{z \in \mathcal{X}} \{\langle g_{t-1} + \tilde{g}_t - \tilde{g}_{t-1}, z \rangle + D_\psi(z, z_{t-1}; \phi_{t-1})\}$ for $\phi_{t-1} \in \partial\psi(z_{t-1})$, which we assume it exists. Also define $z_{T+1} = u$ as an arbitrary point $u \in \mathcal{X}$. Then, the regret $R_T(u)$ satisfies:*

$$\begin{aligned} \sum_{t=1}^T (\ell_t(z_t) - \ell_t(u)) &\leq D_\psi(u, z_0; \phi_0) + \sum_{t=1}^T \left(\langle g_t - \tilde{g}_t, z_t - z_{t+1} \rangle - D_\psi(z_{t+1}, z_t; \phi_t) \right) - D_\psi(z_1, z_0; \phi_0) \\ &\leq D_\psi(u, z_0; \phi_0) + \sum_{t=1}^T \langle g_t - \tilde{g}_t, z_t - z_{t+1} \rangle. \end{aligned} \quad (18)$$

for all subgradients $g_t \in \partial\ell_t(z_t)$, $t \geq 1$, where z_{T+1} is defined as above with $\tilde{g}_{T+1} = 0$.

Proof. Fix the choices $g_t \in \partial\ell_t(z_t)$ for $t \in [T]$ and define $\phi_t \stackrel{\text{def}}{=} g_{t-1} - \tilde{g}_t + \tilde{g}_{t+1}$. Applying Lemma C.3 and adding a term with z_t , we obtain

$$\langle g_t + \tilde{g}_{t+1} - \tilde{g}_t, z_t - u \rangle \leq D_\psi(u, z_t; \phi_t) - D_\psi(z_{t+1}, z_t; \phi_t) - D_\psi(u, z_{t+1}; \phi_{t+1}) + \langle g_t + \tilde{g}_{t+1} - \tilde{g}_t, z_t - z_{t+1} \rangle. \quad (19)$$

Adding up the above from $t = 0$ to T , and taking into account that $g_0 = \tilde{g}_0 = 0$, and setting \tilde{g}_{T+1} we obtain an inequality whose left hand side is

$$\begin{aligned} \sum_{t=0}^T \langle g_t + \tilde{g}_{t+1} - \tilde{g}_t, z_t - u \rangle &= \sum_{t=1}^T \langle g_t, z_t - u \rangle + \langle \tilde{g}_{T+1} - \tilde{g}_0, u \rangle + \sum_{t=0}^T \langle \tilde{g}_{t+1} - \tilde{g}_t, z_t \rangle \\ &\stackrel{\textcircled{1}}{\geq} \sum_{t=1}^T (\ell_t(z_t) - \ell_t(u)) + \sum_{t=0}^T \langle \tilde{g}_{t+1}, z_t - z_{t+1} \rangle. \end{aligned}$$

Note that we can set \tilde{g}_{T+1} to any value since it does not play a role in any of the first T predictions. For simplicity, we used $\tilde{g}_{T+1} = 0$. In $\textcircled{1}$ above, for the first summand we used the subgradient property, the second summand vanished by our choice of hints and then we used an equality for the last term, using again that $\tilde{g}_0 = \tilde{g}_{T+1} = 0$.

Now if we combine the above with the left hand side of the result from adding up (19) from $t = 0$ to T , telescoping and dropping $D_\psi(u, z_{T+1}; \phi_{T+1})$ we obtain

$$\begin{aligned} \sum_{t=1}^T (\ell_t(z_t) - \ell_t(u)) + \sum_{t=0}^T \langle \tilde{g}_{t+1}, z_t - z_{t+1} \rangle &\leq D_\psi(u, z_0; \phi_0) - \sum_{t=0}^T D_\psi(z_{t+1}, z_t; \phi_t) + \sum_{t=0}^T \langle g_t + \tilde{g}_{t+1} - \tilde{g}_t, z_t - z_{t+1} \rangle \\ &\stackrel{\textcircled{1}}{\leq} D_\psi(u, z_0; \phi_0) + \sum_{t=1}^T \langle g_t - \tilde{g}_t, z_t - z_{t+1} \rangle, \end{aligned}$$

where in $\textcircled{1}$ we drop the Bregman divergence terms and we start the sum from $t = 1$ since it is 0 for $t = 0$. The inequalities above equal the one in the statement. \square

Remark C.5 (Constant step size FTRL as an instance of OMD). *It is well known that FTRL and Mirror Descent are equivalent for constant step sizes in the unconstrained setting, and so are their optimistic variants. With our non-differentiable extension we can see that actually in the constrained*

case these algorithms are also equivalent, for a specific choice of subgradients in the MD algorithm. Indeed, because of the optimality of z_t , we have $0 \in (g_{t-1} + \tilde{g}_t - \tilde{g}_{t-1}) + \partial\psi(z_t) - \phi_{t-1}$, and so $\phi_t \stackrel{\text{def}}{=} \phi_{t-1} - (g_{t-1} + \tilde{g}_t - \tilde{g}_{t-1}) \in \partial\psi(z_t)$ can be our next subgradient. In fact, if we choose $z_0 \in \arg \min_{z \in \mathcal{X}} \psi(z)$, we can make the choice $\phi_0 = 0$, and using the recurrence in the definition of ϕ_t , we obtain $\phi_t = -\tilde{g}_t - \sum_{i=1}^{t-1} g_i$ by using $g_0 = \tilde{g}_0 = 0$. So the update rule becomes

$$z_t \in \arg \min_{z \in \mathcal{X}} \{ \langle g_{t-1} + \tilde{g}_t - \tilde{g}_{t-1}, z \rangle + \psi(z) - \langle \phi_{t-1}, z \rangle \} = \arg \min_{z \in \mathcal{X}} \left\{ \langle \tilde{g}_t + \sum_{i=1}^{t-1} g_i, z \rangle + \psi(z) \right\},$$

so under this choice the algorithm becomes equivalent to OFTRL with constant regularizer ψ . However, note that other choices of subgradients ϕ_t yield different update rules.

C.2 CONVERGENCE RATE OF OPTIMISTIC FW

Given the theory developed so far in this section, we can now show the convergence of the two variants in our Algorithm 2. We denote G_t^{OP} the primal-dual gap in (15) when we use either of our two optimistic update rules and where the upper bound of the regret $R_t(x^*)$ is $\hat{R}_t(x^*) \stackrel{\text{def}}{=} \max_{v \in \mathcal{X}} R_t(v)$.

Theorem 3.1. \downarrow Let \mathcal{X} be compact and convex, and let $\psi : \mathcal{X} \rightarrow \mathbb{R}$ be a closed convex function, subdifferentiable in \mathcal{X} . Let f be convex and L -smooth in the set \mathcal{X} of diameter D with respect to a norm $\|\cdot\|$. The iterates x_t of Algorithm 2 satisfy:

$$f(x_t) - f(x^*) \leq G_t^{\text{OP}} \leq \frac{\psi(x^*) - \psi(x_1)}{t(t+1)} + \frac{4LD^2}{t+1},$$

for the variant in Line 5. For the variant in Line 6 we obtain the same except that $\psi(x^*) - \psi(x_1)$ is substituted by $D_\psi(x^*, x_0; \phi_0)$, where $\phi_0 \in \partial\psi(x_0)$.

Generally FW algorithms are applied to optimization problems with compact feasible sets, in which case x^* above exists. However, the generalized FW framework does not assume the feasible set is compact or that f has a minimizer. We note that in the proof, the value of the parameters a_t is not used until the last inequality in the theorem statement. One can also set $a_t = \Theta(t^c)$ for a constant $c > 1$, thus obtaining $A_t = \Theta(t^{c+1})$ and a rate $f(x_t) - f(u) = O(\frac{\psi(u) - \psi(x_1)}{t^{c+1}} + \frac{1}{t^{c+1}} \sum_{i=1}^t \frac{LD^2 i^{2c}}{i^{c+1}}) = O(\frac{\psi(u) - \psi(x_1)}{t^{c+1}} + \frac{LD^2}{t})$. Thus, we can reduce the influence of ψ on the convergence rate polynomially.

We also note that the most common use of Frank-Wolfe corresponds to ψ being the indicator function of a convex compact set \mathcal{X} , in which case we obtain $f(x_t) - f(x^*) \leq \frac{4LD^2}{t+1}$.

Proof. (Theorem 3.1) We start proving the statement for the choice of OFTRL in Line 5 of Algorithm 2 that prescribes how the points v_i are defined. Denoting $R_t(u)$ the corresponding regret at u after t steps, taking into account that we use $a_t g_i \stackrel{\text{def}}{=} a_i \nabla f(x_{i-1})$ as hints for $i \geq 1$, and defining $\tilde{v}_{i+1} = v_{i+1}$ for $i < t$ and $\tilde{v}_{t+1} = u$, we obtain

$$\begin{aligned} f(x_t) - f(u) &\leq G_t^{\text{OP}} \stackrel{\textcircled{1}}{\leq} \frac{R_t(u)}{A_t} \\ &\stackrel{\textcircled{2}}{\leq} \frac{\psi(u) - \psi(x_1)}{A_t} + \frac{1}{A_t} \sum_{i=1}^t a_i \langle \nabla f(x_i) - g_i, v_i - \tilde{v}_{i+1} \rangle \\ &\stackrel{\textcircled{3}}{\leq} \frac{\psi(u) - \psi(x_1)}{A_t} + \frac{1}{A_t} \sum_{i=1}^t a_i L \|x_i - x_{i-1}\| D \\ &\stackrel{\textcircled{4}}{\leq} \frac{\psi(u) - \psi(x_1)}{A_t} + \frac{1}{A_t} \sum_{i=1}^t \frac{a_i^2 L D^2}{A_i} \\ &\stackrel{\textcircled{5}}{\leq} \frac{\psi(u) - \psi(x_1)}{t(t+1)} + \frac{1}{t(t+1)} \sum_{i=1}^t \frac{4i L D^2}{i+1} \\ &\stackrel{\textcircled{6}}{<} \frac{\psi(u) - \psi(x_1)}{t(t+1)} + \frac{4LD^2}{t+1}. \end{aligned} \tag{20}$$

where ① uses (15), which holds since our algorithm has the form described at the beginning of Appendix C for the points v_t running the OFTRL algorithm. Inequality ② holds by Corollary C.2, since our points v_i are computed according to the OFTRL algorithm. Then, ③ uses $g_t = \nabla f(x_{t-1})$, L -Lipschitzness of $\nabla f(\cdot)$, the general Cauchy-Schwarz inequality and bounds $\|v_i - \tilde{v}_{i+1}\| \leq D$ for all $i \in [t]$, and ④ uses that by definition of x_i , we have $x_i - x_{i-1} = \frac{a_i}{A_i}(v_i - x_{i-1})$ for $i \geq 1$ and bounds $\|v_i - x_{i-1}\| \leq D$. This yields the first part of the theorem statement. Now, substituting the choices of $a_t = 2t$, and thus $A_t = \sum_{i=0}^t a_i = t(t+1)$ we obtain ⑤, and a simple bound gives ⑥.

The proof for the OMD variant in Line 6 of Algorithm 2 is identical, except that in ② above we use Theorem C.4 and so \tilde{v}_{t+1} changes to be a point in $\arg \min_{v \in \mathbb{R}^d} \{\langle g_t - \tilde{g}_t, v \rangle + D_\psi(z, v_t; \phi_t)\}$, that is, it corresponds to the next point computed with the update rule when we choose $\tilde{g}_{t+1} = 0$. And also, we have $D_\psi(u, x_0)$ instead of $\psi(u) - \psi(x_1)$. The rest of the inequalities in (20) are the same.

Above, we bounded $\|v_i - x_{i-1}\| \cdot \|v_i - \tilde{v}_{i+1}\| \leq D^2$. If we assume f satisfies the cocoercive inequality $D_f(x, y) \geq \frac{1}{L} \|\nabla f(x) - \nabla f(y)\|^2$, which holds if f is convex and smooth in \mathbb{R}^d and not just in $\text{dom}(\psi)$, cf. (? , Theorem 2.1.5), we can obtain an alternative proof where we will end up having a similar bound, but the D^2 comes from $\|v_i - \tilde{v}_{i+1}\|^2$, except for a term in the first iteration. We do it for OFTRL only for simplicity. Indeed, using the right hand side of ③ in (14) as lower bound, we have some extra terms that allow for this task

$$\begin{aligned}
 f(x_t) - f(u) &\leq G_t \stackrel{\textcircled{1}}{\leq} \frac{R_t(u)}{A_t} \\
 &\stackrel{\textcircled{2}}{\leq} \frac{\psi(u) - \psi(x_1)}{A_t} + \frac{1}{A_t} \left(\sum_{i=1}^t a_i \langle \nabla f(x_i) - g_i, v_i - \tilde{v}_{i+1} \rangle - A_{i-1} D_f(x_{i-1}, x_i) \right) \\
 &\stackrel{\textcircled{3}}{\leq} \frac{\psi(u) - \psi(x_1)}{A_t} + \frac{1}{A_t} \left(a_1 \|\nabla f(x_1) - \nabla f(x_0)\| D + \sum_{i=2}^t \frac{a_i^2 L}{A_{i-1}} \|v_i - \tilde{v}_{i+1}\|^2 \right) \\
 &\stackrel{\textcircled{4}}{\leq} \frac{\psi(u) - \psi(x_1)}{A_t} + \frac{1}{A_t} \left(2LD^2 + \sum_{i=2}^t \frac{4iLD^2}{i-1} \right) \\
 &\stackrel{\textcircled{5}}{\leq} \frac{\psi(u) - \psi(x_1)}{t(t+1)} + \frac{LD^2(2 + 4(t-1) + 4\log(t))}{t(t+1)} \\
 &\stackrel{\textcircled{6}}{=} O\left(\frac{\psi(u) - \psi(x_1)}{t(t+1)} + \frac{LD^2}{t}\right).
 \end{aligned} \tag{21}$$

□

Remark C.6 (Computable primal-dual gap). *Note that for ψ being the indicator of a compact set, and for $u = x^* \in \arg \min_{x \in \text{dom}(\psi)} f(x)$, we have for the OFTRL variant the upper bound on the primal-dual gap $\frac{1}{A_t} \sum_{i=1}^t a_i \langle \nabla f(x_i) - g_i, v_i - \tilde{v}_{i+1} \rangle - A_{i-1} D_f(x_{i-1}, x_i)$, which depends on the unknown point x^* , since $\tilde{v}_{t+1} = x^*$. For a computable primal-dual gap, we can apply an analogous bound to ③ in (20) but for the last summand only, that is*

$$G_t \leq \frac{1}{A_t} \left(\sum_{i=1}^{t-1} a_i \langle \nabla f(x_i) - g_i, v_i - v_{i+1} \rangle + \|\nabla f(x_t) - \nabla f(x_{t-1})\|_* D - \sum_{i=1}^t A_{i-1} D_f(x_{i-1}, x_i) \right).$$

Alternatively, using OMD and taking one more linear minimization oracle to compute \tilde{v}_{t+1} , we already have that our bound is a computable primal-dual gap: $\frac{1}{A_t} \sum_{i=1}^t a_i \langle \nabla f(x_i) - g_i, v_i - \tilde{v}_{i+1} \rangle - A_{i-1} D_f(x_{i-1}, x_i)$. We note that it is also possible to obtain an analogous different \tilde{v}_{t+1} for FTRL that does not depend on x^* , at the expense of computing a linear minimization oracle. Just take the equivalence of OFTRL and OMD for specific choices of the subgradients ϕ_t in Remark C.5.

D PROOFS FOR PRIMAL-DUAL SHORT-STEPS

Proof. (Proposition 4.1) In (8) or (11), after their respective inequalities ①, isolating G_t , using $\gamma_t \stackrel{\text{def}}{=} \frac{a_t}{A_t}$, which ranges from $[0, 1)$ as $a_t \in [0, \infty)$, we obtain

$$\begin{aligned} G_t &\leq \frac{A_{t-1}G_{t-1}}{A_t} - \frac{a_t}{A_t} \langle \nabla f(x_t), v_t - x_t \rangle + (f(x_{t+1}) - f(x_t)) \\ &= (1 - \gamma_t)G_{t-1} - \gamma_t \langle \nabla f(x_t), v_t - x_t \rangle + f((1 - \gamma_t)x_t + \gamma_t v_t) - f(x_t). \end{aligned}$$

Differentiating the right hand side twice with respect to γ_t , we obtain

$$\langle \nabla^2 f((1 - \gamma_t)x_t + \gamma_t v_t)(v_t - x_t), v_t - x_t \rangle \geq 0,$$

hence the expression is convex with respect to γ_t , which proves the first statement.

For the second one, we already performed some steps of the proof in the main paper. Using the primal-dual analysis of FW and HB-FW Appendix B and a few computations we arrive to (4). Plugging the choice of our primal-dual short step (5) into (4) and simplifying leads to

$$G_t \leq \left(1 - \frac{\gamma_t}{2}\right) G_{t-1} = \left(1 - \frac{\min\{1, \frac{G_{t-1}}{L\|v_t - x_t\|^2}\}}{2}\right) G_{t-1},$$

or equivalently

$$G_{t-1} - G_t \geq \frac{1}{2} G_{t-1} \min\{1, \frac{G_{t-1}}{L\|v_t - x_t\|^2}\}.$$

Now, one can apply (Braun et al., 2022, Lemma 2.21) (or similar results; see (Garber & Hazan, 2015)), which converts the contraction inequality into a convergence guarantee, so we obtain:

$$G_t \leq \frac{4LD^2}{t+2},$$

for the primal-dual gap convergence rate after bounding $\|v_i - x_i\|^2 \leq D^2$ for all $i \leq t$.

The progress made by the line search in terms of primal-dual gap is greater than the guaranteed progress (4) used by this second approach and so the line search variant also enjoys the same rates of convergence. \square

Remark D.1 (Using $f(x^*)$ for the gap). *The primal-dual short step is a generalization of the standard short steps, since if we choose the best possible lower bound $L_t \stackrel{\text{def}}{=} f(x^*)$, which is a value that we do not know in general, and if we define the gap accordingly $G_t \stackrel{\text{def}}{=} f(x_{t+1}) - f(x^*)$, then we obtain*

$$A_t G_t - A_{t-1} G_{t-1} = A_t f(x_{t+1}) - A_{t-1} f(x_t) - a_t f(x^*),$$

which after using smoothness and reorganizing yields

$$\begin{aligned} G_t &\leq (1 - \gamma_t)G_{t-1} + \gamma_t(f(x_t) - f(x^*)) + \gamma_t \langle \nabla f(x_t), v_t - x_t \rangle + \gamma_t^2 \frac{L}{2} \|v_t - x_t\|^2 \\ &\stackrel{\text{①}}{=} G_{t-1} + \gamma_t \langle \nabla f(x_t), v_t - x_t \rangle + \gamma_t^2 \frac{L}{2} \|v_t - x_t\|^2 \end{aligned} \tag{22}$$

where ① holds by definition of the gap $G_{t-1} = f(x_t) - f(x^*)$. Note that optimizing the right hand side of the last expression results into regular short steps, which can be computed even if we do not know the value of $L_t \stackrel{\text{def}}{=} f(x^*)$. This computation is simply saying the natural fact that if our primal-dual gap becomes the primal gap, then these new short steps that greedily maximize guaranteed progress on G_t , become the regular short steps, that greedily maximize primal progress.

D.1 DETAILS ON THE PRIMAL-DUAL STEP SIZE FOR GRADIENT DESCENT

First recall a few properties of the classical gradient descent (GD) algorithm, with arbitrary step sizes a_t :

$$\begin{aligned} x_{t+1} &\stackrel{\text{def}}{=} x_t - a_t \nabla f(x_t) = \arg \min_{x \in \mathbb{R}^d} \left\{ a_t \langle \nabla f(x_t), x - x_t \rangle + \frac{1}{2} \|x - x_t\|_2^2 \right\} \\ &= \arg \min_{x \in \mathbb{R}^d} \left\{ \sum_{i=1}^t a_i \langle \nabla f(x_i), x - x_i \rangle + \frac{1}{2} \|x - x_1\|_2^2 \right\} = x_0 - \sum_{i=1}^t a_i \nabla f(x_i). \end{aligned}$$

Proof. (Proposition 4.2)

Recall our notation of positive weights a_t and $A_t \stackrel{\text{def}}{=} \sum_{i=1}^t a_i$. The lower bound that we use to define the primal-dual gap is obtained from

$$\begin{aligned} A_t f(x^*) &\stackrel{\textcircled{1}}{\geq} \sum_{i=1}^t a_i f(x_i) + \sum_{i=1}^t a_i \langle \nabla f(x_i), x^* - x_i \rangle + \frac{1}{2} \|x^* - x_1\|_2^2 \\ &\stackrel{\textcircled{2}}{\geq} \sum_{i=1}^t a_i f(x_i) + \min_{x \in \mathbb{R}^d} \left\{ \sum_{i=1}^t a_i \langle \nabla f(x_i), x - x_i \rangle + \frac{1}{2} \|x - x_1\|_2^2 \right\} - \frac{1}{2} D^2 \\ &\stackrel{\textcircled{3}}{=} \sum_{i=1}^t a_i f(x_i) + \sum_{i=1}^t a_i \langle \nabla f(x_i), x_{t+1} - x_i \rangle + \frac{1}{2} \|x_{t+1} - x_1\|_2^2 - \frac{1}{2} D^2 \stackrel{\text{def}}{=} A_t L_t. \end{aligned}$$

Where $\textcircled{1}$ is due to convexity, and in $\textcircled{2}$ and $\textcircled{3}$ we take a minimum and use the definition of x_{t+1} as a minimizer of $\Lambda_t(x) \stackrel{\text{def}}{=} \sum_{i=1}^t a_i \langle \nabla f(x_i), x - x_i \rangle + \frac{1}{2} \|x - x_1\|_2^2$. We also used the bound $D \geq \|x^* - x_1\|_2$. Now define the gap $G_t \stackrel{\text{def}}{=} f(x_{t+1}) - L_t$. We have

$$\begin{aligned} A_t G_t - A_{t-1} G_{t-1} - 1_{\{t=1\}} \frac{1}{2} D^2 &\stackrel{\textcircled{1}}{=} A_t (f(x_{t+1}) - f(x_t)) + \cancel{a_t f(x_t)} \\ &\quad - \left(\sum_{i=1}^t \cancel{a_i f(x_i)} + \sum_{i=1}^{t-1} a_i \langle \nabla f(x_i), x_{t+1} - x_i \rangle + \frac{1}{2} \|x_{t+1} - x_1\|_2^2 \right) - a_t \langle \nabla f(x_t), x_{t+1} - x_t \rangle \\ &\quad + \left(\sum_{i=1}^{t-1} \cancel{a_i f(x_i)} + \sum_{i=1}^{t-1} a_i \langle \nabla f(x_i), x_t - x_i \rangle + \frac{1}{2} \|x_t - x_1\|_2^2 \right) \\ &\stackrel{\textcircled{2}}{\leq} \langle \nabla f(x_t), A_t (x_{t+1} - x_t) - a_t (x_{t+1} - x_t) \rangle + \frac{A_t L}{2} \|x_{t+1} - x_t\|_2^2 - \frac{1}{2} \|x_{t+1} - x_t\|_2^2 \\ &\stackrel{\textcircled{3}}{=} \|\nabla f(x_t)\|_2^2 \left(-a_t A_{t-1} + \frac{a_t^2 A_t L}{2} - \frac{a_t^2}{2} \right) \stackrel{\text{def}}{=} E_t(a_t). \end{aligned} \tag{23}$$

Above, in $\textcircled{1}$ we write out the definitions and cancel some terms. In $\textcircled{2}$, we used smoothness for the first term, and we used the 1-strong convexity of $\Lambda_{t-1}(x)$ and the fact that x_t is its minimizer so $\Lambda_{t-1}(x_t) - \Lambda_{t-1}(x_{t+1}) \leq -\frac{1}{2} \|x_{t+1} - x_t\|_2^2$. In $\textcircled{3}$, we used the definition of x_{t+1} and grouped terms. This time, we have defined the error $E_t(a_t)$ as a function of a_t .

We can apply the same technique as in the primal dual steps for Frank-Wolfe algorithms. Let $\mathcal{G}_{t-1} \stackrel{\text{def}}{=} A_{t-1} G_{t-1} + 1_{\{t=1\}} \frac{1}{2} D^2 \geq A_{t-1} (f(x_{t+1}) - f(x^*)) \geq 0$. Hence, for $t \geq 1$, we aim to minimize the right hand side of the following that comes reorganizing the above

$$G_t \leq \frac{\mathcal{G}_{t-1}}{A_t} + \|\nabla f(x_t)\|_2^2 \left(-\frac{a_t A_{t-1}}{A_t} + \frac{a_t^2 L}{2} - \frac{a_t^2}{2A_t} \right). \tag{24}$$

Differentiating the RHS with respect to a_t and equating to 0 (recall $A_t = A_{t-1} + a_t$) we obtain

$$\begin{aligned} 0 &= -\frac{\mathcal{G}_{t-1}}{A_t^2} + \|\nabla f(x_t)\|_2^2 \left(-\frac{A_{t-1} A_t - a_t A_{t-1}}{A_t^2} + a_t L - \frac{2a_t A_t - a_t^2}{2A_t^2} \right) \\ \iff 0 &= -2\mathcal{G}_{t-1} + \|\nabla f(x_t)\|_2^2 (-2A_{t-1}^2 + 2a_t A_t^2 L - 2a_t A_t + a_t^2), \end{aligned} \tag{25}$$

which gives a cubic equation for a_t . From now on, assume a_t has the value of this minimizer. Note that this solution is always greater than $\frac{1}{2L}$, for all $t \geq 1$. Indeed, first notice that the right hand side of (24) is convex on a_t for $a_t > 0$, since if we differentiate the right hand side a second time we obtain:

$$2 \frac{\mathcal{G}_{t-1}}{A_t^3} + \|\nabla f(x_t)\|_2^2 \left(\frac{2A_{t-1}^2}{A_t^3} + L - \frac{A_{t-1}^2}{A_t^3} \right) \geq 0.$$

Given this convexity, in order to show that the solution of (25) is $a_t > \frac{1}{2L}$, it is enough to check that the derivative is negative at $a_t = \frac{1}{2L}$ which is immediate after substitution. This fact yields the convergence rate:

$$\begin{aligned} f(x_{t+1}) - f(x_t) &\leq G_t \stackrel{\textcircled{1}}{\leq} \frac{\mathcal{G}_{t-1}}{A_t} + \|\nabla f(x_t)\|_2^2 \left(-\frac{a_t A_{t-1}}{A_t} + \frac{a_t^2 L}{2} - \frac{a_t^2}{2A_t} \right) \\ &\stackrel{\textcircled{2}}{\leq} \frac{\mathcal{G}_{t-1}}{A_{t-1} + 1/(2L)} \stackrel{\textcircled{3}}{\leq} \frac{\mathcal{G}_0}{t/(2L)} \\ &\stackrel{\textcircled{4}}{=} \frac{LD^2}{t}, \end{aligned}$$

where $\textcircled{1}$ is (24), $\textcircled{2}$ holds by substituting the minimizer a_t by $1/(2L)$ and dropping $E_t(1/(2L))$ which after substitution it is immediate to see that it is nonpositive. Inequality $\textcircled{3}$ holds by applying recursively $\textcircled{1}$ and $\textcircled{2}$, taking into account that $\mathcal{G}_k = A_k G_k$ for all $k > 1$. And finally $\textcircled{4}$ substitutes the value of \mathcal{G}_0 .

Line search for minimizing the primal-dual gap Recall that we define $\mathcal{G}_{t-1} \stackrel{\text{def}}{=} A_{t-1} G_{t-1} + 1_{\{t=1\}} \frac{1}{2} D^2 \geq 0$. If for $t \geq 1$, after $\textcircled{1}$ in (23) we do not apply smoothness but only apply the inequality $\Lambda_{t-1}(x_t) - \Lambda_{t-1}(x_{t+1}) \leq -\frac{1}{2} \|x_{t+1} - x_t\|^2$, use $x_{t+1} - x_t = -a_t \nabla f(x_t)$ and then isolate G_t , we obtain

$$G_t \leq \frac{\mathcal{G}_{t-1}}{A_t} + \frac{a_t^2}{2A_t} \|\nabla f(x_t)\|^2 + (f(x_{t+1}) - f(x_t)) \quad (26)$$

If we differentiate twice with respect to a_t , we obtain (taking into account A_t and x_{t+1} depend on a_t):

$$2 \frac{\mathcal{G}_{t-1}}{A_t^3} + \frac{A_{t-1}^2}{A_t^3} \|\nabla f(x_t)\|^2 + \langle \nabla^2 f(x_{t+1}) \nabla f(x_t), \nabla f(x_t) \rangle \geq 0.$$

Thus, the right hand side of (26) is convex with respect to the variable a_t , which means that we can do a line search in order to greedily minimize the primal-dual gap progress at each iteration. Above, we used twice differentiability of f for the last summand but in fact this is not required, since $f(x_{t+1}) = f(x_t - a_t \nabla f(x_t))$ is clearly convex on a_t due to the convexity of f restricted to the line $a_t \mapsto x_t - a_t \nabla f(x_t)$.

Note that the upper bound in (26) is tighter than the one in (24), and thus, the convergence rate from the line search on the primal-dual gap bound is at least the one from the primal-dual short step that minimizes (24). \square

E EXTENDING PRIMAL ANALYSES OF KNOWN FW CONVERGENCE RESULTS TO PRIMAL-DUAL ANALYSES

In this section we demonstrate via some examples how primal-dual analyses can be derived from existing primal analyses. We picked examples that demonstrate relevant aspects of the argument, so that the interested reader should be able to carry them over to other settings. We also refer the interested reader to Braun et al. (2022) for full details on the primal convergence analyses.

In line with Section 2, we let $g(x_\ell) = \max_{v \in \mathcal{X}} \langle \nabla f(x_\ell), x_\ell - v \rangle = \langle \nabla f(x_\ell), x_\ell - v_\ell \rangle$ be the FW gap and

$$A_\ell L_\ell = \sum_{i=0}^{\ell} a_i f(x_i) + \sum_{i=0}^{\ell} a_i \langle \nabla f(x_i), v_i - x_i \rangle,$$

be the associated lower bound function for the standard FW gap. Together with $G_\ell = f(x_{\ell+1}) - L_\ell$ we can now rewrite the crucial term $A_\ell G_\ell - A_{\ell-1} G_{\ell-1} \leq E_\ell$ appearing in primal-dual analyses as follows:

$$\begin{aligned} &A_\ell(f(x_{\ell+1}) - L_\ell) - A_{\ell-1}(f(x_\ell) - L_{\ell-1}) \\ &= A_\ell f(x_{\ell+1}) - A_{\ell-1} f(x_\ell) - (A_\ell L_\ell - A_{\ell-1} L_{\ell-1}) \\ &= A_\ell(f(x_{\ell+1}) - f(x_\ell)) + a_\ell g(x_\ell) \leq E_\ell. \end{aligned}$$

This can be rearranged to the following fundamental bound inequality:

$$\underbrace{a_\ell g(x_\ell)}_{\text{weighted gap}} - \underbrace{A_\ell(f(x_\ell) - f(x_{\ell+1}))}_{\text{weighted progress}} \leq E_\ell. \quad (27)$$

From this inequality there are various ways to proceed. The most natural way is often to plug-in the progress guarantee from the short-step for a direction d_t (typically $d_t = x_t - v_t$ but not always), that is, we use smoothness and optimize the bound:

$$f(x_{\ell+1}) - f(x_\ell) \leq -\frac{\langle \nabla f(x_\ell), d_\ell \rangle^2}{2L\|d_\ell\|^2}.$$

Chaining this bound into (27) makes the left-hand side only larger, so that we obtain:

$$\underbrace{a_\ell g(x_\ell)}_{\text{weighted gap}} - \underbrace{A_\ell \frac{\langle \nabla f(x_\ell), d_\ell \rangle^2}{2L\|d_\ell\|^2}}_{\text{weighted progress}} \leq E_\ell. \quad (28)$$

Inequalities (27) and (28) will be key in the following to simply transfer primal convergence results into primal-dual convergence results. To this end the following lemma will be useful, which turns the relation between a_ℓ and A_ℓ into a convergence rate.

Lemma E.1 (Linear rate conversion). *Suppose that $a_\ell(\kappa - 1) = A_{\ell-1}$ (or equivalently: $a_\ell\kappa = A_\ell$) then it holds*

$$1/A_\ell = \left(1 - \frac{1}{\kappa}\right) 1/A_{\ell-1}.$$

Proof. The proof is by straightforward rearranging:

$$\begin{aligned} a_\ell(\kappa - 1) &= A_{\ell-1} \\ \Rightarrow A_\ell(\kappa - 1) - A_{\ell-1}(\kappa - 1) &= A_{\ell-1} \\ \Rightarrow A_\ell(\kappa - 1) &= A_{\ell-1}\kappa \\ \Rightarrow A_\ell &= \frac{\kappa}{\kappa - 1} A_{\ell-1} \\ \Rightarrow 1/A_\ell &= \left(1 - \frac{1}{\kappa}\right) 1/A_{\ell-1} \end{aligned}$$

□

E.1 LINEAR CONVERGENCE RATES: THE CASE $E_\ell = 0$

We will first consider the case where we can set $E_\ell = 0$. This case, which usually comes with a simpler analysis basically captures all linear convergence rate results for FW algorithms. To this end we present a generic argument that then allows us to simply reuse already known analyses for primal convergence to establish primal-dual convergence simply by plugging in the bounds on the primal progress and the dual gap; it is important to note that these are usually known already from the primal analysis. We start from (27):

$$a_\ell g(x_\ell) - A_\ell(f(x_\ell) - f(x_{\ell+1})) \leq E_\ell,$$

and we set $E_\ell = 0$ and solve for equality. For the sake of exposition here and in the following let $p_\ell \stackrel{\text{def}}{=} f(x_\ell) - f(x_{\ell+1})$ denote the *primal progress* in iteration ℓ . We obtain:

$$a_\ell g(x_\ell) - A_\ell p_\ell = 0 \quad (29)$$

$$\Rightarrow a_\ell g(x_\ell) = A_\ell p_\ell \quad (30)$$

$$\Rightarrow a_\ell (g(x_\ell) - p_\ell) = A_{\ell-1} p_\ell \quad (31)$$

$$\Rightarrow a_\ell \left(\frac{g(x_\ell)}{p_\ell} - 1 \right) = A_{\ell-1}, \quad (32)$$

and combining this with Lemma E.1, we obtain:

Lemma E.2 (Transfer template for bounds). *Let $g(x)$ be a dual bound and let p_ℓ denote the primal progress in iteration ℓ . If we set $E_\ell = 0$, then we obtain:*

$$1/A_\ell = \left(1 - \frac{p_\ell}{g(x_\ell)}\right) 1/A_{\ell-1}.$$

Also observe that we can use (32) to dynamically compute the a_ℓ from an actual step which allows us to report an adaptive gap if desired.

E.1.1 f STRONGLY CONVEX AND $x^* \in \text{rel.int}(P)$

For the sake of completeness, we will present two ways of deriving primal-dual convergence rates for the case where f is strongly convex and $x^* \in \text{rel.int}(P)$. First, we will go the “complicated” route, without relying on Lemma E.2 and instead directly solving (28) for a_ℓ . Then, we will use Lemma E.2 to obtain the (same) convergence rate. The primal analysis is due to Guélat & Marcotte (1986).

We assume that f is μ -strongly convex and $B(x^*, r) \subseteq \mathcal{X}$. We start from the primal gap bound via strong convexity:

$$f(x) - f(x^*) \leq \frac{\langle \nabla f(x), x - x^* \rangle^2}{2\mu \|x - x^*\|^2},$$

where $\mu > 0$ is the strong convexity constant, which we then combine with the *scaling inequality* for $x^* \in \text{rel.int}(P)$ from Guélat & Marcotte (1986):

$$\frac{r}{D} \frac{\langle \nabla f(x), x - x^* \rangle}{\|x - x^*\|} \leq \frac{\langle \nabla f(x), x - v \rangle}{\|x - v\|},$$

where v is the FW-vertex for x , D is the diameter, and r is radius of the ball around the optimal solution x^* contained in P , to obtain the bound on the primal gap:

$$f(x) - f(x^*) \leq \frac{D^2}{r^2} \frac{\langle \nabla f(x), x - v \rangle^2}{2\mu \|x - v\|^2},$$

and we set the gap function $g(x_\ell)$ to be the right hand side of the above inequality:

$$g(x_\ell) \stackrel{\text{def}}{=} \frac{D^2}{r^2} \frac{\langle \nabla f(x_\ell), x_\ell - v_\ell \rangle^2}{2\mu \|x_\ell - v_\ell\|^2},$$

where x_ℓ, v_ℓ are the iterate and FW-vertex in the ℓ -th iteration, respectively. Now for the direction d_ℓ we pick the standard FW direction, i.e., $d_\ell = x_\ell - v_\ell$. With this choice (28) becomes

$$\begin{aligned} a_\ell g(x_\ell) - A_\ell \frac{\langle \nabla f(x_\ell), d_\ell \rangle^2}{2L \|d_\ell\|^2} &\leq E_\ell \\ \Rightarrow a_\ell \frac{D^2}{r^2} \frac{\langle \nabla f(x_\ell), x_\ell - v_\ell \rangle^2}{2\mu \|x_\ell - v_\ell\|^2} - A_\ell \frac{\langle \nabla f(x_\ell), x_\ell - v_\ell \rangle^2}{2L \|x_\ell - v_\ell\|^2} &\leq E_\ell \\ \Rightarrow \frac{\langle \nabla f(x_\ell), x_\ell - v_\ell \rangle^2}{\|x_\ell - v_\ell\|^2} \left(a_\ell \frac{1}{2\mu} \frac{D^2}{r^2} - A_\ell \frac{1}{2L} \right) &\leq E_\ell. \end{aligned}$$

We want $a_\ell \frac{1}{2\mu} \frac{D^2}{r^2} - A_\ell \frac{1}{2L} = 0$, and in particular then we can choose $E_\ell = 0$. We obtain:

$$\begin{aligned} a_\ell \frac{1}{2\mu} \frac{D^2}{r^2} - A_\ell \frac{1}{2L} &= 0 \\ \Rightarrow a_\ell \frac{1}{2\mu} \frac{D^2}{r^2} &= A_\ell \frac{1}{2L} \\ \Rightarrow a_\ell \left(\frac{1}{2\mu} \frac{D^2}{r^2} - \frac{1}{2L} \right) &= A_{\ell-1} \frac{1}{2L} \\ \Rightarrow a_\ell \left(\frac{L}{\mu} \frac{D^2}{r^2} - 1 \right) &= A_{\ell-1}, \end{aligned}$$

and applying [Lemma E.1](#), we obtain immediately that

$$1/A_\ell = \left(1 - \frac{\mu}{L} \frac{r^2}{D^2}\right) 1/A_{\ell-1},$$

and hence the expected linear rate of convergence follows for the primal-dual gap G_t , i.e.,

$$G_t \leq \frac{A_{t-1}}{A_t} G_{t-1} = \left(1 - \frac{\mu}{L} \frac{r^2}{D^2}\right) G_{t-1}.$$

Alternatively, we could have directly applied [Lemma E.2](#) to

$$p_\ell = \frac{\langle \nabla f(x_\ell), x_\ell - v_\ell \rangle^2}{2L \|x_\ell - v_\ell\|^2} \quad \text{and} \quad g(x_\ell) = \frac{D^2}{r^2} \frac{\langle \nabla f(x_\ell), x_\ell - v_\ell \rangle^2}{2\mu \|x_\ell - v_\ell\|^2},$$

to obtain the same result.

E.1.2 AWAY-STEP VARIANTS OVER POLYTOPES AND STRONGLY CONVEX FUNCTIONS

Let P be a polytope and f be μ -strongly convex. The proof is similar to [Appendix E.1.1](#). We pick

$$g(x_\ell) \stackrel{\text{def}}{=} \frac{\langle \nabla f(x_\ell), s_\ell - v_\ell \rangle^2}{2\mu w(P)^2},$$

which is the *geometric strong convexity* that arises from combining the strong convexity bound in [Appendix E.1.1](#) with a different scaling inequality, namely the one that appears in active set based, away-step inducing variants:

$$\langle \nabla f(x_\ell), s_\ell - v_\ell \rangle \geq w(P) \frac{\langle \nabla f(x_\ell), x_\ell - x^* \rangle}{\|x_\ell - x^*\|},$$

where $w(P)$ is the *pyramidal width* of P . Moreover, we pick $d_\ell = s_\ell - v_\ell$ and use a slightly different bound on the primal progress, namely

$$p_\ell = \frac{\langle \nabla f(x_\ell), s_\ell - v_\ell \rangle^2}{4L \|s_\ell - v_\ell\|^2},$$

note the extra factor of 2 in the denominator, which is due to the selection between the away-step or the FW step; we refer the reader to [Lacoste-Julien & Jaggi \(2015\)](#) and [Braun et al. \(2022\)](#) for more details on both the geometric strong convexity as well as the modified primal progress bound; see also [Peña & Rodríguez \(2018\)](#) for a unified perspective on the different notions of geometric conditioning arising from the polytope. We also ignore drop steps here for the sake of simplicity, however they can be added easily and deteriorate the number of required steps only by a constant factor of 2; see [Braun et al. \(2022\)](#) for more details.

We apply [Lemma E.2](#) to

$$p_\ell = \frac{\langle \nabla f(x_\ell), s_\ell - v_\ell \rangle^2}{4L \|s_\ell - v_\ell\|^2} \quad \text{and} \quad g(x_\ell) = \frac{\langle \nabla f(x_\ell), s_\ell - v_\ell \rangle^2}{2\mu w(P)^2},$$

to obtain

$$1/A_\ell = \left(1 - \frac{\mu}{L} \frac{w(P)^2}{2D^2}\right) 1/A_{\ell-1},$$

as expected but again for the primal-dual gap.

E.1.3 STRONGLY CONVEX FEASIBLE REGION AND $\|\nabla f(x)\| \geq c > 0$

We now consider the case where the feasible region is strongly convex and the unconstrained minimizer of f lies outside of the feasible region. Since the feasible region is strongly convex we obtain the scaling inequality ([Garber & Hazan, 2015](#); [Braun et al., 2022](#)):

$$\frac{\alpha}{4} \|\nabla f(x_t)\| \leq \frac{\langle \nabla f(x_t), x_t - v_t \rangle}{\|x_t - v_t\|^2}.$$

Moreover we have

$$f(x_t) - f(x^*) \leq \langle \nabla f(x_t), x_t - v_t \rangle.$$

Combining the two leads to:

$$(f(x_t) - f(x^*)) \frac{\alpha}{4} \|\nabla f(x_t)\| \leq \frac{\langle \nabla f(x_t), x_t - v_t \rangle^2}{\|x_t - v_t\|^2},$$

which gives the dual gap bound:

$$f(x_t) - f(x^*) \leq \frac{4}{\alpha \|\nabla f(x_t)\|} \frac{\langle \nabla f(x_t), x_t - v_t \rangle^2}{\|x_t - v_t\|^2}, \quad (33)$$

and we set

$$g(x_t) \stackrel{\text{def}}{=} \frac{4}{\alpha \|\nabla f(x_t)\|} \frac{\langle \nabla f(x_t), x_t - v_t \rangle^2}{\|x_t - v_t\|^2}.$$

Moreover, for the primal progress bound we simply pick the short-step with the standard direction $x_t - v_t$, i.e.,

$$p_t \stackrel{\text{def}}{=} \frac{\langle \nabla f(x_t), x_t - v_t \rangle^2}{2L \|x_t - v_t\|^2}.$$

Applying [Lemma E.2](#) we immediately obtain:

$$1/A_\ell = \left(1 - \frac{\alpha \|\nabla f(x_t)\|}{8L}\right) 1/A_{\ell-1},$$

and using $\|\nabla f(x)\| \geq c > 0$ we have

$$1/A_\ell \leq \left(1 - \frac{\alpha c}{8L}\right) 1/A_{\ell-1},$$

which completes the argument.

E.2 VARIANTS UTILIZING BREGMAN DIVERGENCES: THE CASE WHERE $E_\ell \neq 0$

We will now consider the slightly more involved case where $E_\ell \neq 0$. This means that we cannot simply apply [Lemma E.2](#) but rather have to account for the error term E_ℓ . To this end, we first collect alternative choices of a_t , A_t , and γ_t , assuming linear coupling $\gamma_t = a_t/A_t$, that will become useful in the rest of the section.

Lemma E.3 (Open-loop step-sizes). *Let $1 \leq \ell \in \mathbb{N}$ and define $\gamma_t = \frac{\ell}{t+\ell}$. Then*

$$a_t = a_0 \binom{t+\ell-1}{t} \quad \text{and} \quad A_t = a_0 \binom{t+\ell}{t},$$

is a valid choice and moreover it holds:

1. $A_t = a_t \frac{t+\ell}{\ell}$;
2. $\frac{a_t^2}{A_t} = \gamma_t a_t = \frac{\ell}{t+\ell} a_t$;
3. *With the choice $a_0 = \ell!$, we recover for $\ell = 2$ the known values $a_t = 2(t+2)$ and $A_t = (t+1)(t+2)$.*

Proof. Verified by straightforward calculation. □

The next lemma is folklore and we include it for completeness.

Lemma E.4 (Power Sum Estimations). *The following holds:*

1.

$$\sum_{t=0}^T t^\alpha \begin{cases} \approx \ln(T) + \kappa \leq \ln(T) + 1 & \text{for } \alpha = -1 \\ = T + 1 & \text{for } \alpha = 0 \\ \approx \frac{T^{\alpha+1}}{\alpha+1} + \frac{T^\alpha}{2} + O(T^{\alpha-1}) & \text{for } \alpha > 0 \end{cases},$$

where $\kappa \approx 0.577$ is the Euler-Mascheroni constant.

2.

$$\int_0^T x^\alpha dx = \frac{T^{\alpha+1}}{\alpha+1}$$

3.

$$\sum_{t=0}^T t^\alpha \leq \int_0^{T+1} x^\alpha dx = \frac{(T+1)^{\alpha+1}}{\alpha+1}$$

Proof. The approximation is via Faulhaber's formula. \square

As an example, in this section we specifically consider convergence rates arising from the so-called strong growth property as studied in Peña (2023) and Wirth et al. (2024) for FW algorithms. Our starting point is the following simple Bregman expansion of f :

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

We use the standard FW gap (or alternatively the Heavy-Ball variant) as a lower bound as before, so that the critical estimation becomes, with the usual $x_{t+1} = (1 - \gamma_t)x_t + \gamma_t v_t$:

$$\begin{aligned} & A_t(f(x_{t+1}) - f(x_t)) - a_t \langle \nabla f(x_t), v_t - x_t \rangle \\ &= A_t(\langle \nabla f(x_t), x_{t+1} - x_t \rangle + D_f(x_{t+1}, x_t)) - a_t \langle \nabla f(x_t), v_t - x_t \rangle. \end{aligned}$$

Instead of using the standard linear coupling $\gamma_t = a_t/A_t$, we use the modified linear coupling $\gamma_t = 2a_t/A_t$, so that we obtain:

$$A_t(\langle \nabla f(x_t), x_{t+1} - x_t \rangle + D_f(x_{t+1}, x_t)) - a_t \langle \nabla f(x_t), v_t - x_t \rangle \quad (34)$$

$$= A_t D_f(x_{t+1}, x_t) + 2a_t \langle \nabla f(x_t), v_t - x_t \rangle - a_t \langle \nabla f(x_t), v_t - x_t \rangle \quad (35)$$

$$= A_t D_f(x_{t+1}, x_t) - a_t \langle \nabla f(x_t), x_t - v_t \rangle. \quad (36)$$

The (M, r) -strong growth property asserts

$$D_f(x_{t+1}, x_t) \leq \gamma_t^2 M / 2 \langle \nabla f(x_t), x_t - v_t \rangle^r,$$

with $M > 0$ and $r \in [0, 1]$.

E.2.1 PRIMAL-DUAL CONVERGENCE $r = 1$

We start with the most obvious case $r = 1$, then (36) becomes:

$$\begin{aligned} A_t D_f(x_{t+1}, x_t) - a_t \langle \nabla f(x_t), x_t - v_t \rangle &\leq 2M \frac{a_t^2}{A_t} \langle \nabla f(x_t), x_t - v_t \rangle - a_t \langle \nabla f(x_t), x_t - v_t \rangle \\ &= \langle \nabla f(x_t), x_t - v_t \rangle a_t (2M \frac{a_t}{A_t} - 1), \end{aligned}$$

and hence choosing $a_t/A_t = \frac{1}{2M}$, which implies $\gamma_t = \frac{1}{M}$, the right-hand side becomes 0 and we obtain linear convergence with a rate $1/A_{t-1}(1 - \frac{1}{2M}) = 1/A_t$.

E.2.2 PRIMAL-DUAL CONVERGENCE $r \in [0, 1)$

The situation is much more complicated for $r \in (0, 1)$. In this case (36) becomes:

$$A_t D_f(x_{t+1}, x_t) - a_t \langle \nabla f(x_t), x_t - v_t \rangle \leq 2M \frac{a_t^2}{A_t} \langle \nabla f(x_t), x_t - v_t \rangle^r - a_t \langle \nabla f(x_t), x_t - v_t \rangle$$

We now apply Young's inequality with $p = \frac{1}{r}$, $q = \frac{1}{1-r}$, rebalancing with $\frac{a_t^r}{r}$, so that we obtain:

$$\begin{aligned}
& 2M \frac{a_t^2}{A_t} \langle \nabla f(x_t), x_t - v_t \rangle^r - a_t \langle \nabla f(x_t), x_t - v_t \rangle \\
&= 2M \frac{a_t^2}{A_t} \frac{r^r}{a_t^r} \langle \nabla f(x_t), x_t - v_t \rangle^r - a_t \langle \nabla f(x_t), x_t - v_t \rangle \\
&\leq \frac{\left(2M \frac{a_t^2}{A_t} \frac{r^r}{a_t^r}\right)^q}{q} + \frac{\left(\frac{a_t^r}{r^r} \langle \nabla f(x_t), x_t - v_t \rangle^r\right)^{1/r}}{1/r} - a_t \langle \nabla f(x_t), x_t - v_t \rangle \\
&= (1-r) \left(2M \frac{a_t^2}{A_t} \frac{r^r}{a_t^r}\right)^{\frac{1}{1-r}} = (1-r) (2Mr^r)^{\frac{1}{1-r}} \left(\frac{a_t^{2-r}}{A_t}\right)^{\frac{1}{1-r}}
\end{aligned}$$

Summing the errors as customary, we obtain:

$$\frac{1}{A_T} \sum_{\ell=0}^T E_\ell \leq (1-r) (2Mr^r)^{\frac{1}{1-r}} \frac{1}{A_T} \sum_{t=0}^T \left(\frac{a_t^{2-r}}{A_t}\right)^{\frac{1}{1-r}},$$

which now needs to be estimated to derive the convergence rate.

Estimation in the order We will first do an approximate estimation, to recover the order of convergence and provide intuition. Suppose that $a_t = t^c$ and hence $A_t \approx \frac{t^{c+1}}{c+1} = \Theta(t^{c+1})$. With this we obtain:

$$\begin{aligned}
& (1-r) (2Mr^r)^{\frac{1}{1-r}} \frac{c+1}{T^{c+1}} \sum_{t=0}^T \left(\frac{t^{c(2-r)}}{t^{c+1}}\right)^{\frac{1}{1-r}} \\
&= (1-r) (2Mr^r)^{\frac{1}{1-r}} \frac{c+1}{T^{c+1}} \sum_{t=0}^T \left(t^{c(2-r)-c-1}\right)^{\frac{1}{1-r}} \\
&\approx (1-r) (2Mr^r)^{\frac{1}{1-r}} \frac{c+1}{T^{c+1}} \frac{T^{c-\frac{1}{1-r}+1}}{c-\frac{1}{1-r}+1} = (1-r) (2Mr^r)^{\frac{1}{1-r}} \frac{c+1}{c-\frac{1}{1-r}+1} T^{-\frac{1}{1-r}},
\end{aligned}$$

assuming $c > \frac{1}{1-r}$. Therefore we obtain for the total convergence rate:

$$G_t \leq \frac{A_0 G_0}{A_T} + (1-r) (2Mr^r)^{\frac{1}{1-r}} \frac{c+1}{c-\frac{1}{1-r}+1} T^{-\frac{1}{1-r}},$$

together with $1/A_T \approx \frac{c+1}{T^{c+1}} \leq \frac{c+1}{T^{\frac{1}{1-r}+1}}$, implies via a rather weak estimation

$$G_t \leq 2(1-r) (2Mr^r)^{\frac{1}{1-r}} \frac{c+1}{c-\frac{1}{1-r}+1} T^{-\frac{1}{1-r}},$$

Specific choice of a_t We will now refine the analysis from above via a specific choice of a_t and A_t . We choose $a_t = a_0 \binom{t+\ell-1}{t}$ and $A_t = a_0 \binom{t+\ell}{t}$ similar to [Lemma E.3](#). Note however the changed linear coupling $2a_t/A_t = \gamma_t$ that we use here:

$$\begin{aligned}
& (1-r) (2Mr^r)^{\frac{1}{1-r}} \frac{1}{A_T} \sum_{t=0}^T \left(\frac{a_t^{2-r}}{A_t} \right)^{\frac{1}{1-r}} \\
&= (1-r) (2Mr^r)^{\frac{1}{1-r}} \frac{1}{A_T} \sum_{t=0}^T \left(\frac{a_t}{A_t} \right)^{\frac{1}{1-r}} a_t \\
&= (1-r) (2Mr^r)^{\frac{1}{1-r}} a_0^{-1} \binom{T+\ell}{T}^{-1} \sum_{t=0}^T \left(\frac{\ell}{t+\ell} \right)^{\frac{1}{1-r}} a_0 \binom{t+\ell-1}{t} \\
&= (1-r) (2Mr^r)^{\frac{1}{1-r}} \binom{T+\ell}{T}^{-1} \sum_{t=0}^T \left(\frac{\ell}{t+\ell} \right)^{\frac{1}{1-r}} \binom{t+\ell-1}{t} \\
&\leq (1-r) (2Mr^r)^{\frac{1}{1-r}} \binom{T+\ell}{T}^{-1} \frac{\ell^{\frac{1}{1-r}}}{(\ell-1)!} \sum_{t=0}^T t^{\ell-1-\frac{1}{1-r}} \\
&= (1-r) (2Mr^r)^{\frac{1}{1-r}} \binom{T+\ell}{T}^{-1} \frac{\ell^{\frac{1}{1-r}}}{(\ell-1)!} \left(\frac{T^{\ell-\frac{1}{1-r}}}{\ell-\frac{1}{1-r}} + O\left(T^{\ell-1-\frac{1}{1-r}}\right) \right) \\
&\leq (1-r) (2Mr^r)^{\frac{1}{1-r}} \frac{\ell^\ell}{T^\ell} \frac{\ell^{\frac{1}{1-r}}}{(\ell-1)!} \left(\frac{T^{\ell-\frac{1}{1-r}}}{\ell-\frac{1}{1-r}} + O\left(T^{\ell-1-\frac{1}{1-r}}\right) \right) \\
&= (1-r) (2Mr^r)^{\frac{1}{1-r}} \frac{\ell^{\ell+\frac{1}{1-r}}}{(\ell-1)!} \left(\frac{T^{-\frac{1}{1-r}}}{\ell-\frac{1}{1-r}} + O\left(T^{-1-\frac{1}{1-r}}\right) \right) \\
&= (1-r) (2Mr^r)^{\frac{1}{1-r}} \frac{\ell^{\ell+\frac{1}{1-r}}}{(\ell-\frac{1}{1-r})(\ell-1)!} T^{-\frac{1}{1-r}} + O\left(T^{-1-\frac{1}{1-r}}\right),
\end{aligned}$$

with $\ell > 1 + \frac{1}{1-r}$. Therefore we obtain for the total convergence rate:

$$G_t \leq \frac{A_0 G_0}{A_T} + (1-r) (2Mr^r)^{\frac{1}{1-r}} \frac{\ell^{\ell+\frac{1}{1-r}}}{(\ell-\frac{1}{1-r})(\ell-1)!} T^{-\frac{1}{1-r}} + O\left(T^{-1-\frac{1}{1-r}}\right),$$

together with $1/A_T \leq \ell^\ell T^{-\ell} \leq \ell^\ell T^{-1-\frac{1}{1-r}}$, this implies

$$(1-r) (2Mr^r)^{\frac{1}{1-r}} \frac{\ell^{\ell+\frac{1}{1-r}}}{(\ell-\frac{1}{1-r})(\ell-1)!} T^{-\frac{1}{1-r}} + O\left(T^{-1-\frac{1}{1-r}}\right).$$

Remark E.5 (Step-size choice and short steps). *Note, that we can always use short-steps instead of the $\frac{\ell}{t+\ell}$ -step size by monotonicity and obtain the same rates.*

F EXPERIMENTS: APPENDIX

We performed additional experiments to test convergence and isolate the effect of the different components. Note that for the additional experiments below, we sometimes have used additional LMO calls for computing evaluation metrics for comparisons, so that the timings in the right columns are not as reliable and not comparable to the main experiments. Either way, the performance in iterations gives a clear picture here as the cost per iteration of all algorithms is quite comparable, apart from the adaptive variant which behaves more like line search but serves as a baseline only.

PRIMAL-DUAL SHORT-STEPS

We conducted a series of experiments to evaluate the performance of the primal-dual short step method described in [Section 4](#). Unfortunately, our findings indicate that this variant does not outperform the vanilla FW variant with short-steps (or line search). The primary reason for this effect appears to be the improved primal-dual gap measure, which results in a *smaller* step size, simply by how the step

size is computed. Consequently, this leads to slower convergence rates. This behavior is reminiscent of the lower bound discussed in (Guélat & Marcotte, 1986), and thus, it is not entirely unexpected. Moreover, by focusing on large primal-dual progress, the algorithm can overshoot in terms of primal progress, so that while the primal-dual progress is relatively large, we are still converging slowly (assuming that the primal-dual gap converges at a certain rate).

To illustrate this point, we present a typical experiment in Figure 3. This figure is representative of the results we observed across various tests, consistently showing that the primal-dual short step method converges at a similar rate to the vanilla FW variant. For completeness, we also include the results for the optimistic variant, which indeed outperforms the other variants.

HEAVY BALL LOWER BOUND VS OPTIMISM

In our experiments, we also investigated whether the heavy ball lower bound is inherently stronger than the vanilla FW lower bound, which might have implied that optimism does not actually help at all. However, our findings indicate that this is not the case; see Figure 4 for a representative experiment. Specifically, when the heavy ball lower bound is applied to the vanilla FW variant, it actually performs worse than the standard FW gap in our tests. This observation strongly suggests that the superior performance of the optimistic variant is due to its trajectory rather than the strength of the lower bound itself.

ADDITIONAL EXPERIMENTS

We present two additional experiments in Figures 5 and 6 that had to be relegated to the appendix due to space constraints. Also here the results are consistent with our findings in the main text.




Figure 3: Comparison over the probability simplex of dimension $n = 1000$ with objective $f(x) = \|x - x_0\|_2^2$, where x_0 is a random point outside the probability simplex. We see that primal-dual short steps converge exactly as fast as the vanilla variant.

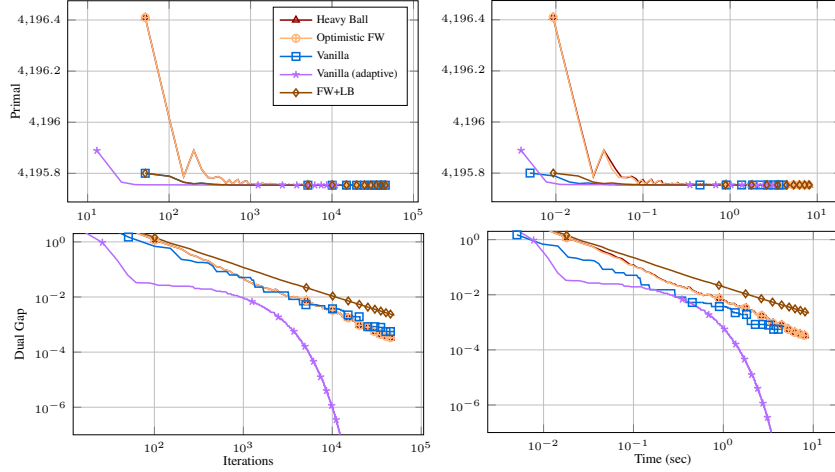


Figure 4: Comparison over the probability simplex of dimension $n = 1000$ with objective $f(x) = \|x - x_0\|_2^2$, where x_0 is a random point outside the probability simplex. FW+LB denotes the standard FW variant with the heavy ball lower bound. We can see that FW with the heavy ball lower bound converges no faster than the normal FW in dual gap.

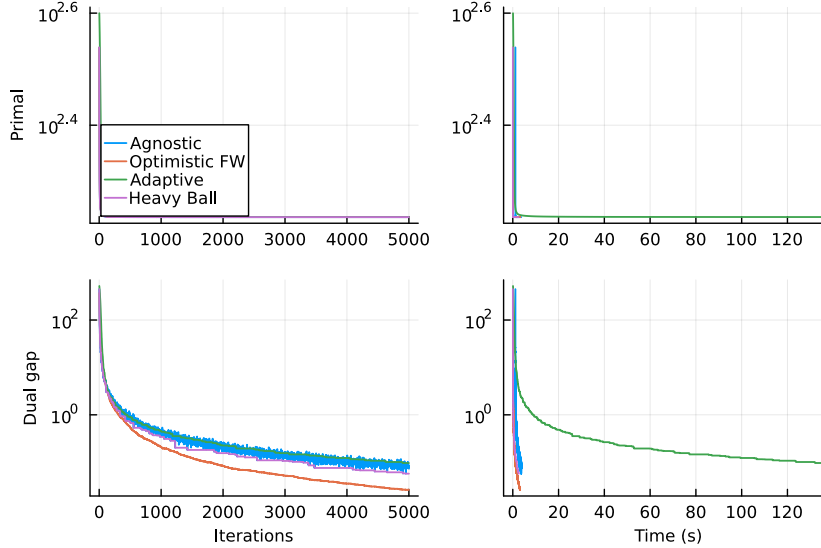


Figure 5: Comparison over the probability simplex and with a 1000×1500 matrix, with the negative log-likelihood of the portfolio returns. We can see that the optimistic variant converges faster than the other variants both in iterations and time.




Figure 6: Comparison over the probability simplex and with a 1000×1500 matrix, with the negative log-likelihood of the portfolio returns. We can see that the optimistic variant converges faster than the other variants both in iterations and time.

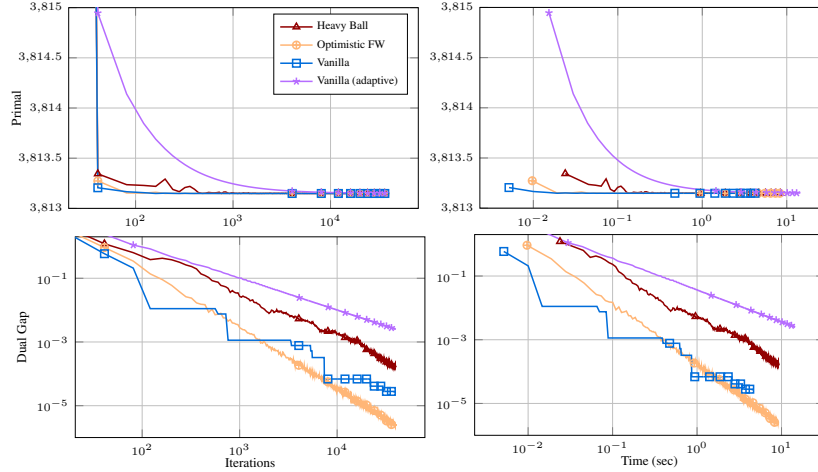


Figure 7: Comparison over the probability simplex of dimension $n = 100$ with objective $f(x) = \|Ax - b\|_2^2$, where A and b are random. We can see that the optimistic variant converges faster than the other variants both in iterations and time.

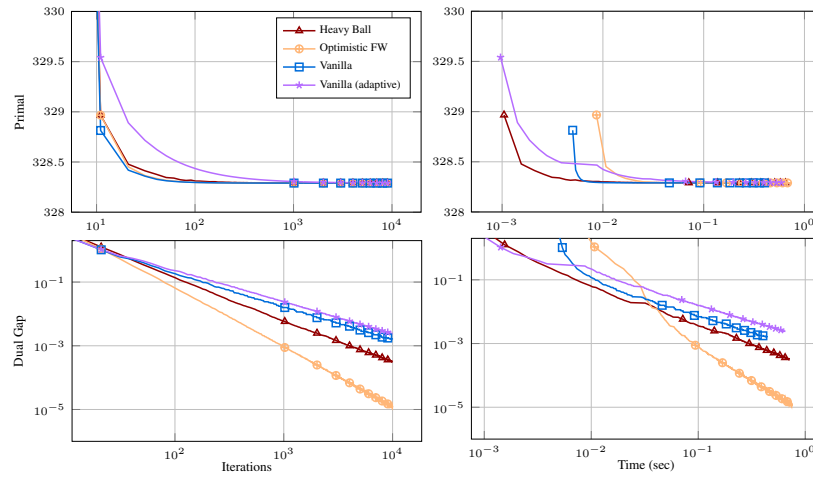


Figure 8: Comparison over k -sparse polytope of dimension $n = 1000$ and $k = 10$ with objective $f(x) = \|x - x_0\|_2^2$, where x_0 is a random point outside the k -sparse polytope. We can see that the optimistic variant converges faster than the other variants both in iterations and time.