

Analysis of the Frank-Wolfe Method for Convex Composite Optimization involving a Logarithmically-Homogeneous Barrier*

Renbo Zhao[†]

Robert M. Freund[‡]

July 2, 2021

Abstract

We present and analyze a new generalized Frank-Wolfe method for the composite optimization problem $(P) : \min_{x \in \mathbb{R}^n} f(Ax) + h(x)$, where f is a θ -logarithmically-homogeneous self-concordant barrier, A is a linear operator and the function h has bounded domain but is possibly non-smooth. We show that our generalized Frank-Wolfe method requires $O((\delta_0 + \theta + R_h) \ln(\delta_0) + (\theta + R_h)^2/\varepsilon)$ iterations to produce an ε -approximate solution, where δ_0 denotes the initial optimality gap and R_h is the variation of h on its domain. This result establishes certain intrinsic connections between θ -logarithmically homogeneous barriers and the Frank-Wolfe method. When specialized to the D -optimal design problem, we essentially recover the complexity obtained by Khachiyan [35] using the Frank-Wolfe method with exact line-search. We also study the (Fenchel) dual problem of (P) , and we show that our new method is equivalent to an adaptive-step-size mirror descent method applied to the dual problem. This enables us to provide iteration complexity bounds for the mirror descent method despite that the dual objective function is non-Lipschitz and has unbounded domain. In addition, we present computational experiments that point to the potential usefulness of our generalized Frank-Wolfe method on Poisson image de-blurring problems with TV regularization, and on simulated PET problem instances.

Keywords: Frank-Wolfe method, composite optimization, logarithmic-homogeneity, self-concordance, barrier, complexity analysis.

1 Introduction

We present and analyze a new generalized Frank-Wolfe method [8, 14, 17–19, 23, 38] for the following composite optimization problem:

$$(P) : F^* := \min_{x \in \mathbb{R}^n} [F(x) := f(Ax) + h(x)] \quad (1.1)$$

where $f : \mathbb{R}^m \rightarrow \mathbb{R} \cup \{+\infty\}$ is an extended real-valued convex function that is differentiable on its domain $\text{dom } f := \{u \in \mathbb{R}^m : f(u) < +\infty\}$, $A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a linear operator (though not

*Research supported by AFOSR Grant No. FA9550-19-1-0240.

[†]MIT Operations Research Center, 77 Massachusetts Avenue, Cambridge, MA 02139 (mailto: renboz@mit.edu).

[‡]MIT Sloan School of Management, 77 Massachusetts Avenue, Cambridge, MA 02139 (mailto: rfreund@mit.edu).

necessarily invertible or surjective) and the function $h : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$ is proper, closed and convex (but possibly non-smooth), for which $\text{dom } h$ is a nonempty compact convex set. Furthermore, and in contrast to the standard setting where f is assumed to be L -smooth on $\text{dom } F$ (i.e., its gradient is L -Lipschitz on $\text{dom } F$), our focus is on the setting where f belongs to a particularly special and important class of functions that arise in practice, namely θ -logarithmically-homogeneous self-concordant barrier functions (whose definition and properties will be reviewed below). For convenience we distinguish between $u \mapsto f(u)$ and $x \mapsto f(Ax)$ by defining

$$\bar{f}(x) := f(Ax) . \tag{1.2}$$

The Frank-Wolfe method was developed in 1956 in the seminal paper of Frank and Wolfe [23] for the case $h = \iota_{\mathcal{X}}$, where $\iota_{\mathcal{X}}$ denotes the indicator function of \mathcal{X} (i.e., $\iota_{\mathcal{X}}(x) = 0$ for $x \in \mathcal{X}$ and $+\infty$ otherwise), and \mathcal{X} is a (bounded) polytope. (In particular, (P) then is the constrained problem $\min_{x \in \mathcal{X}} \bar{f}(x)$.) The Frank-Wolfe method was a significant focus of research up through approximately 1980, during which time it was generalized to handle more general compact sets \mathcal{X} , see e.g., [14, 17–19]. Each iteration of the Frank-Wolfe method computes the minimizer of the first-order (gradient) approximation of $\bar{f}(x)$ on \mathcal{X} , and constructs the next iterate by moving towards the minimizer. Just in the last decade, and due to the importance of optimization in modern computational statistics and machine learning, the Frank-Wolfe method has again attracted significant interest (see [24–26, 30, 33, 44] among many others), for at least two reasons. First, in many modern application settings, \mathcal{X} is a computationally “simple” set for which the linear optimization sub-problem in the Frank-Wolfe method is easier to solve compared to the standard projection sub-problems required by other first-order methods, see e.g., [43]. Secondly, the Frank-Wolfe method naturally produces “structured” (such as sparse, or low-rank) solutions in several important settings, which is very useful in the high-dimensional regime in machine learning. This is because each iterate of the Frank-Wolfe method is a convex combination of all the previous linear optimization sub-problem solutions, and hence if the extreme points of \mathcal{X} are unit coordinate vectors (for example), then the k -th iterate will have at most k non-zeroes. This statement can be made more precise in specific settings, and also generalizes to the matrix setting where \mathcal{X} is the nuclear norm ball [30] — in this setting the k -th (matrix) iterate will have rank at most k .

More recently, the Frank-Wolfe method has been generalized to the composite setting, where the function h is a general convex non-smooth function with compact domain \mathcal{X} , see e.g., [3, 26, 44]. In this generalized framework, the sub-problem solved at iteration k is

$$\text{minimize}_{x \in \mathbb{R}^n} \langle \nabla \bar{f}(x^k), x \rangle + h(x) , \tag{1.3}$$

which specializes to the standard Frank-Wolfe sub-problem in the case when $h = \iota_{\mathcal{X}}$. In certain situations, this minimization problem admits (relatively) easily computable solutions despite the presence of the non-smooth function h . For example, if $h = \bar{h} + \iota_{\mathcal{P}}$, where \bar{h} is a polyhedral function and \mathcal{P} is a polytope, then (1.3) can be reformulated as a linear optimization problem (LP), which can be solved efficiently if it has moderate size or a special structure, e.g., network flow structure [30]. For more such examples we refer the reader to [44].

When analyzing the convergence of the standard or generalized Frank-Wolfe method, almost all such analyses rely on the L -smooth assumption of the function f . Perhaps accidentally, the first specific attempt to extend the Frank-Wolfe method and analysis beyond the case of L -smooth functions is due to Khachiyan [35]. In the specific case of the D -optimal design problem [22], Khachiyan [35] developed a “barycentric coordinate descent” method with an elegant computational

complexity analysis, and it turns out that this method is none other than the Frank-Wolfe method with exact line-search [53, 57]. Khachiyan’s proof of his complexity result (essentially $O(n^2/\varepsilon)$ iterations) used clever arguments that do not easily carry over elsewhere, and hence begged the question of whether or not any of the arguments in [35] might underly any general themes beyond D -optimal design, and if so what might be the mathematical structures driving any such themes? In this work, we provide affirmative answers to these questions, by considering the D -optimal design problem as a special instance of the broader class of composite optimization problem (P) . The second attempt was the recent paper of Dvurechensky et al. [20], which presented and analyzed an adaptive step-size Frank-Wolfe method for tackling the problem $\min_{x \in \mathcal{X}} \bar{f}(x)$ where \bar{f} is assumed to be a non-degenerate (i.e., with positive-definite Hessians) self-concordant function and \mathcal{X} is a compact convex set, and was the first paper to study the Frank-Wolfe method for these special functions. The set-up in [20] can be seen as an instance of (P) with h being the indicator function of \mathcal{X} , namely, $h = \iota_{\mathcal{X}}$, and the additional assumption that \bar{f} is non-degenerate, which we do not need. (Note that in our setting, this amounts to assuming that $\mathbf{A} = \mathbf{I}$, namely the identity operator or that the linear operator \mathbf{A} is invertible.) However, unlike [20], we additionally assume that f is θ -logarithmically homogeneous. As our analysis will show, this last property — which holds true for all applications that we are aware of — is the key property that leads to relatively simple and natural computational guarantees for the Frank-Wolfe method in this expanded relevant setting.

Let us now review the formal definition of a θ -logarithmically-homogeneous self-concordant barrier function. Let $\mathcal{K} \subsetneq \mathbb{R}^m$ be a regular cone, i.e., \mathcal{K} is closed, convex, pointed (\mathcal{K} contains no line), and has nonempty interior ($\text{int } \mathcal{K} \neq \emptyset$). We say that f is a θ -logarithmically-homogeneous (non-degenerate) self-concordant barrier on \mathcal{K} for some $\theta \geq 1$ and we write “ $f \in \mathcal{B}_{\theta}(\mathcal{K})$ ”, if f is three-times differentiable and strictly convex on $\text{int } \mathcal{K}$ and satisfies the following three properties:

$$(P1) \quad |D^3 f(u)[w, w, w]| \leq 2(\langle H(u)w, w \rangle)^{3/2} \quad \forall u \in \text{int } \mathcal{K}, \forall w \in \mathbb{R}^m,$$

$$(P2) \quad f(u_k) \rightarrow \infty \text{ for any } \{u_k\}_{k \geq 1} \subseteq \text{int } \mathcal{K} \text{ such that } u_k \rightarrow u \in \text{bd } \mathcal{K}, \text{ and}$$

$$(P3) \quad f(tu) = f(u) - \theta \ln(t) \quad \forall u \in \text{int } \mathcal{K}, \forall t > 0,$$

where $H(u)$ denotes the Hessian of f at $u \in \text{int } \mathcal{K}$. For details on these properties, we refer readers to Nesterov and Nemirovski [45, Section 2.3.3] and Renegar [48, Section 2.3.5]. Properties (P1) and (P2) correspond to f being a (standard, strongly) self-concordant function on $\text{int } \mathcal{K}$ (cf. [45, Remark 2.1.1]), and property (P3) corresponds to f being a θ -logarithmically-homogeneous barrier function on \mathcal{K} . Here θ is called the *complexity parameter* of f in the terminology of Renegar [48]. The two prototypical examples of such functions are (i) $-\ln \det(U)$ for $U \in \mathcal{K} := \mathbb{S}_+^k$ and $\theta = k$, and (ii) $-\sum_{j=1}^m w_j \ln(u_j)$ for $u \in \mathcal{K} := \mathbb{R}_+^m$ and $\theta = \sum_{j=1}^m w_j$ where $w_1, \dots, w_n \geq 1$, see [45, 48].

We now present some application examples of (P) where $f \in \mathcal{B}_{\theta}(\mathcal{K})$, including the aforementioned D -optimal design problem.

1.1 Applications

1. *Poisson image de-blurring with total variation (TV) regularization* [10, 15, 32]. In image processing let the $m \times n$ matrix X be the true representation of an image, such that each entry $X_{ij} \geq 0$ represents the intensity of the pixel at location $(i, j) \in [m] \times [n]$, and $X_{ij} \in \{0, 1, \dots, M\}$, where $M := 2^b - 1$ for b -bit images. In many applications, ranging from microscopy to astronomy, we observe a blurred image contaminated by Poisson noise, which we denote by Y , and we wish

to estimate the true image X from Y . The generative model of Y from X is presumed to be as follows. Let $\mathbf{A} : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}$ denote the 2D discrete convolutional (linear) operator with periodic boundary conditions, which is assumed to be known. This convolutional operator is defined by a $p \times p$ 2D convolutional kernel with a size $q := p^2$ that is typically much smaller than N . (For an illustration of the 2D convolution, see [31] for example.) The blurred image \tilde{Y} is obtained by passing X through \mathbf{A} , i.e., $\tilde{Y} := \mathbf{A}(X)$, and the observed image Y results from adding independent entry-wise Poisson noise to \tilde{Y} , i.e., $Y_{ij} \sim \text{Poiss}(\tilde{Y}_{ij})$, for all $(i, j) \in [m] \times [n]$, and $\{Y_{ij}\}_{(i,j) \in [m] \times [n]}$ are assumed to be independent.

It will be preferable to work with vectors in addition to matrices, whereby we equivalently describe the above model using vector notation as follows. We denote $X = [x_1 \cdots x_m]^\top$, where x_i^\top denotes the i -th row of X , and let $\text{vec} : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{mn}$ denote the vectorization operator that sequentially concatenates X into the column vector $\text{vec}(X) := x := [x_1^\top \cdots x_m^\top]^\top$, and let $\text{mat}(x)$ denote the inverse operator of vec , so that $\text{mat}(x) = X$. Define $y := \text{vec}(Y)$ and $\tilde{y} := \text{vec}(\tilde{Y})$. In addition, we represent \mathbf{A} in its matrix form $A \in \mathbb{R}^{N \times N}$, where $N := mn$, such that $\tilde{y} := Ax$. Furthermore, let us represent $A := [a_1 \cdots a_N]^\top$, where a_l^\top denotes the l -th row of A for $l \in [N]$. Note that A is a sparse doubly-block-circulant matrix, such that each row a_l^\top of A has at most q non-zeros, where $q \ll N$ denotes the size of the 2D convolution kernel. Finally, we have $y_l \sim \text{Poiss}(\tilde{y}_l)$ for all $l \in [N]$, and $\{y_l\}_{l \in [n]}$ are independent.

The maximum likelihood (ML) estimator of X from the observed image Y is the optimal solution of the following optimization problem:

$$\min_{x \in \mathbb{R}^N} - \sum_{l=1}^N y_l \ln(a_l^\top x) + (\sum_{l=1}^N a_l)^\top x \quad \text{s. t. } 0 \leq x \leq Me, \quad (1.4)$$

where e denotes the vector with all entries equal to one. In addition, following [49], in order to recover a smooth image with sharp edges, we add Total Variation (TV) regularization to the objective function in (1.4), which yields the following regularized ML estimation problem:

$$\min_{x \in \mathbb{R}^N} \bar{F}(x) := - \sum_{l=1}^N y_l \ln(a_l^\top x) + (\sum_{l=1}^N a_l)^\top x + \lambda \text{TV}(x) \quad \text{s. t. } 0 \leq x \leq Me, \quad (1.5)$$

where

$$\begin{aligned} \text{TV}(x) &:= \sum_{i=1}^m \sum_{j=1}^{n-1} |[\text{mat}(x)]_{i,j} - [\text{mat}(x)]_{i,j+1}| + \sum_{i=1}^{m-1} \sum_{j=1}^n |[\text{mat}(x)]_{i,j} - [\text{mat}(x)]_{i+1,j}| \\ &:= \sum_{i=1}^m \sum_{j=1}^{n-1} |X_{i,j} - X_{i,j+1}| + \sum_{i=1}^{m-1} \sum_{j=1}^n |X_{i,j} - X_{i+1,j}| \end{aligned}$$

is a standard formulation of the total variation. Here we see that (1.5) is an instance of (P) with $f(u) := - \sum_{l=1}^N y_l \ln(u_l)$, $\mathcal{K} := \mathbb{R}_+^N$, $h(x) := (\sum_{l=1}^N a_l)^\top x + \lambda \text{TV}(x) + \iota_{\mathcal{X}}$ where $\mathcal{X} = \{x \in \mathbb{R}^N : 0 \leq x \leq Me\}$, \mathbf{A} is defined by $(Ax)_l := a_l^\top x$, $l = 1, \dots, N$, and $\theta = \sum_{l=1}^N y_l$. We note that $y_l \geq 1$ whenever $y_l \neq 0$ for all $l \in [N]$, and hence $f \in \mathcal{B}_\theta(\mathcal{K})$. In Section 4.1 we will discuss how the Frank-Wolfe sub-problem (1.3) associated with (1.5) can be efficiently solved.

2. Positron emission tomography (PET) [6, 51]. PET is a medical imaging technique that measures the metabolic activities of human tissues and organs. Typically, radioactive materials are injected into the organ of interest, and these materials emit (radioactive) events that can be detected by PET scanners. The mathematical model behind this process is described as follows. Suppose that an emission object (e.g., a human organ) has been discretized into n voxels. The number of events emitted by voxel i ($i \in [n]$) is a Poisson random variable \tilde{X}_i with *unknown* mean $x_i \geq 0$ and so $\tilde{X}_i \sim \text{Poiss}(x_i)$, and furthermore $\{\tilde{X}_i\}_{i=1}^n$ are assumed to be independent. We also have a scanner

array with m bins. Each event emitted by voxel i has a *known* probability p_{ij} of being detected by bin j ($j \in [m]$), and we assume that $\sum_{j=1}^m p_{ij} = 1$, i.e., the event will be detected by exactly one bin. Let \tilde{Y}_j denote the total number of events detected by bin j , whereby

$$\mathbb{E}[\tilde{Y}_j] := y_j := \sum_{i=1}^n p_{ij} x_i . \quad (1.6)$$

By Poisson thinning and superposition, it follows that $\{\tilde{Y}_j\}_{j=1}^m$ are independent random variables and $\tilde{Y}_j \sim \text{Poiss}(y_j)$ for all $j \in [m]$.

We seek to perform maximum-likelihood (ML) estimation of the unknown means $\{x_i\}_{i=1}^n$ based on observations $\{Y_j\}_{j=1}^m$ of the random variables $\{\tilde{Y}_j\}_{j=1}^m$. From the model above, we easily see that the log-likelihood of observing $\{Y_j\}_{j=1}^m$ given $\{\tilde{X}_i\}_{i=1}^n$ is (up to some constants)

$$l(x) := - \sum_{i=1}^n x_i + \sum_{j=1}^m Y_j \ln \left(\sum_{i=1}^n p_{ij} x_i \right) , \quad (1.7)$$

and therefore an ML estimate of $\{x_i\}_{i=1}^n$ is given by an optimal solution x^* of

$$\max_{x \geq 0} l(x) . \quad (1.8)$$

It follows from the first-order optimality conditions that any optimal solution x must satisfy

$$\sum_{i=1}^n x_i = S := \sum_{j=1}^m Y_j , \quad (1.9)$$

and by incorporating (1.9) into (1.8) and defining the re-scaled variable $z := x/S$, (1.8) can be equivalently written as

$$\min L(z) := - \sum_{j=1}^m Y_j \ln \left(\sum_{i=1}^n p_{ij} z_i \right) \quad \text{s. t. } z \in \Delta_n , \quad (1.10)$$

where $\Delta_n := \{z \in \mathbb{R}^n : \sum_{i=1}^n z_i = 1, z \geq 0\}$ is the unit simplex in \mathbb{R}^n . Here we see that (1.10) is an instance of (1.1) with $f(u) := - \sum_{j=1}^m Y_j \ln(u_j)$, $\mathcal{K} := \mathbb{R}_+^m$, $h := \iota_{\Delta_n}$, \mathbf{A} defined by $(\mathbf{A}z)_j := \sum_{i=1}^n p_{ij} z_i$, $j = 1, \dots, m$, and $\theta = \sum_{j=1}^m Y_j$. We note that $Y_j \geq 1$ whenever $Y_j \neq 0$ for all $j \in [m]$, and hence $f \in \mathcal{B}_\theta(\mathcal{K})$.

3. Optimal expected log investment [12, 55]. In this problem we consider n stocks in the market, and let R_i denote the random per-unit capital return on investing in stock i , for $i \in [n]$. The random vector $R := (R_1, \dots, R_n)$ has unknown distribution P . An investor allocates her investment capital over these n stocks, and let w_i denote the (nonnegative) proportion of capital invested in stock i , whereby $w_i \geq 0$ for all $i \in [n]$ and $\sum_{i=1}^n w_i = 1$. Define $w := (w_1, \dots, w_n)$. The goal of the investor is to maximize her expected log return $f(w) := \mathbb{E}_{R \sim P}[\ln(w^\top R)]$ subject to the constraint $w \in \Delta_n$ where $\Delta_n := \{w \in \mathbb{R}^n : w \geq 0, e^\top w = 1\}$. The naturalness of this objective can be justified from several perspectives involving the principle that money compounds multiplicatively rather than additively, see the discussion and references in Cover [12]. Since P is unknown, one can use a (historical) data-driven empirical distribution such as $\hat{P}_m := \sum_{j=1}^m p_j \delta_{r_j}$, where $p_j > 0$, $\sum_{j=1}^m p_j = 1$, $r_j \in \mathbb{R}^n$ is a realization of R and δ_{r_j} denotes the unit point mass at r_j for $j \in [m]$. Under this empirical distribution, the investor instead solves the problem:

$$\min_{w \in \Delta_n} - \sum_{j=1}^m p_j \ln(r_j^\top w) . \quad (1.11)$$

Define $p_{\min} := \min_{j \in [m]} \{p_j\} > 0$ and consider re-scaling the objective function of (1.11) by $1/p_{\min}$, which ensures the coefficient in front of each $\ln(\cdot)$ term is at least 1. Then this re-scaled problem

is an instance of (P) with $f(u) := -\sum_{j=1}^m (p_j/p_{\min}) \ln(u_j)$, $\mathcal{K} := \mathbb{R}_+^m$, $h := \iota_{\Delta_n}$, \mathbf{A} defined by $(\mathbf{A}w)_j := r_j^\top w$ for $j \in [m]$, $\theta = 1/p_{\min}$, and $f \in \mathcal{B}_\theta(\mathcal{K})$.

4. *Computing the analytic center* [42]. Given a nonempty solid polytope $\mathcal{Q} = \{x \in \mathbb{R}^n : a_i^\top x \geq d_i, i = 1, \dots, m\}$, the function

$$b(x) := -\sum_{i=1}^m \ln(a_i^\top x - d_i), \quad x \in \mathcal{Q}, \quad (1.12)$$

is an m -self-concordant barrier \mathcal{Q} , see [45]. We wish to compute the analytic center of \mathcal{Q} under b , which is the optimal solution to the problem $\min_{x \in \mathcal{Q}} b(x)$. We can transform this problem into an instance of (P) as follows. Define

$$f(x, t) := -\sum_{i=1}^m \ln(a_i^\top x - td_i) = -\sum_{i=1}^m \ln(a_i^\top (x/t) - d_i) - m \ln(t), \quad (1.13)$$

which is a m -logarithmically-homogeneous self-concordant barrier on the conic hull of \mathcal{Q} , denoted by $\text{cone}(\mathcal{Q})$ and defined by

$$\text{cone}(\mathcal{Q}) := \text{cl} \{(x, t) \in \mathbb{R}^{n+1} : x/t \in \mathcal{Q}, t > 0\} = \{x \in \mathbb{R}^n : a_i^\top x \geq td_i, i = 1, \dots, m, t \geq 0\}.$$

Then we can formulate the analytic center problem as

$$\min_{(x,t) \in \mathbb{R}^{n+1}} f(x, t) \quad \text{s. t.} \quad (x, t) \in \text{cone}(\mathcal{Q}), t = 1,$$

which is an instance of (P) with $f(u) := -\sum_{i=1}^m \ln(u_i)$, $\mathcal{K} := \mathbb{R}_+^m$, $h = \iota_{\mathcal{X}}$ where $\mathcal{X} = \{(x, t) \in \mathbb{R}^{n+1} : x \in \mathcal{Q}, t = 1\}$, \mathbf{A} defined by $(\mathbf{A}(x, t))_i := a_i^\top x - td_i$ for $i \in [m]$, and $\theta = m$.

The above formulation can be generalized to any nonempty convex compact set $\mathcal{Q} \subseteq \mathbb{R}^n$ equipped with a (standard strongly) ϑ -self-concordant barrier b on \mathcal{Q} . From [45, Proposition 5.1.4], there exists a constant $c \leq 20$ for which

$$f(x, t) := c^2(b(x/t) - 2\vartheta \ln t) \quad (1.14)$$

is a $(2c^2\vartheta)$ -logarithmically-homogeneous self-concordant barrier on $\text{cone}(\mathcal{Q})$. Therefore using (1.14) the analytic center problem can be formulated as an instance of (P) in a similar way as above.

5. *D-optimal design* [22]. Given m points $a^1, \dots, a^m \in \mathbb{R}^n$ whose affine hull is \mathbb{R}^n , the D -optimal design problem is:

$$\min h(x) := -\ln \det \left(\sum_{i=1}^m x_i a_i a_i^\top \right) \quad \text{s. t.} \quad x \in \Delta_m. \quad (1.15)$$

In the domain of statistics the D -optimal design problem corresponds to maximizing the determinant of the Fisher information matrix $\mathbb{E}(aa^\top)$, see [36], [1], as well as the exposition in [7]. And in computational geometry, D -optimal design arises as a Lagrangian dual problem of the minimum volume covering ellipsoid (MVCE) problem, which dates back at least 70 years to [34], see Todd [54] for a modern treatment. Indeed, (1.15) is useful in a variety of different application areas, for example, computational statistics [13] and data mining [37].

A recent extension of (1.15) is the design of a supervised learning pipeline for new datasets to maximize the Fisher information, see the recent paper by Yang et al. [56]. The optimization problem they consider is the following variant of (1.15):

$$\min h(x) := -\ln \det \left(\sum_{i=1}^m x_i a_i a_i^\top \right) \quad (1.16)$$

$$\text{s. t.} \quad \sum_{i=1}^n \bar{t}_i x_i \leq \tau \quad (1.17)$$

$$x_i \in [0, 1] \text{ for } i \in [m], \quad (1.18)$$

where the decision variable x_i models the decision to fit model i or not, \bar{t}_i is the estimated pipeline running time of pipeline i , τ is the runtime limit, and a_i is the vector of latent meta-features of model i for $i \in [m]$. The constraints (1.18) are a linear relaxation of the (computationally unattractive) desired combinatorial constraints $x_i \in \{0, 1\}$, $i \in [m]$. We refer the interested reader to [56] for further details and model variations. Here we see that (1.16)-(1.18) is an instance of (1.1) with $f(U) := -\ln \det(U)$, $\theta = n$, $Ax := \sum_{i=1}^m x_i a_i a_i^T$, and h is the indicator function of the feasible region of the constraints (1.17)-(1.18).

As mentioned earlier, the D -optimal design problem was one of the primary motivators for the research in this paper. Indeed, Khachiyan proved that his “barycentric coordinate descent” method for this problem – which it turns out is precisely the Frank-Wolfe method (with exact line search) – has a computational guarantee that is essentially $O(n^2/\varepsilon)$ iterations to produce an ε -approximate solution. What has been surprising about this result is that the D -optimal design problem violates the basic assumption underlying the premise for the analysis of the Frank-Wolfe method, namely L -smoothness. Khachiyan’s proof used original and rather clever arguments that do not easily carry over elsewhere, which has begged the question of whether or not any of Khachiyan’s arguments might underlie any general themes beyond D -optimal design, and if so what might be the mathematical structures driving any such themes? We will answer these questions in the affirmative in Section 2 by showing that the Frank-Wolfe method achieves essentially $O((\theta + R_h)^2/\varepsilon)$ iteration complexity when used to tackle any problem of the form (1.1), where R_h is the variation of h on its domain ($R_h := \max_{x,y \in \mathcal{X}} |h(x) - h(y)|$) and f is a θ -logarithmically homogeneous self-concordant barrier. When specialized to the D -optimal design problem, we have $R_h = 0$ (since h is an indicator function), and $\theta = n$, whereby we recover Khachiyan’s $O(n^2/\varepsilon)$ dependence on ε . In this respect, our results reveal certain intrinsic connections between θ -logarithmic homogeneity and the Frank-Wolfe method.

By way of concluding this discussion, we note that historically the theory of self-concordant functions was initially developed to present a general underlying theory for Newton’s method for barrier methods in convex optimization. However, the results in this paper indicate that a subclass of self-concordant functions, namely the class of θ -logarithmically homogeneous self-concordant barriers, are also tied to an underlying theory for the Frank-Wolfe method (and perhaps other first-order methods as well).

1.2 Contributions

We summarize our contributions as follows:

1. We propose a generalized Frank-Wolfe method for solving (P) with $f \in \mathcal{B}_\theta(\mathcal{K})$. We show that the Frank-Wolfe method requires $O((\delta_0 + \theta + R_h) \ln(\delta_0) + (\theta + R_h)^2/\varepsilon)$ iterations to produce an ε -approximate solution of (P) , namely $x \in \text{dom } F$ such that $F(x) - F^* \leq \varepsilon$, where δ_0 denotes the initial optimality gap and R_h denotes the variation of h on its domain. This iteration complexity bound depends on just three (natural) quantities associated with (P) : (i) the initial optimality gap δ_0 , (ii) the complexity parameter θ of f , and (iii) the variation of h on \mathcal{X} . When h is the sum of a convex quadratic function and an indicator function, our algorithm specializes to that in Dvurechensky et al. [20]. However, our iteration complexity bounds are quite different from that in [20] – in particular our bounds are affine invariant, more naturally interpretable, and are typically easy to estimate and can yield an *a priori* bound on the number of iterations needed to guarantee a desired optimality tolerance ε .

These issues are discussed in details in Remark 2.3.

2. Our analysis also yields $O((\delta_0 + \theta + R_h) \ln(\delta_0) + (\theta + R_h)^2/\varepsilon)$ iteration complexity to produce $x \in \text{dom } F$ whose Frank-Wolfe gap (defined in (2.3) below) is no larger than ε . Since the Frank-Wolfe gap is constructed at each iteration and is often used as the stopping criterion for the Frank-Wolfe method, our result provides a further constructive bound on the number of iterations required to detect a desired optimality tolerance.
3. We present a mirror descent method with adaptive step-size applied to the (Fenchel) dual problem (D) of (P) . The dual problem (D) shares a somewhat similar structure to (P) in that its objective function is non-smooth and non-Lipschitz. However, unlike (P) , the objective function has unbounded domain. Although these features make the direct analysis of mirror descent rather difficult, through the duality of mirror descent and the Frank-Wolfe method we provide a computational complexity bound for this mirror descent method via the Frank-Wolfe method applied to (P) . An application of our mirror descent method for (D) arises in the sub-problem in Bregman proximal-point methods.
4. We present computational experiments that point to the potential usefulness of our generalized Frank-Wolfe method on Poisson image de-blurring problems with TV regularization.

1.3 Outline and Notation

Outline. The paper is organized as follows. In Section 2 we present and analyze our generalized Frank-Wolfe method for (P) when $f \in \mathcal{B}_\theta(\mathcal{K})$, using an adaptive step-size strategy that is a natural extension of the strategy developed in [20]. In Section 3 we study the (Fenchel) dual (D) of (P) and derive and analyze a dual mirror descent method for solving (D) based on the generalized Frank-Wolfe method for solving (P) . In Section 4 we present computational experiments that point to the potential usefulness of our generalized Frank-Wolfe method on Poisson image de-blurring problems with TV regularization, and we also present computational experiments on the PET problem.

Notation. Let $\mathbb{R}_+^n := \{x \in \mathbb{R}^n : x \geq 0\}$ and $\mathbb{R}_{++}^n := \{x \in \mathbb{R}^n : x > 0\}$. The set of integers $\{1, \dots, n\}$ is denoted by $[n]$. The domain of a convex function f is denoted by $\text{dom } f := \{x \in \mathbb{R}^n : f(x) < \infty\}$. We use $H(x)$ to denote the Hessian of the function f . The interior and relative interior of a set \mathcal{S} are denoted by $\text{int } \mathcal{S}$ and $\text{ri } \mathcal{S}$, respectively. We use e to denote the vector with entries all equal to 1, $\text{Diag}(x)$ to denote the diagonal matrix whose diagonal entries correspond to those of x , and Δ_n to denote the standard $(n - 1)$ -dimensional simplex in \mathbb{R}^n , namely $\Delta_n := \{x \in \mathbb{R}^n : \sum_{i=1}^n x_i = 1, x \geq 0\}$. We use \mathbb{S}_+^n (\mathbb{S}_{++}^n) to denote the set of $n \times n$ symmetric positive semidefinite (positive definite) matrices, and write $B \in \mathbb{S}_+^n$ as $B \succeq 0$ and $B \in \mathbb{S}_{++}^n$ as $B \succ 0$. The p -norm of a vector $x \in \mathbb{R}^n$ is denoted and defined by $\|x\|_p = (\sum_{i=1}^n |x_i|^p)^{1/p}$.

2 A generalized Frank-Wolfe method for (P) when f is a θ -logarithmically-homogeneous self-concordant barrier

In this section we present and analyze a generalized Frank-Wolfe method for the composite optimization problem (P) in the case when $f \in \mathcal{B}_\theta(\mathcal{K})$, using an adaptive step-size strategy that is a natural extension of the strategy developed in Dvurechensky et al. [20]. We assume throughout that $\mathcal{X} := \text{dom } h$ is a convex and compact set, and that $A(\mathcal{X}) \cap \text{dom } f \neq \emptyset$. These two assumptions

together with the differentiability of f on $\text{int}\mathcal{K}$ ensure that (P) has at least one optimal solution which we denote by x^* and hence $F^* = F(x^*)$.

Let us introduce some important notation. For any $u \in \text{int}\mathcal{K}$, the Hessian $H(u)$ of f is used to define the local (Hilbert) norm $\|\cdot\|_u$ defined by:

$$\|w\|_u := \sqrt{\langle w, H(u)w \rangle} \quad \text{for all } w \in \mathbb{R}^m .$$

The Dikin ball $\mathcal{D}(u, 1)$ at $u \in \text{int}\mathcal{K}$ is defined by

$$\mathcal{D}(u, 1) := \{v \in \mathcal{K} : \|v - u\|_u < 1\} ,$$

and it can be shown that $\mathcal{D}(u, 1) \subseteq \text{int}\mathcal{K}$, see Nesterov and Nemirovski [45, Theorem 2.1.1]. The self-concordant function f is well-behaved inside the Dikin ball as we will review shortly. Define the univariate function ω and its Fenchel conjugate ω^* as follows:

$$\omega(a) := -a - \ln(1 - a) \quad \forall a < 1 , \quad \text{and} \quad \omega^*(a) := a - \ln(1 + a) \quad \forall a > -1 . \quad (2.1)$$

It turns out that f can be nicely upper-bounded inside the Dikin ball, namely:

$$f(v) \leq f(u) + \langle \nabla f(u), v - u \rangle + \omega(\|v - u\|_u) \quad \forall u \in \text{int}\mathcal{K}, \forall v \in \mathcal{D}(u, 1) , \quad (2.2)$$

see Nesterov [42, Theorem 4.1.8].

We now develop our generalized Frank-Wolfe method for the composite optimization problem (P) under the condition that $f \in \mathcal{B}_\theta(\mathcal{K})$, and whose formal rules are presented in Algorithm 1. Given the current iterate x^k , the method first computes the gradient $\nabla f(\mathbf{A}x^k)$ and then solves for a minimizer v^k of the generalized Frank-Wolfe sub-problem given by

$$v^k \in \arg \min_{x \in \mathbb{R}^n} \langle \nabla f(\mathbf{A}x^k), \mathbf{A}x \rangle + h(x) .$$

The next iterate is then determined as a convex combination of x^k and v^k : $x^{k+1} \leftarrow x^k + \alpha_k(v^k - x^k)$ for some $\alpha_k \in [0, 1]$, where α_k is the step-size at iteration k . For L -smooth functions f , the step-size can be chosen by a variety of strategies, including simple rules such as $\alpha_k = \frac{2}{(k+2)}$, exact line-search to minimize $f(x^k + \alpha(v^k - x^k))$ on $\alpha \in [0, 1]$, or an adaptive strategy based on curvature information, etc. Here we present an adaptive strategy based on the upper bound model (2.2), which can also be viewed as an extension of the adaptive strategy used in [20] and which itself is an extension of the adaptive strategy in Demyanov and Rubinov [14]. Define the Frank-Wolfe gap (“FW-gap”) G_k by

$$G_k := \langle \nabla f(\mathbf{A}x^k), \mathbf{A}(x^k - v^k) \rangle + h(x^k) - h(v^k) , \quad (2.3)$$

and the optimality gap $\delta_k := F(x^k) - F^*$. Note that $G_k \geq 0$ and in fact by the convexity of f it holds that

$$\begin{aligned} \delta_k &= (f(\mathbf{A}x^k) - f(\mathbf{A}x^*)) + (h(x^k) - h(x^*)) \\ &\leq \langle \nabla f(\mathbf{A}x^k), \mathbf{A}(x^k - x^*) \rangle + (h(x^k) - h(x^*)) \\ &\leq \langle \nabla f(\mathbf{A}x^k), \mathbf{A}(x^k - v^k) \rangle + (h(x^k) - h(v^k)) = G_k , \end{aligned} \quad (2.4)$$

hence G_k is indeed an upper bound on the optimality gap δ_k . Denoting $D_k := \|\mathbf{A}(v^k - x^k)\|_{\mathbf{A}x^k}$, we then have that for any $\alpha \geq 0$,

$$f(\mathbf{A}x^k + \alpha \mathbf{A}(v^k - x^k)) \leq f(\mathbf{A}x^k) - \alpha \langle \nabla f(\mathbf{A}x^k), \mathbf{A}(x^k - v^k) \rangle + \omega(\alpha D_k) . \quad (2.5)$$

Algorithm 1 (generalized) Frank-Wolfe Method for composite optimization involving $f \in \mathcal{B}_\theta(\mathcal{K})$ with adaptive step-size

Input: Starting point $x^0 \in \mathcal{X}$

At iteration $k \in \{0, 1, \dots\}$:

1. Compute $\nabla f(\mathbf{A}x^k)$ and $v^k \in \arg \min_{x \in \mathbb{R}^n} \langle \nabla f(\mathbf{A}x^k), \mathbf{A}x \rangle + h(x)$
2. Compute $G_k := \langle \nabla f(\mathbf{A}x^k), \mathbf{A}(x^k - v^k) \rangle + h(x^k) - h(v^k)$ and $D_k := \|\mathbf{A}(v^k - x^k)\|_{\mathbf{A}x^k}$, and compute the step-size:

$$\alpha_k := \min \left\{ \frac{G_k}{D_k(G_k + D_k)}, 1 \right\} \quad (2.8)$$

3. Update $x^{k+1} := x^k + \alpha_k(v^k - x^k)$
-

(Note that if $\alpha < 1/D_k$, then (2.5) follows from (2.2); otherwise, by the definition of ω in (2.1), we have $\omega(\alpha D_k) = +\infty$ and (2.5) still holds.) Also, by the convexity of h , we have

$$h(x^k + \alpha(v^k - x^k)) \leq (1 - \alpha)h(x^k) + \alpha h(v^k) = h(x^k) - \alpha(h(x^k) - h(v^k)). \quad (2.6)$$

Adding (2.5) and (2.6) together, we obtain

$$F(x^k + \alpha(v^k - x^k)) \leq F(x^k) - \alpha G_k + \omega(\alpha D_k), \quad \forall \alpha \geq 0, \quad (2.7)$$

and optimizing the right-hand-side over $\alpha \in [0, 1]$ yields the step-size:

$$\alpha_k := \min \left\{ \frac{G_k}{D_k(G_k + D_k)}, 1 \right\}.$$

Notice that this step-size specializes precisely to the adaptive step-size developed in [20] in the case when the function h is the indicator function $\iota_{\mathcal{X}}$ of a compact convex set \mathcal{X} , since in this case $h(x^k) = h(v^k) = 0$ for all k and hence G_k turns out to be the standard ‘‘gap function’’ $\langle \nabla f(\mathbf{A}x^k), \mathbf{A}(x^k - v^k) \rangle$ as used in [20].

Before presenting our analysis of Algorithm 1 we make two remarks. First, notice that the complexity parameter θ of f is not needed to run Algorithm 1, but it will play a central role in analyzing the iteration complexity of the algorithm. In a sense, Algorithm 1 automatically adapts to the value of θ . Second, for most applications — including all of the applications discussed in Section 1.1 — the computational cost of computing D_k (in Step 2) is of the same order as computing G_k , and grows linearly in the ambient dimension of the variable x . (Note that here our discussion focuses on the dependence of the computational cost on the dimension of x only.) To see this, let us fix any $v, x \in \text{dom } F$. For the first application in Section 1.1 (where N denotes the dimension), both $u := \mathbf{A}x$ and $w := \mathbf{A}v$ can be computed in $O(qN)$ time, due to the fact that the matrix representation of \mathbf{A} has only $O(qN)$ nonzeros. Since $D_k^2 = \sum_{l=1}^N (u_l - w_l)^2 / u_l^2$, it can be computed in $O(qN)$ time. Using similar reasoning, we easily see that for the second, third and fourth applications (where n denotes the dimension), D_k can be computed in $O(mn)$ time. For the last application (namely the D -optimal design problem), where the dimension is denoted by m , D_k can be computed in $O(mn^2 + n^3)$ time for $k = 0$ and $O(n^2)$ time for $k \geq 1$; for details see Appendix A.

2.1 Computational guarantees for Algorithm 1

We now present our computational guarantees for Algorithm 1. These guarantees depend on only three natural quantities associated with (P): (i) the initial optimality gap δ_0 , (ii) the complexity parameter θ of f , and (iii) the *variation* of h on \mathcal{X} , which is defined as:

$$R_h := \max_{x,y \in \mathcal{X}} |h(x) - h(y)| . \quad (2.9)$$

Regarding the variation R_h we mention two cases in particular:

1. when $h = \iota_{\mathcal{X}}$, we have $R_h = 0$, and
2. when h is L -Lipschitz on \mathcal{X} with respect to some norm $\|\cdot\|$, we have $R_h \leq LD_{\mathcal{X},\|\cdot\|}$, where

$$D_{\mathcal{X},\|\cdot\|} := \sup_{x,x' \in \mathcal{X}} \|x - x'\| < +\infty \quad (2.10)$$

is the diameter of \mathcal{X} under $\|\cdot\|$. And in particular if $h = \|\cdot\|$, then $R_h \leq D_{\mathcal{X},\|\cdot\|}$.

Theorem 2.1. *Suppose that $f \in \mathcal{B}_\theta(\mathcal{K})$ and that Algorithm 1 is initiated at $x^0 \in \mathcal{X}$. Let δ_0 denote the initial optimality gap.*

1. *At iteration k of Algorithm 1 the following hold:*

(a) *If $G_k > \theta + R_h$, then the optimality gap decreases at least linearly at the iteration:*

$$\delta_{k+1} \leq \left(1 - \frac{1}{5.3(\delta_0 + \theta + R_h)}\right) \delta_k , \quad (2.11)$$

(b) *If $G_k \leq \theta + R_h$, then the inverse optimality gap increases by at least a constant at the iteration:*

$$\frac{1}{\delta_{k+1}} \geq \frac{1}{\delta_k} + \frac{1}{12(\theta + R_h)^2} , \quad \text{and} \quad (2.12)$$

(c) *The number of iterations K_{Lin} where $G_k > \theta + R_h$ occurs is bounded from above as follows: $K_{\text{Lin}} \leq \lceil 5.3(\delta_0 + \theta + R_h) \ln(10.6\delta_0) \rceil$.*

2. *Let K_ε denote the number of iterations required by Algorithm 1 to obtain $\delta_k \leq \varepsilon$. Then:*

$$K_\varepsilon \leq \lceil 5.3(\delta_0 + \theta + R_h) \ln(10.6\delta_0) \rceil + \left\lceil 12(\theta + R_h)^2 \max \left\{ \frac{1}{\varepsilon} - \frac{1}{\delta_0}, 0 \right\} \right\rceil . \quad (2.13)$$

3. *Let FWGAP_ε denote the number of iterations required by Algorithm 1 to obtain $G_k \leq \varepsilon$. Then:*

$$\text{FWGAP}_\varepsilon \leq \lceil 5.3(\delta_0 + \theta + R_h) \ln(10.6\delta_0) \rceil + \left\lceil \frac{24(\theta + R_h)^2}{\varepsilon} \right\rceil . \quad (2.14)$$

□

Before we present our proof, let us make a few remarks about the results in Theorem 2.1.

Remark 2.1 (Discussion of complexity results). Theorem 2.1 indicates that if $f \in \mathcal{B}_\theta(\mathcal{K})$, then the iteration complexity to obtain an ε -optimal solution using Algorithm 1 is

$$O((\delta_0 + \theta + R_h) \ln(\delta_0) + (\theta + R_h)^2/\varepsilon) , \quad (2.15)$$

which only depends on θ , δ_0 , R_h , and the desired optimality gap ε . It is instructive to compare this result with the one obtained by assuming the L -smoothness of f (with respect to some given norm $\|\cdot\|$) and its consequent upper-bound model of f (see e.g., [24]):

$$f(y) \leq f(x) + \langle \nabla f(x), y - x \rangle + \frac{L}{2} \|y - x\|^2 \quad \forall x, y \in \mathcal{X} . \quad (2.16)$$

Under the condition (2.16), the complexity of the generalized Frank-Wolfe method to obtain an ε -optimal solution is

$$O(LD_{\mathbf{A}(\mathcal{X}), \|\cdot\|}^2/\varepsilon) , \quad (2.17)$$

where $D_{\mathbf{A}(\mathcal{X}), \|\cdot\|}$ denotes the diameter of $\mathbf{A}(\mathcal{X})$ under $\|\cdot\|$, and is defined in (2.10). Although both (2.15) and (2.17) have the $O(1/\varepsilon)$ dependence on ε , we observe at least two main differences. First, unlike (2.17), (2.15) has no dependence on any given norm structure imposed on \mathcal{X} , which is in concert with the intrinsic properties of self-concordance. Second, (2.15) explicitly depends on the initial sub-optimality gap δ_0 and the variation of h on \mathcal{X} , which is not the case for (2.17) (although [24] shows how a variation of (2.17) can be developed that accounts for bounds on δ_0).

Remark 2.2 (Comparison with Khachiyan [35] for D -optimal design). Let us specialize the complexity bound in (2.15) to the D -optimal design problem in (1.15), and compare it with the complexity bound in Khachiyan [35]. Note that for the problem in (1.15) we have $\theta = n$, i.e., the dimension of the ambient space of the points a_1, \dots, a_m . In addition, if we choose the starting point $p^0 = (1/m)e$, where $e := (1, \dots, 1) \in \mathbb{R}^m$, then

$$\delta_0 \leq n \ln(m/n) \quad (2.18)$$

(which we show in Appendix D), and then based on (2.18) and $R_h = 0$, the iteration complexity bound in (2.15) becomes

$$O(n \ln(m/n)(\ln n + \ln \ln(m/n)) + n^2/\varepsilon) . \quad (2.19)$$

Using the same starting point, Khachiyan's Frank-Wolfe method [35] uses exact line-search (based on a clever observation from the Inverse Matrix Update formula [29]), and attains the complexity bound

$$O(n(\ln n + \ln \ln(m/n)) + n^2/\varepsilon) . \quad (2.20)$$

Observe that (2.19) has the exact same dependence on ε as (2.20), namely $O(n^2/\varepsilon)$, but its first term is inferior to (2.20) by the factor $O(\ln(m/n))$. The improvement in the first term of Khachiyan's bound over the bound in (2.19) is due to his improved estimate of the linear convergence rate in the case $G_k > \theta$ in Algorithm 1, which arises from exploiting an exact line-search on f . This is in contrast with our method, which only does an exact line-search on the upper bound model of f in (2.7). A detailed analysis of this last point is given in Remark 2.4 at the end of this section.

Remark 2.3 (Comparison with Dvurechensky et al. [20]). The recent work of Dvurechensky et al. [20] considers the Frank-Wolfe method for the problem $\min_{x \in \mathcal{X}} \bar{F}(x)$, where \mathcal{X} is nonempty, convex and compact and \bar{F} is a non-degenerate (strongly) M -self-concordant function for some

$M > 0$. This means \bar{F} is convex and three-times differentiable on $\text{dom } \bar{F}$, $\nabla^2 \bar{F}(x) \succ 0$ for all $x \in \text{dom } \bar{F}$ (this is the definition that \bar{F} is non-degenerate), and

$$|D^3 \bar{F}(x)[z, z, z]| \leq 2M^{-1/2} (\langle \nabla^2 \bar{F}(x) z, z \rangle)^{3/2} \quad \forall x \in \text{dom } \bar{F}, \quad \forall z \in \mathbb{R}^n. \quad (2.21)$$

For convenience of comparison, henceforth let $M = 1$. When h is the sum of a convex quadratic function and an indicator function, i.e., $h(x) = \frac{1}{2} \langle x, Qx \rangle + \xi^\top x + \iota_{\mathcal{X}}(x)$ for some $Q \succeq 0$ and $\xi \in \mathbb{R}^n$, then our method coincides with that in Dvurechensky et al. [20] with $\bar{F}(x) = f(\mathbf{A}x) + \frac{1}{2} \langle x, Qx \rangle + \xi^\top x$. The complexity bound developed in [20] for computing an ε -optimal solution is:

$$O \left(\sqrt{L(x^0)} D_{\mathcal{X}, \|\cdot\|_2} \ln \left(\frac{\delta_0}{\sqrt{L(x^0)} D_{\mathcal{X}, \|\cdot\|_2}} \right) + \frac{L(x^0) D_{\mathcal{X}, \|\cdot\|_2}^2}{\varepsilon} \right), \quad (2.22)$$

where

$$\mathcal{S}(x^0) := \{x \in \text{dom } \bar{F} \cap \mathcal{X} : \bar{f}(x) \leq \bar{f}(x^0)\}, \text{ and} \quad (2.23)$$

$$L(x^0) := \max_{x \in \mathcal{S}(x^0)} \lambda_{\max}(\nabla^2 \bar{F}(x)) < +\infty, \quad (2.24)$$

where in (2.24) $\lambda_{\max}(\nabla^2 \bar{F}(x))$ denotes the largest eigenvalue of $\nabla^2 \bar{F}(x)$, and in [20] it is further assumed that \bar{F} is non-degenerate which necessarily presumes that $\text{rank}(\mathbf{A}) = n$.

It is instructive to compare the two iteration complexity bounds (2.15) and (2.22), and we note the following advantageous features of our bound (2.15) as follows:

- *Affine invariance.* Affine invariance is an important structural property certain algorithms; for instance Newton's method is affine invariant whereas the steepest descent method is not. It is well known that the Frank-Wolfe method is affine invariant. Current state-of-the-art complexity analysis of the Frank-Wolfe method for L -smooth functions yields an appropriate affine-invariant complexity bound by utilizing the so-called *curvature constant* $C_{\bar{F}}$ of Clarkson [11] defined by

$$C_{\bar{F}} := \max_{x, y \in \mathcal{X}, \alpha \in [0, 1]} \frac{2}{\alpha^2} (\bar{F}(x + \alpha(y - x)) - \bar{F}(x) - \alpha \langle \nabla \bar{F}(x), y - x \rangle), \quad (2.25)$$

which is a finite affine-invariant quantity [11] when \bar{F} is L -smooth. Of course $C_{\bar{F}}$ is not typically finite of when \bar{F} is self-concordant. The complexity bound in (2.22) depends on measures that are tied to the Euclidean inner product and norm, namely $D_{\mathcal{X}, \|\cdot\|_2}$ and $L(x^0)$, and are not affine-invariant, even though the Euclidean norm plays no part in the algorithm. (Under an invertible affine transformation of the variables these measures will change but the performance of the Frank-Wolfe method will not change.) In contrast, all the quantities in our complexity bounds, namely δ_0 , θ and R_h are affine-invariant, therefore the complexity bounds in (2.15) are affine-invariant.

- *Interpretability.* Apart from ε , our complexity result only depends on δ_0 , θ , and R_h , all of which admit natural behavioral interpretations. Specifically, δ_0 measures the initial sub-optimality gap, θ is the ‘‘complexity parameter’’ of the barrier f (in the lexicon of Renegar [48]), and R_h measures the variation of h over the set \mathcal{X} .

- *Ease of parameter estimation.* Note that all of the three parameters δ_0 , θ , and R_h in our complexity bound (2.15) are either easy to know or easy to appropriately bound. Given a logarithmically-homogeneous self-concordant barrier f , its complexity value θ is typically known *a priori* or can be easily determined using (P7) of Lemma 2.1. Since $\delta_0 \leq G_0$, a natural upper bound on δ_0 is the initial FW-gap G_0 , which is computed in the second step at iteration $k = 0$ of Algorithm 1. Regarding R_h , note that $R_h = 0$ when h is the indicator function of a convex set \mathcal{X} , namely $h = \iota_{\mathcal{X}}$. In the case when $h(x) := \xi^\top x + \iota_{\mathcal{X}}(x)$ then R_h can be computed exactly as the difference of two linear optimization optimal values on \mathcal{X} or can be upper bounded using

$$R_h = \max_{x, x' \in \mathcal{X}} |\xi^\top (x - x')| \leq \|\xi\|_* D_{\mathcal{X}, \|\cdot\|} ,$$

in the case when the norm $\|\cdot\|$ can possibly be chosen to yield easily computable values of $D_{\mathcal{X}, \|\cdot\|}$. Apart from this case, there also exist many other cases where the simple nature of h and \mathcal{X} yield easily computable upper-bounds on R_h .

2.2 Proof of Theorem 2.1

We first state some facts about θ -logarithmically homogeneous self-concordant barrier functions.

Lemma 2.1 (see Nesterov and Nemirovskii [45, Corollary 2.3.1, Proposition 2.3.4, and Corollary 2.3.3]). *If $f \in \mathcal{B}_\theta(\mathcal{K})$, then for any $u \in \text{int } \mathcal{K}$, we have*

$$(P4) \quad |\langle \nabla f(u), w \rangle| \leq \sqrt{\theta} \|w\|_u \quad \forall w \in \mathbb{R}^m ,$$

$$(P5) \quad \|v\|_u \leq -\langle \nabla f(u), v \rangle \quad \forall v \in \mathcal{K} ,$$

$$(P6) \quad \langle \nabla f(u), w \rangle = -\langle H(u)u, w \rangle \quad \forall w \in \mathbb{R}^m ,$$

$$(P7) \quad \langle \nabla f(u), u \rangle = -\theta, \text{ and}$$

$$(P8) \quad \theta \geq 1. \quad \square$$

We also introduce some properties of the function ω^* (cf. (2.1)) and present an “old” property of the logarithm function.

Proposition 2.1. *The function ω^* is strictly increasing on $[0, +\infty)$, and*

$$\omega^*(s) \geq s^2/3 \quad \forall s \in [0, 1/2] , \text{ and} \quad (2.26)$$

$$\omega^*(s) \geq s/5.3 \quad \forall s \geq 1/2 . \quad (2.27)$$

Proof. See Appendix B. □

Proposition 2.2.

$$\ln(1 + s) \geq s - \frac{s^2}{2(1 - |s|)} \quad \forall s \in (-1, 1) . \quad (2.28)$$

Proof. See Appendix C. □

We have the following inequality concerning values of D_k and G_k computed in Step 2 of Algorithm 1. For convenience, in the following, define

$$\tilde{G}_k := \langle \nabla f(\mathbf{A}x^k), \mathbf{A}(x^k - v^k) \rangle \quad \text{and} \quad \beta_k := h(x^k) - h(v^k),$$

so that $G_k = \tilde{G}_k + \beta_k$. Also, by the definition of R_h , we know that $|\beta_k| \leq R_h$.

Proposition 2.3. *For all $k \geq 0$ it holds that*

$$D_k \leq G_k + \theta + R_h. \quad (2.29)$$

Proof. We have:

$$D_k^2 = \langle H(\mathbf{A}x^k)\mathbf{A}v^k, \mathbf{A}v^k \rangle - 2\langle H(\mathbf{A}x^k)\mathbf{A}x^k, \mathbf{A}v^k \rangle + \langle H(\mathbf{A}x^k)\mathbf{A}x^k, \mathbf{A}x^k \rangle. \quad (2.30)$$

By (P5) we see that

$$\langle H(\mathbf{A}x^k)\mathbf{A}v^k, \mathbf{A}v^k \rangle \leq \langle \nabla f(\mathbf{A}x^k), \mathbf{A}v^k \rangle^2 = (-\tilde{G}_k + \langle \nabla f(\mathbf{A}x^k), \mathbf{A}x^k \rangle)^2 = (\tilde{G}_k + \theta)^2, \quad (2.31)$$

where the last equality above uses (P7). In addition, from (P6) and (P7) we have

$$\begin{aligned} -2\langle H(\mathbf{A}x^k)\mathbf{A}x^k, \mathbf{A}v^k \rangle + \langle H(\mathbf{A}x^k)\mathbf{A}x^k, \mathbf{A}x^k \rangle &= 2\langle \nabla f(\mathbf{A}x^k), \mathbf{A}v^k \rangle - \langle \nabla f(\mathbf{A}x^k), \mathbf{A}x^k \rangle \\ &= -2\tilde{G}_k + \langle \nabla f(\mathbf{A}x^k), \mathbf{A}x^k \rangle \\ &= -2\tilde{G}_k - \theta. \end{aligned} \quad (2.32)$$

Combining (2.30), (2.31) and (2.32), we have

$$\begin{aligned} D_k^2 &\leq (\tilde{G}_k + \theta)^2 - 2\tilde{G}_k - \theta \\ &= (G_k - \beta_k + \theta)^2 + 2\beta_k - (2G_k + \theta) \\ &\leq (G_k + \theta)^2 + \beta_k^2 - 2\beta_k(G_k + \theta - 1) \end{aligned} \quad (2.33)$$

$$\leq (G_k + \theta)^2 + R_h^2 + 2R_h(G_k + \theta) \quad (2.34)$$

$$= (G_k + \theta + R_h)^2, \quad (2.35)$$

where in (2.33) we use $2G_k + \theta \geq 0$ and in (2.34) we use $|\beta_k| \leq R_h$ and $G_k + \theta \geq \theta \geq 1$. \square

The basic iteration improvement inequality for the Frank-Wolfe was presented in (2.5), and the step-size in Algorithm 1 is given by (2.8). In the case when $\alpha_k < 1$, it follows from substituting $\alpha_k = \frac{G_k}{D_k(G_k + D_k)}$ from (2.8) into (2.7) that the iteration improvement bound is

$$F(x^{k+1}) \leq F(x^k) - \omega^* \left(\frac{G_k}{D_k} \right). \quad (2.36)$$

Using the notation $\Delta_k := F(x^k) - F(x^{k+1}) = \delta_k - \delta_{k+1}$, we can write this improvement as:

$$\Delta_k = \delta_k - \delta_{k+1} = F(x^k) - F(x^{k+1}) \geq \omega^* \left(\frac{G_k}{D_k} \right) \geq 0 \quad \text{when } \alpha_k < 1. \quad (2.37)$$

Let us now prove (2.11) and part 1c of Theorem 2.1. Since $G_k > \theta + R_h$, by (P4) in Lemma 2.1 we have

$$\begin{aligned} D_k &= \|Av^k - Ax^k\|_{Ax^k} \geq |\langle \nabla f(Ax^k), Av^k - Ax^k \rangle| / \sqrt{\theta} \\ &\geq \tilde{G}_k / \sqrt{\theta} = (G_k - \beta_k) / \sqrt{\theta} \geq (G_k - R_h) / \sqrt{\theta} > \sqrt{\theta} \geq 1. \end{aligned} \quad (2.38)$$

As a result, $\alpha_k = \frac{G_k}{D_k(G_k + D_k)} < 1$. Consequently, by (2.37) we have

$$F(x^{k+1}) \leq F(x^k) - \omega^* \left(\frac{G_k}{D_k} \right) \leq F(x^k) - \omega^* \left(\frac{G_k}{G_k + \theta + R_h} \right), \quad (2.39)$$

where the last inequality uses (2.29) and the monotonicity of ω^* . Now notice from the condition $G_k > \theta + R_h$ that $G_k / (G_k + \theta + R_h) > 1/2$, whereby invoking (2.27) we have

$$\Delta_k = \delta_k - \delta_{k+1} = F(x^k) - F(x^{k+1}) \geq \omega^* \left(\frac{G_k}{G_k + \theta + R_h} \right) \geq \frac{G_k}{5.3(G_k + \theta + R_h)}. \quad (2.40)$$

In addition we have

$$\frac{G_k}{5.3(G_k + \theta + R_h)} \geq \frac{\delta_k}{5.3(\delta_k + \theta + R_h)} \geq \frac{\delta_k}{5.3(\delta_0 + \theta + R_h)}, \quad (2.41)$$

where the first inequality uses the strict monotonicity of the function $c \mapsto c / (c + \theta + R_h)$ on $[0, +\infty)$, and the second inequality uses the monotonicity of the sequence $\{\delta_k\}_{k \geq 0}$ (see (2.37)). Combining (2.40) and (2.41), we obtain

$$\delta_{k+1} \leq \left(1 - \frac{1}{5.3(\delta_0 + \theta + R_h)} \right) \delta_k, \quad (2.42)$$

which proves (2.11). Furthermore, $G_k > \theta + R_h$ implies that

$$\delta_k \geq \Delta_k \geq \frac{G_k}{5.3(G_k + \theta + R_h)} > \frac{G_k}{5.3(G_k + G_k)} = \frac{1}{10.6}. \quad (2.43)$$

Now let K_{Lin} denote the number of iterations of Algorithm 1 where $G_k > \theta + R_h$ occurs. By (2.42) and (2.43) it follows that

$$\frac{1}{10.6} < \delta_0 \left(1 - \frac{1}{5.3(\delta_0 + \theta + R_h)} \right)^{K_{\text{Lin}} - 1}, \quad (2.44)$$

which then implies that $K_{\text{Lin}} \leq \lceil 5.3(\delta_0 + \theta + R_h) \ln(10.6\delta_0) \rceil$ and thus proving part 1c of Theorem 2.1.

Let us now prove (2.12) of Theorem 2.1. Towards doing so, we will establish:

$$G_k \leq \theta + R_h \implies \Delta_k \geq \frac{G_k^2}{12(\theta + R_h)^2}. \quad (2.45)$$

We first consider the case where $\alpha_k = 1$, whereby $D_k(G_k + D_k) \leq G_k$, which implies that $D_k < 1$, and also can be rearranged to yield

$$G_k \geq \frac{D_k^2}{1 - D_k}. \quad (2.46)$$

In addition, by (2.7) we obtain

$$\Delta_k = F(x^k) - F(x^{k+1}) \geq G_k - \omega(D_k) = G_k + D_k + \ln(1 - D_k) . \quad (2.47)$$

By (2.47), Proposition 2.2, and (2.46), we have

$$\Delta_k \geq G_k - \frac{D_k^2}{2(1 - D_k)} \geq \frac{G_k}{2} , \quad (2.48)$$

which then implies that

$$\Delta_k \geq \frac{G_k}{2} \geq \frac{G_k^2}{2(\theta + R_h)} \geq \frac{G_k^2}{2(\theta + R_h)^2} \geq \frac{G_k^2}{12(\theta + R_h)^2} , \quad (2.49)$$

where the second inequality used $G_k \leq \theta + R_h$ and the third inequality used $\theta + R_h \geq 1$. This establishes (2.45) for the case when $\alpha_k = 1$.

We next consider the case where $\alpha_k < 1$, whereby $\alpha_k = \frac{G_k}{D_k(G_k + D_k)}$, and then by (2.37) we have

$$\begin{aligned} \Delta_k &= F(x^k) - F(x^{k+1}) \\ &\geq \omega^* \left(\frac{G_k}{D_k} \right) \geq \omega^* \left(\frac{G_k}{G_k + \theta + R_h} \right) \geq \frac{G_k^2}{3(G_k + \theta + R_h)^2} \geq \frac{G_k^2}{12(\theta + R_h)^2} , \end{aligned} \quad (2.50)$$

where the second inequality uses (2.29) and the monotonicity of ω^* , the third inequality uses (2.26) in conjunction with $G_k/(G_k + \theta + R_h) \leq 1/2$, and the fourth inequality uses $G_k \leq \theta + R_h$. This establishes (2.45) for the case when $\alpha_k < 1$, completing the proof of (2.45). It thus follows for $G_k \leq \theta + R_h$ that

$$\delta_k - \delta_{k+1} = \Delta_k \geq \frac{G_k^2}{12(\theta + R_h)^2} \geq \frac{\delta_k \delta_{k+1}}{12(\theta + R_h)^2} ,$$

where the last inequality follows from $\delta_{k+1} \leq \delta_k \leq G_k$, and dividing both sides by $\delta_k \delta_{k+1}$ and rearranging yields the inequality (2.12).

We next prove (2.13) and (2.14). If $\delta_0 \leq \varepsilon$ the result follows trivially; thus we assume that $\delta_0 > \varepsilon$. Let \bar{K} denote the expression on the right-side of (2.13), and suppose Algorithm 1 has been run for \bar{K} iterations. Let $N := \lceil 12(\theta + R_h)^2 (1/\varepsilon - 1/\delta_0) \rceil$, whereby it follows from part 1c of Theorem 2.1 that among the first \bar{K} iterations, the number of iterations where $G_k \leq \theta + R_h$ is at least N . Thus from (2.12) it follows that

$$\frac{1}{\delta_{K\varepsilon}} \geq \frac{1}{\delta_0} + \frac{N}{12(\theta + R_h)^2} \geq \frac{1}{\delta_0} + \left(\frac{1}{\varepsilon} - \frac{1}{\delta_0} \right) = \frac{1}{\varepsilon} ,$$

and rearranging yields (2.13).

Let $k_0 < k_1 < k_2 < \dots$ denote indices where $G_k \leq \theta + R_h$. From (2.45), (2.4), and the monotonicity of the sequence $\{\delta_k\}_{k \geq 0}$, it follows for all $j \geq 0$ that

$$\delta_{k_{j+1}} \leq \delta_{k_j+1} \leq \delta_{k_j} - \frac{G_{k_j}^2}{12(\theta + R_h)^2} \quad \text{and} \quad G_{k_j} \geq \delta_{k_j} .$$

Let $d_j := \delta_{k_j}$ and $g_j := G_{k_j}$ for all $j \geq 0$, then the nonnegative sequences $\{d_j\}_{j \geq 0}$ and $\{g_j\}_{j \geq 0}$ satisfy for all $j \geq 0$:

$$d_{j+1} \leq d_j - \frac{g_j^2}{12(\theta + R_h)^2} \quad \text{and} \quad g_j \geq d_j .$$

Thus $\{d_j\}_{j \geq 0}$ and $\{g_j\}_{j \geq 0}$ satisfy the hypotheses of the following elementary sequence proposition using $M = 12(\theta + R_h)^2$. (This proposition is a slight extension of the standard sequence property for Frank-Wolfe type sequences, and we provide a proof in Appendix E.)

Proposition 2.4. *Suppose the two nonnegative sequences $\{d_j\}_{j \geq 0}$ and $\{g_j\}_{j \geq 0}$ satisfy for all $j \geq 0$:*

- $d_{j+1} \leq d_j - g_j^2/M$ for some $M > 0$, and
- $g_j \geq d_j$.

Then for all $j \geq 0$ the following holds:

$$d_j \leq \frac{M}{j + \frac{M}{d_0}} < \frac{M}{j}, \quad (2.51)$$

and

$$\min\{g_0, \dots, g_j\} < \frac{2M}{j}. \quad (2.52)$$

□

Let FWGAP_ε be as given in part 3 of Theorem 2.1, and let \tilde{K} denote the expression on the right-side of (2.14). Suppose Algorithm 1 has been run for \tilde{K} iterations. Let $\tilde{N} := \lceil 24(\theta + R_h)^2/\varepsilon \rceil$, whereby it follows from part 1c of Theorem 2.1 that among the first \tilde{K} iterations, the number of iterations where $G_k \leq \theta + R_h$ is at least \tilde{N} . Then, it follows that

$$\min\{G_0, \dots, G_{\tilde{K}}\} \leq \min\{G_{k_0}, \dots, G_{k_{\tilde{N}}}\} = \min\{g_0, \dots, g_{\tilde{N}}\} < \frac{2M}{\tilde{N}} = \frac{24(\theta + R_h)^2}{\tilde{N}} \leq \varepsilon,$$

where the strict inequality uses Proposition 2.4. This shows (2.14) and completes the proof of Theorem 2.1. □

Remark 2.4 (Continued discussion from Remark 2.2 comparing Theorem 2.1 with Khachiyan [35] for D -optimal design). Here $h = \iota_{\Delta_n}$, whereby $R_h = 0$, and the rate of linear convergence for iterates where $G_k > \theta$ in (2.11) is order $O(1 - 1/(n + \delta_0))$ as compared to the rate of $O(1 - 1/n)$ proved in [35] specifically for the D -optimal design problem with exact line-search. Due to the very special structure of the D -optimal design problem, the exact line-search is in closed-form, and it enables Khachiyan [35] to show that the optimality gap improvement bound (2.37) is instead

$$\delta_{k+1} \leq \delta_k - \omega\left(\frac{G_k}{G_k + \theta}\right). \quad (2.53)$$

Notice that ω is larger than ω^* , and all the moreso for larger values of its argument, which corresponds to $G_k > \theta$; this then leads to an improved guaranteed linear rate of Khachiyan's algorithm in the case when $G_k > \theta$. However, we stress that the stronger estimate in (2.53) is rather specific to the D -optimal design problem, and we do not expect it to hold in general for $f \in \mathcal{B}_\theta(\mathcal{K})$.

3 A Mirror Descent Method for the Dual Problem

In this section we present a mirror descent method with adaptive step-size applied to the (Fenchel) dual problem of (P) . We denote the dual problem of (P) by (D) , which is given by:

$$(D) : \quad -d^* := -\min_{y \in \mathbb{R}^m} [d(y) := f^*(y) + h^*(-A^*y)] , \quad (3.1)$$

where f^* and h^* are the Fenchel conjugates of f and h , respectively, and $A^* : \mathbb{R}^m \rightarrow \mathbb{R}^n$ denotes the adjoint of A . We observe the following properties related to (D) :

1. f^* is a θ -logarithmically-homogeneous self-concordant barrier on the polar of \mathcal{K} , namely $\mathcal{K}^\circ := \{y \in \mathbb{R}^m : \langle y, u \rangle \leq 0 \ \forall u \in \mathcal{K}\}$, and \mathcal{K}° is also a regular cone. This follows from [45, Theorem 2.4.4].
2. h^* is Lipschitz (but not necessarily differentiable) on \mathbb{R}^n . Indeed, let $\|\cdot\|$ be a given norm on the space of x variables, and define $R_{\mathcal{X}} := \max_{x \in \mathcal{X}} \|x\|$. Since $\mathcal{X} = \text{dom } h$ is compact and h is closed, it follows that $R_{\mathcal{X}} < +\infty$ and h^* is $R_{\mathcal{X}}$ -Lipschitz on \mathbb{R}^n .
3. $F^* = -d^*$ and (D) has at least one optimal solution. Indeed, since $A(\mathcal{X}) \cap \text{dom } f \neq \emptyset$ and f is continuous on $\text{dom } f$, the strong duality and attainment follows from [46, Theorem 3.50].

Although (D) has a similar structure as (P) , certain key differences are unattractive in regards to the application of first-order methods for solving (D) . One key difference is that the domain of the dual function d is unbounded, which is in obvious contrast to the bounded domain of the primal function F . This makes it difficult or prohibitive to apply a Frank-Wolfe type method to solve (D) . Furthermore, and similar to (P) , ∇f^* does not satisfy either uniform boundedness or uniform Lipschitz continuity on \mathcal{K}° , thereby preventing the application of most other types of first-order methods. Nevertheless, below we present a mirror descent method with adaptive step-size for (D) . Although the lack of good properties prevents the direct analysis of mirror descent in the usual manner, through the duality of mirror descent and the Frank-Wolfe method we provide a computational complexity bound for this mirror descent method.

Before presenting our mirror descent method for tackling (D) , we first present an important application of (D) in the Bregman proximal-point method.

Application in the Bregman proximal-point method (BPPM) [2, 9, 21]. Consider the convex non-smooth optimization problem $\min_{y \in \mathbb{R}_+^m} \xi(y)$, where $\xi : \mathbb{R}^m \rightarrow \mathbb{R}$ is assumed to be Lipschitz on \mathbb{R}^m . At the k -th iteration of BPPM, one solves the following problem:

$$y^{k+1} := \arg \min_{y \in \mathbb{R}^m} \xi(y) + \beta_k^{-1} D_\zeta(y, y^k) , \quad (3.2)$$

where y^k is the k -th iterate of BPPM, $\beta_k > 0$ is the step-size, and $\zeta : \mathbb{R}_+^m \rightarrow \mathbb{R}$ is the prox function that induces the Bregman divergence

$$D_\zeta(y, y^k) := \zeta(y) - \zeta(y^k) - \langle \nabla \zeta(y^k), y - y^k \rangle . \quad (3.3)$$

As pointed out in [2], one of the standard choices of ζ is $\zeta(y) = -\sum_{i=1}^m \ln(y_i)$, and under this choice, if $y^0 \in \mathbb{R}_+^m = \text{int } \mathbb{R}_+^m$, then $y^k \in \mathbb{R}_+^m$ for all $k \geq 1$, and so the constraint set \mathbb{R}_+^m is automatically taken care of by the prox-function ζ . From (3.2) and (3.3) we note that (3.2) is in the form of (D) with $f^*(y) := \zeta(y) = -\sum_{i=1}^m \ln(y_i)$, $A = -I$ and $h^*(-A^*y) = \beta_k \xi(y) - \langle \nabla \zeta(y^k), y \rangle$.

Algorithm 2 Mirror descent method for solving (D) using f^* as the prox function

Input: Starting points $y^0 \in \text{int } \mathcal{K}^\circ$ and $z^0 \in \{z \in \mathcal{X} : Az = \nabla f^*(y^0)\}$

At iteration $k \in \{0, 1, \dots\}$:

1. Let $s^k \in \partial h^*(-A^*y^k)$ and define

$$g^k := \nabla f^*(y^k) - As^k \in \partial d(y^k) . \quad (3.4)$$

2. If $k \geq 1$, compute

$$z^k := (1 - \gamma_{k-1})z^{k-1} + \gamma_{k-1}s^{k-1} . \quad (3.5)$$

3. Compute $\bar{G}_k := \langle g^k, y^k \rangle + h(z^k) - h(s^k)$ and $\bar{D}_k := \|g^k\|_{\nabla f^*(y^k)}$, and compute the step-size:

$$\gamma_k := \min \left\{ \frac{\bar{G}_k}{\bar{D}_k(\bar{G}_k + \bar{D}_k)}, 1 \right\} . \quad (3.6)$$

4. Update

$$y^{k+1} := \arg \min_{y \in \mathbb{R}^m} \langle g^k, y \rangle + \gamma_k^{-1} D_{f^*}(y, y^k) . \quad (3.7)$$

Our mirror descent method for (D) is shown in Algorithm 2, and is based on using the function f^* itself as the prox function to induce the Bregman divergence:

$$D_{f^*}(y, y^k) := f^*(y) - f^*(y^k) - \langle \nabla f^*(y^k), y - y^k \rangle ,$$

(When f is L -smooth, similar ideas have appeared in some previous works, for example Grigas [27], Bach [4], as well as Lu and Freund [39].) In step 1, we compute a subgradient of the dual function d at y^k , which is denoted by g^k . In step 2 we update the primal variables z^k which are used in the method to adaptively determine the step-size in the next step. In step 3, we compute the step-size γ_k , which we will show to be same as the step-size α_k in the Frank-Wolfe method (shown in Algorithm 1). Equipped with y^k , g^k and γ_k , in step 4 we perform a Bregman proximal minimization step to obtain y^{k+1} . We emphasize that different from the classical mirror descent method (e.g., [41]), in step 4 we use f^* (which is part of the objective function) as the prox function to induce the Bregman divergence $D_{f^*}(\cdot, \cdot)$. Also notice that the domain of the sub-problem (3.7) is $\text{int } \mathcal{K}^\circ$ and it is perhaps not so obvious without further analysis that (3.7) has an optimal solution.

At first glance it appears that Algorithm 2 might not be efficient to implement, since it involves working with a system of linear equations to determine z^0 in the Input, and also involves solving the minimization sub-problem in step 4. However, as we show below, Algorithm 2 corresponds exactly to the generalized Frank-Wolfe method in Algorithm 1 for solving (P) , which does not involve these computationally expensive steps. This of course implies that Algorithm 2 can be implemented via Algorithm 1 to obtain the primal iterate sequence $\{x^k\}_{k \geq 0}$, and then the dual iterate sequence $\{y^k\}_{k \geq 0}$ is determined by the simple rule $y^k = \nabla f(Ax^k)$ for $k \geq 0$.

Theorem 3.1. *Algorithms 1 and 2 are equivalent in the following sense: If the starting points x^0 in Algorithm 1 and z^0 in Algorithm 2 satisfy $x^0 = z^0$ and $y^0 = \nabla f(Ax^0)$, then an iterate sequence of either algorithm exactly corresponds to an iterate sequences of the other.*

Before we prove this theorem, we first recall some properties of conjugate functions. Let $w : \mathbb{R}^p \rightarrow \mathbb{R} \cup \{+\infty\}$ be a closed convex function and let w^* denote its conjugate function, which is

defined by $w^*(g) := \max_u \{\langle g, u \rangle - w(u)\}$. Then $w^* : \mathbb{R}^p \rightarrow \mathbb{R} \cup \{+\infty\}$ is a closed convex function, and

$$g \in \partial w(u) \iff u \in \partial w^*(g) \iff \langle g, u \rangle = w(u) + w^*(g) . \quad (3.8)$$

Proof of Theorem 3.1. Let $\{y^k\}_{k \geq 0}$ be the sequence of iterates of Algorithm 2, and let us also collect the sequences $\{z^k\}_{k \geq 0}$, $\{s^k\}_{k \geq 0}$, $\{g^k\}_{k \geq 0}$, $\{\bar{G}_k\}_{k \geq 0}$, $\{\bar{D}_k\}_{k \geq 0}$, and $\{\gamma^k\}_{k \geq 0}$ generated in Algorithm 2, and use these sequences to define the following five sequences by the simple assignments $x^k := z^k$, $v^k := s^k$, $\alpha^k := \gamma^k$, $G_k := \bar{G}_k$, $D_k := \bar{D}_k$, for $k \geq 0$. We now show that these five sequences correspond to an iterate sequence of Algorithm 1. Our argument will rely on the following identity:

$$y^k = \nabla f(\mathbf{A}z^k) \quad \forall k \geq 0 , \quad (3.9)$$

which we will prove by induction. Notice that (3.9) is true for $k = 0$ by supposition in the statement of the theorem. Now let us assume that (3.9) holds for a given iteration k , and let us examine the properties of our sequences. We have

$$v^k := s^k \in \partial h^*(-\mathbf{A}^*y^k) \implies -\mathbf{A}^*y^k \in \partial h(v^k) \quad (3.10)$$

$$\implies v^k \in \arg \min_{x \in \mathbb{R}^n} \langle \nabla f(\mathbf{A}x^k), \mathbf{A}x \rangle + h(x) , \quad (3.11)$$

where (3.11) follows from the conjugacy properties (3.8). This shows that v^k satisfies Step 1 of iteration k in Algorithm 1. We also have

$$G_k := \bar{G}_k := \langle g^k, y^k \rangle + h(z^k) - h(s^k) \quad (3.12)$$

$$= \langle \nabla f^*(y^k) - \mathbf{A}s^k, \nabla f(\mathbf{A}z^k) \rangle + h(z^k) - h(s^k) , \quad (3.13)$$

$$= \langle \mathbf{A}x^k - \mathbf{A}v^k, \nabla f(\mathbf{A}x^k) \rangle + h(x^k) - h(v^k) \quad (3.14)$$

satisfies the definition of G_k in Step 2 of iteration k of Algorithm 1. Similarly, we have

$$D_k := \bar{D}_k := \|g^k\|_{\nabla f^*(y^k)} = \|\nabla f^*(y^k) - \mathbf{A}s^k\|_{\mathbf{A}z^k} = \|\mathbf{A}z^k - \mathbf{A}s^k\|_{\mathbf{A}z^k} = \|\mathbf{A}x^k - \mathbf{A}v^k\|_{\mathbf{A}x^k} \quad (3.15)$$

satisfies the the definition of D_k in Step 2 of iteration k in Algorithm 1, which then implies similarly for the formula for α_k in (2.8) of Algorithm 1. Last of all, we prove the inductive step of the equality (3.9). From the optimality conditions of the optimization problem in (3.7) we have

$$\nabla f^*(y^{k+1}) = \nabla f^*(y^k) - \gamma_k g_k \implies \nabla f^*(y^{k+1}) = \nabla f^*(y^k) - \gamma_k g_k = (1 - \gamma_k)\mathbf{A}z^k + \gamma_k \mathbf{A}s^k ,$$

where in the last step we use $\nabla f^*(y^k) = \mathbf{A}z^k$ from (3.9) and (3.4). Since $z^{k+1} := (1 - \gamma_k)z^k + \gamma_k s^k$, we have $\nabla f^*(y^{k+1}) = \mathbf{A}z^{k+1}$ implies $y^{k+1} = \nabla f(\mathbf{A}z^{k+1})$ by conjugacy and completes the proof of (3.9). This then shows that an iterate sequence of Algorithm 2 corresponds to an iterate sequence of Algorithm 1. The reverse implication can also be proved using identical logic as above. \square

We now leverage the equivalence of Algorithms 1 and 2 to analyze the computational complexity of Algorithm 2. The following proposition relating the duality gap to the Frank-Wolfe gap will be useful.

Proposition 3.1. $G_k = d(y^k) + F(x^k)$ for all $k \geq 0$.

Proof. We have for all $k \geq 0$:

$$G_k = \langle \nabla f(\mathbf{A}x^k), \mathbf{A}x^k \rangle - \langle \nabla f(\mathbf{A}x^k), \mathbf{A}v^k \rangle + h(x^k) - h(v^k) \quad (3.16)$$

$$= f(\mathbf{A}x^k) + h(x^k) + f^*(y^k) + \langle -\mathbf{A}^*y^k, v^k \rangle - h(v^k) \quad (3.17)$$

$$= f(\mathbf{A}x^k) + h(x^k) + f^*(y^k) + h^*(-\mathbf{A}^*y^k) = d(y^k) + F(x^k). \quad (3.18)$$

where in (3.17) we use the conjugacy property in (3.8) and $y^k = \nabla f(\mathbf{A}x^k)$, and in (3.18) we use $-\mathbf{A}^*y^k \in \partial h(v^k)$ (by (3.10) and $s^k = v^k$) and (3.8). \square

Since G_k upper bounds the dual optimality gap $d_k := d(y^k) - d(y^*)$, from part 3 of Theorem 2.1 we immediately have the following corollary.

Corollary 1. *Let DGAP_ε denote the number of iterations required by Algorithm 2 to obtain $d_k \leq \varepsilon$. Then:*

$$\text{DGAP}_\varepsilon \leq \lceil 5.3(\delta_0 + \theta + R_h) \ln(10.6\delta_0) \rceil + \left\lceil \frac{24(\theta + R_h)^2}{\varepsilon} \right\rceil. \quad (3.19)$$

\square

We end this section with some remarks. First, if one considers (D) directly then its “primitive” objects are f^* and h^* , and implementing Algorithm 2 via Algorithm 1 requires knowing $h = (h^*)^*$ and also the Hessian of $f = (f^*)^*$ (used to compute the step-size α_k). However, these objects are not part of the primitive description of (D). Of course, for standard logarithmically-homogeneous barriers f^* such as $f^*(y) = -\sum_{i=1}^m \ln y_i$, where $y \in \mathbb{R}_{++}^m$ or $f^*(Y) = -\ln \det Y$, where $Y \in \mathbb{S}_{++}^p$ (and $m = p(p+1)/2$), its Fenchel conjugate f and the Hessian of f are well-known. In addition, for many simple non-smooth functions h^* , e.g., $h^*(w) = \|w\|_p$ ($p \in [1, +\infty]$), its Fenchel conjugate either is well-known or can be easily computed. Therefore, for many problems of interest implementing Algorithm 2 via Algorithm 1 is likely to be quite viable. Second, note that the step-size γ_k used in Algorithm 2 is adaptive, and is different from a standard step-size that is monotone decreasing in k , e.g., $\gamma_k = O(1/\sqrt{k})$ or $\gamma_k = O(1/k)$ (cf. [41]). This poses difficulties in attempting to directly analyze Algorithm 2 via standard approaches which involves using $D_{f^*}(y^k, y^*)$ as the potential function (see e.g., [3]). Nevertheless, the convergence guarantee for G_k derived from Algorithm 1 enables us to analyze the converge of the dual optimality gap d_k in Algorithm 2.

4 Computational Experiments

In this section we present the results of some basic numerical experiments where we evaluate the performance of our generalized Frank-Wolfe method in Algorithm 1 on the Poisson image de-blurring problem with TV regularization (Application 1 in Section 1.1), and also on the PET problem (Application 2 of Section 1.1).

4.1 First numerical experiment: Poisson image de-blurring with TV regularization

We consider the Poisson image de-blurring problem with TV regularization as described in Application 1 of Section 1.1, where the formulation was presented in equation (1.5). Observe that $f : u \mapsto -\sum_{l=1}^N y_l \ln(u_l)$ is neither Lipschitz nor L -smooth on the set $\{u \in \mathbb{R}^N : u = \mathbf{A}x, 0 \leq x \leq$

$Me\}$, and $\text{TV}(\cdot)$ does not have an efficiently computable proximal operator, which prevents most standard first-order methods from being applicable to tackle (1.5). As a result, very few methods have been proposed to solve (1.5) in the literature. In [15] an ad-hoc expectation-maximization (EM) method was proposed to solve (1.5), however the method is not guaranteed to converge due to the non-smoothness of the function $\text{TV}(\cdot)$ (see e.g., [47]). Both [32] and [10] proposed methods to solve a “perturbed” version of (1.5) by adding a small offset to each $\ln(\cdot)$ term, which of course, compromises the original objective function $\bar{F}(x)$ near $x = 0$. Using the generalized Frank-Wolfe method of Algorithm 1, we are able to directly solve formulation (1.5), the details of which we now describe.

4.1.1 Implementation of generalized Frank-Wolfe method for solving (1.5)

We first re-describe the total variation function $\text{TV}(x)$ by introducing some network definitions and terminology. The TV function penalizes potential differences on the horizontal and vertical grid arcs of the standard $m \times n$ pixel grid. Considering each pixel location as a node, let \mathcal{V} denote these nodes enumerated as $\mathcal{V} = [N] = \{1, \dots, N = m \times n\}$, and consider the undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where \mathcal{E} is the set of horizontal and vertical edges of the grid. These horizontal and vertical edges can be described as \mathcal{E}_h and \mathcal{E}_v , respectively, where

$$\mathcal{E}_h := \{\{l, l + 1\} : l \in [N], l \bmod n \neq 0\} \quad \text{and} \quad \mathcal{E}_v := \{\{l, l + n\} : l \in [N - m]\} .$$

With this notation we have

$$\text{TV}(x) = \sum_{\{i,j\} \in \mathcal{E}} |x_i - x_j| . \quad (4.1)$$

Based on \mathcal{G} , we define the directed graph $\tilde{\mathcal{G}} = (\mathcal{V}, \tilde{\mathcal{E}})$ where $\tilde{\mathcal{E}}$ is obtained by replacing each (undirected) edge in \mathcal{E} with two directed edges of opposite directions. Then from (4.1) we have

$$\begin{aligned} \text{TV}(x) &= \sum_{(i,j) \in \tilde{\mathcal{E}}} \max\{x_i - x_j, 0\} = \min e^\top r \\ &\quad \text{s. t. } r_{ij} \geq x_i - x_j, \quad \forall (i, j) \in \tilde{\mathcal{E}} \end{aligned} \quad (4.2)$$

$$\begin{aligned} &= \min e^\top r \\ &\quad \text{s. t. } r \geq B^\top x, \end{aligned} \quad (4.3)$$

where B is the node-arc incidence matrix of the network $\tilde{\mathcal{E}}$.

The Frank-Wolfe subproblem (1.3) associated with (1.5) is:

$$\min_{x \in \mathbb{R}^N} \langle \nabla \bar{f}(x^k), x \rangle + (\sum_{l=1}^N a_l)^\top x + \lambda \text{TV}(x) \quad \text{s. t. } 0 \leq x \leq Me, \quad (4.4)$$

where

$$\bar{f}(x) := -\sum_{l=1}^N y_l \ln(a_l^\top x) . \quad (4.5)$$

Based on (4.3), we can rewrite (4.4) as the following linear optimization problem:

$$\begin{aligned} \min_{(x,r) \in \mathbb{R}^N \times \mathbb{R}^{2|\mathcal{E}|}} \quad &\langle \nabla \bar{f}(x^k), x \rangle + (\sum_{l=1}^N a_l)^\top x + \lambda e^\top r \\ \text{s. t. } \quad &0 \leq x \leq Me, \quad r \geq B^\top x . \end{aligned} \quad (4.6)$$

This linear problem can be solved as a linear program (LP) using available software, or as a constrained dual network flow problem using available network flow software. In our implementation

we solved (4.6) using the primal simplex method in Gurobi [28]. Note that in (4.6), the variable r in the constraint set appears to be unbounded. However, from the form of the objective function and the definition of B , it is easy to see that any optimal solution (x^*, r^*) must satisfy $0 \leq r^* \leq Me$, and hence (4.6) always admits an optimal solution.

Given the representation of $\text{TV}(\cdot)$ in (4.3), we can also rewrite (1.5) as

$$\begin{aligned} \min_{x \in \mathbb{R}^N, r \in \mathbb{R}^{2|\mathcal{E}|}} \quad & \bar{f}(x) + (\sum_{l=1}^N a_l)^\top x + \lambda e^\top r \\ \text{s. t.} \quad & 0 \leq x \leq Me, \quad r \geq B^\top x. \end{aligned} \quad (4.7)$$

(Note that the Frank-Wolfe sub-problem (1.3) associated with (4.7) is the same as that associated with (1.5), which is shown in (4.6).) In the following, we will apply our Frank-Wolfe method to solve the reformulated problem (4.7). The advantage of (4.7) lies in that its structure yields an efficient procedure for an exact line-search to compute the step-size α_k . Specifically, given (x^k, r^k) , let (v^k, w^k) be an optimal solution of (4.6); then the exact line-search problem using (4.7) is:

$$\alpha_k = \arg \min_{\alpha \in [0,1]} \bar{f}(x^k + \alpha(v^k - w^k)) + \alpha((\sum_{l=1}^N a_l)^\top (v^k - x^k) + \lambda e^\top (w^k - r^k)). \quad (4.8)$$

The detailed description of the exact line-search procedure for problems of the format (4.8) are presented in Appendix F. Henceforth, we denote our generalized Frank-Wolfe method for Poisson de-blurring with the adaptive step-size in (2.8) as **FW-Adapt**, and we denote our Frank-Wolfe method with exact line-search of (4.8) as **FW-Exact**.

4.1.2 Results

We tested **FW-Adapt** and **FW-Exact** on the Shepp-Logan phantom image [50] of size 100×100 (hence $N = 10,000$). The true image X is shown in Figure 1(a); this image acts as a 2D slice of a representative 3D human brain image, and is a standard test image in image reconstruction algorithms. We generated the blurred noisy image Y shown in Figure 1(b) using the methodology exactly as described in Section 1.1. For both **FW-Adapt** and **FW-Exact**, we chose the starting point $x^0 = \text{vec}(Y)$, and we set $\lambda = 0.01$. In order to have an accurate computation of optimality gaps, we used CVXPY [16] to (approximately) find the optimal objective value \bar{F}^* of (1.5). All computations were performed in Python 3.8 on a laptop computer.

Figures 2(a) and 2(b) show (in log-log scale) the empirical optimality gaps $\bar{F}(x^k) - \bar{F}^*$ obtained by **FW-Adapt** and **FW-Exact** both as a function of the wall-clock time (in seconds) and as a function of the iteration counter k , respectively. From the figure we observe that **FW-Exact** converges faster than **FW-Adapt**, although the difference between the empirical optimality gaps of these two methods gradually lessens. This is expected since with exact line-search, **FW-Exact** can take a potentially larger step at each iteration than **FW-Adapt**, and likely therefore has faster convergence. The recovered images computed using **FW-Adapt** and **FW-Exact** are shown in Figure 3(a) and Figure 3(b), respectively. We observe that the image recovered by **FW-Adapt** has similar but slightly inferior quality compared to that recovered by **FW-Exact**. This is consistent with the algorithms' performance shown Figure 2, as **FW-Adapt** has a larger empirical optimality gap at termination compared to that of **FW-Exact**.

4.2 Second numerical experiment: positron emission tomography

We consider the positron emission tomography (PET) problem as described in Application 2 of Section 1.1, where the formulation was presented in equation (1.10). We generated artificial data




Figure 1: True and noisy 100×100 versions of the Shepp-Logan phantom image.




Figure 2: Comparison of empirical optimality gaps of FW-Adapt (FW-A) and FW-Exact (FW-E) for image recovery of the Shepp-Logan phantom image [50] of size 100×100 .

instances of this problem according to the following data generation process. Because the events emitted from each voxel i can only be detected by a small proportion of bins, it follows that the probability matrix P should be highly sparse. We chose a sparsity value of 5%. Given the number of voxels n and the number of bins m , for each $i \in [n]$ we randomly chose (without replacement) a subset of $[m]$ denoted by \mathcal{J}_i , such that $|\mathcal{J}_i| = \lfloor m/20 \rfloor$. Next, for all $j \in \mathcal{J}_i$ we then generated i.i.d. entries $\bar{p}_{ij} \sim U[0, 1]$ and normalized the values to obtain $p_{ij} := \bar{p}_{ij} / \sum_{j' \in \mathcal{J}_i} \bar{p}_{ij'}$. For all $j \notin \mathcal{J}_i$ we set $p_{ij} = 0$. We generated the mean values x_i for $i \in [n]$ by first generating i.i.d. values $\bar{x}_i \sim N(100, 9)$ and then set $x_i = \lfloor \bar{x}_i \rfloor$ for $i \in [n]$. We then simulated the event counts X_i at each voxel i by generating $X_i \sim \text{Poisson}(x_i)$ for $i \in [n]$. Finally, using P and $\{X_i\}_{i \in [n]}$, we generated the number of observed events Y_j detected at bin j by generating i.i.d. values $\tilde{Y}_j \sim \text{Poisson}(y_j)$ with $y_j := \sum_{i=1}^n p_{ij} X_i$ for $j \in [m]$. Since $Y_j \in \{0, 1, 2, \dots\}$, by omitting bins for which $Y_j = 0$ we ensure that $Y_j \geq 1$ for all $j \in [m]$ in the PET problem (1.10).




Figure 3: Recovered images computed using FW-Adapt and FW-Exact.

4.2.1 Comparison of Algorithms

We solved instances of the PET problem (1.10) using the following five algorithms/variants:

- **FW-Adapt** – our generalized Frank-Wolfe method in Algorithm 1 with the adaptive step-size as stated in (2.8),
- **FW-Exact** – our generalized Frank-Wolfe method in Algorithm 1 with the adaptive step-size (2.8) replaced by an exact line-search as described in detail in Appendix F,
- **RSGM-Fixed** – relatively smooth gradient method with fixed step-size [5, 40],
- **RSGM-LSBack** – relatively smooth gradient method with backtracking line-search [52],
- **EM** – a simple algorithm developed by Cover in 1984 specifically for a problem that is equivalent to the normalized PET problem [12].

We excluded mirror descent from our computational comparisons because the sparsity of P violates the basic assumption needed to apply the mirror descent (MD) method to (1.10) (see e.g., [6]). We now review the relevant details of the three algorithms **RSGM-F**, **RSGM-LS** and **EM**.

1. **RSGM-Fixed** [5, 40]. Although the objective function L of (1.10) is differentiable on its domain, its gradient ∇L is not Lipschitz on the constraint set Δ_n . Therefore standard gradient methods (or accelerated versions) [43] are not applicable. As a remedy for this situation, we can use the relatively-smooth gradient method [5, 40] to solve (1.10), which is designed in particular for problems whose objective functions have certain types of structured non-smooth gradients. Indeed, as shown in Bauschke et al. [5], L is \bar{Y} -smooth relative to the reference function

$$r(z) := -\sum_{i=1}^n \ln(z_i) \text{ for } z \in \mathbb{R}_{++}^n := \{x \in \mathbb{R}^n : x > 0\}, \quad (4.9)$$

where $\bar{Y} := \sum_{j=1}^m Y_j$. Specifically, this means that

$$\nabla^2 L(z) \preceq \bar{Y} \nabla^2 r(z) \quad \forall z \in \mathbb{R}_{++}^n. \quad (4.10)$$

Algorithm 3 RSGM for solving the PET problem (1.10)

Input: Starting point $z^0 \in \text{ri } \Delta_n := \{z > 0 : \sum_{i=1}^n z_i = 1\}$.

At iteration k :

1. Compute $\nabla L(z^k) = \sum_{j=1}^m (Y_j / \langle p_j, z^k \rangle) p_j$, where p_j is the vector corresponding to the j -th column of P for $j \in [m]$.
2. Choose step-size $\alpha_k > 0$ and compute

$$z^{k+1} = \arg \max_{z \in \Delta_n} \langle \nabla L(z^k), z \rangle - \alpha_k^{-1} D_r(z, z^k). \quad (4.11)$$

Algorithm 3 describes RSGM specified to solve the PET problem (1.10) using the reference function r with relative-smoothness parameter \bar{Y} , where $\text{ri } \Delta_n$ denotes the relative interior of Δ_n in the Input, and $D_r(\cdot, \cdot)$ denotes the Bregman divergence in (4.11) which is defined similarly as in (3.3). We set the step-size $\alpha_k = 1/\bar{Y}$ for all $k \geq 0$ in the fixed-step-size version of the method. Regarding the sub-problem (4.11) that needs to be solved at each iteration, its optimal solution is unique and lies in $\text{ri } \Delta_n$ since the reference function r is Legendre with domain \mathbb{R}_{++}^n (see [5, Section 2.1] for details). Therefore we have $z^k \in \text{ri } \Delta_n$ for all $k \geq 0$. To efficiently solve (4.11), we used the approach in [40, pp. 341-342], which reduces (4.11) to finding the unique positive root of a strictly decreasing univariate function on $(0, +\infty)$.

2. **RSGM-LSBack** [52]. This method is a version of the relatively smooth gradient method shown in Algorithm 3, with the extension that a backtracking line-search procedure is employed to compute the local relative-smoothness parameter \bar{Y}_k at z^k and then the step-size is chosen as $\alpha_k = 1/\bar{Y}_k$ at each iteration. (The details of this procedure can be found in [52, Algorithm 1].) Note that depending on the location of z^k , \bar{Y}_k may be (significantly) smaller than the (global) relative-smoothness parameter \bar{Y} .

3. **EM** [12]. This surprisingly simple algorithm was developed by Cover in 1984 specifically for a problem that is equivalent to the following normalized PET problem (see [12]):

$$\max \bar{L}(z) := \sum_{j=1}^m \bar{Y}_j \ln \left(\sum_{i=1}^n p_{ij} z_i \right) \quad \text{s. t. } z \in \Delta_n, \quad (4.12)$$

where $\bar{Y}_j := Y_j / \sum_{j'=1}^m Y_{j'} > 0$ for all $j \in [m]$ (recall that we have assumed without loss of generality that $Y_j \geq 1$ for all $j \in [m]$). The method starts with $z^0 \in \text{ri } \Delta_n$ and at each iteration k updates z^k as follows:

$$z_i^{k+1} := z_i^k \nabla_i \bar{L}(z^k) = \sum_{j=1}^m \bar{Y}_j \frac{p_{ij} z_i^k}{\sum_{l=1}^n p_{lj} z_l^k}, \quad \forall i \in [n]. \quad (4.13)$$

Note that since $\bar{Y}_j > 0$ for all $j \in [m]$, we easily see that $z^k \in \text{ri } \Delta_n$ for all $k \geq 0$.

4.2.2 Results

We report results for problems of dimensions $n = m = 1,000$, as we observed that the algorithms' relative performance were not particularly sensitive to the dimensions m and n . We ran the algorithms using two different choices of starting points: (i) $z^0 = z^{\text{bd}} \in \text{ri } \Delta_n$ chosen very close to $\text{bd } \mathbb{R}_+^n$ (the boundary of \mathbb{R}_+^n), and (ii) $z^0 = z^{\text{ct}} := \frac{1}{n} e$ which is the barycenter of Δ_n . To determine

z^{bd} we first used a greedy heuristic to determine a low- (or lowest-)cardinality index set $\mathcal{I} \subseteq [n]$ for which $\sum_{i \in \mathcal{I}} p_j > 0$, where p_j is the j -th column of P . Define $\bar{\delta} := 10^{-6}/n$, and we then set


$$z_i^{\text{bd}} := \begin{cases} \bar{\delta} & \text{for } i \notin \mathcal{I}, \\ (1 - (n - |\mathcal{I}|)\bar{\delta})/|\mathcal{I}| & \text{for } i \in \mathcal{I}. \end{cases}$$

Note that this ensures $z^{\text{bd}} \in \text{ri } \Delta_n$. Similar to Section (4.1), we used CVXPY [16] to (approximately) compute the optimal objective function value L^* of (1.10) in order to accurately report optimality gaps. Again, all computations were performed in Python 3.8 on a laptop computer.


Figures 4 and 5 show plots of the empirical optimality gaps $L(z^k) - L^*$ of all five methods with the starting points z^{bd} and z^{ct} , respectively. From Figure 4 we observe the following:

- (i) The relatively smooth gradient methods, namely **RSGM-Fixed** and **RSGM-LSBack**, make very little progress in reducing the empirical optimality gap. For **RSGM-Fixed**, this is because the relative smoothness parameter $\bar{Y} = \sum_{j=1}^m Y_j$ is typically very large, implying that the step-size $1/\bar{Y}$ is very small. Since the starting point z^{bd} is very close to $\text{bd } \mathbb{R}_+^n$, where the local relative smoothness parameter of the objective function L is close to \bar{Y} , **RSGM-LSBack** exhibits similar behavior to **RSGM-Fixed**. (In other words, the backtracking line-search does not help much in this case.)
- (ii) The two versions of our generalized Frank-Wolfe methods, namely **FW-Adapt** and **FW-Exact**, outperform the relatively smooth gradient methods. In addition, **FW-Exact** converges faster than **FW-Adapt** initially, but when close to the optimum (for example, when the empirical optimality gap falls below 10), these two methods have the same convergence behavior. Indeed, this observation also appears on the Poisson image de-blurring problem with TV-regularization (see Section 4.1.2).
- (iii) The **EM** algorithm outperforms all the other methods which is rather surprising at first glance. (And in fact it is unknown from the literature whether the method has any type of non-asymptotic convergence guarantee.) However, we note that this method, which uses a multiplicative form of update (see equation (4.13)), is specifically designed for problems with the PET problem structure in (1.10), and it is not clear how to suitably generalize this method to more general problems of the form (P) such as the Poisson image de-blurring problem with TV-regularization in (1.5).

Figure 5 shows the performance of the five algorithms when $z^0 = z^{\text{ct}}$ is the barycenter. We observe that the performance of the methods is mostly similar to when started at z^{bd} except for one significant difference, namely that **RSGM-LSBack** exhibits significantly faster convergence compared to when started at z^{ct} . Indeed, **RSGM-LSBack** outperforms all the other general purpose methods, namely **RSGM-Fixed**, **FW-Adapt** and **FW-Exact**. This is likely because on the ‘‘central region’’ of Δ_n the local relative smoothness parameter of L is probably much smaller than the global bound \bar{Y} , and therefore **RSGM-LSBack** is able to take a much larger steps. This also indicates that the convergence behavior of **RSGM-LSBack** is likely to be sensitive to the choice of starting point, which is not the case for the other methods (including our generalized Frank-Wolfe methods).




(a) Time plot: $n = 1000, m = 1000$




(b) Iteration plot: $n = 1000, m = 1000$

Figure 4: Comparison of optimality gaps of FW-Adapt (FW-A), FW-Exact (FW-E), RSGM-Fixed (RSGM-F), RSGM-LSBack (RSGM-LS) and EM, with $z^0 = z^{\text{bd}}$.



(a) Time plot: $n = 1000, m = 1000$



(b) Iteration plot: $n = 1000, m = 1000$

Figure 5: Comparison of optimality gaps of FW-Adapt (FW-A), FW-Exact (FW-E), RSGM-Fixed (RSGM-F), RSGM-LSBack (RSGM-LS) and EM, with $z^0 = z^{\text{ct}}$.

A Complexity of computing D_k in the D -optimal design problem

Let the objective function \bar{f} be defined in (1.15), and denote its Hessian at $x \in \text{dom } \bar{f}$ as $\bar{H}(x)$. Also denote $C := [a_1 \dots a_m] \in \mathbb{R}^{n \times m}$ and $X := \text{diag}(x)$. From standard results (e.g., [40, Proposition 2.2]), we know that $\nabla_i f(x) = -Q(x)_{ii}$ and $\bar{H}(x)_{ii} = Q(x)_{ii}^2$ for $i \in [m]$, where $Q(x) := C^\top(CXC^\top)^{-1}C$ and hence $Q(x)_{ii} = a_i^\top(CXC^\top)^{-1}a_i$. Since the constraint set in (1.15) is Δ_m , we can choose $v^k = e^{i_k}$, where $i_k \in \arg \min_{i \in [m]} \nabla_i f(x^k)$ and e^i denotes the i -th standard coordinate vector. Let $Q^k := Q(x^k)$. Now, note that we can write

$$\begin{aligned} D_k^2 &= \langle \bar{H}(x^k)(x^k - v^k), x^k - v^k \rangle = \langle \bar{H}(x^k)x^k, x^k \rangle - 2\langle \bar{H}(x^k)x^k, v^k \rangle + \langle \bar{H}(x^k)v^k, v^k \rangle \\ &\stackrel{(a)}{=} -\langle \nabla \bar{f}(x^k), x^k \rangle + 2\langle \nabla \bar{f}(x^k), e^{i_k} \rangle + \langle \bar{H}(x^k)e^{i_k}, e^{i_k} \rangle \end{aligned}$$

$$\begin{aligned}
&\stackrel{(b)}{=} n + 2\nabla_{i_k} \bar{f}(x^k) + \bar{H}(x^k)_{i_k, i_k} \\
&= n - 2Q_{i_k, i_k}^k + (Q_{i_k, i_k}^k)^2.
\end{aligned}$$

where in (a) we use (P6) and $v^k = e^{i_k}$ and in (b) we use (P7) and that the complexity parameter of \bar{f} is n . Therefore, the computational complexity of D_k is the same as that of $Q_{i_k, i_k}^k = a_{i_k}^\top B^k a_{i_k}$, where $B^k := (CX^k C^\top)^{-1}$. If $k = 0$, this can be computed in $O(mn^2 + n^3)$ time. In addition, since we have $\alpha_k < 1$ for all $k \geq 0$ (because $e^i \notin \text{dom } \bar{f}$ for any $i \in [m]$), the following holds:

$$\begin{aligned}
Q_{i_{k+1}, i_{k+1}}^k &= a_{i_{k+1}}^\top \left((1 - \alpha_k) CX^k C^\top + \alpha_k a_{i_k} a_{i_k}^\top \right)^{-1} a_{i_{k+1}} \\
&\stackrel{(a)}{=} (1 - \alpha_k)^{-1} a_{i_{k+1}}^\top \left(B^k - \frac{\beta_k B^k a_{i_k} a_{i_k}^\top B^k}{1 + \beta_k a_{i_k}^\top B^k a_{i_k}} \right) a_{i_{k+1}} \quad [\text{where } \beta_k := \alpha_k / (1 - \alpha_k)], \\
&= (1 - \alpha_k)^{-1} \left(a_{i_{k+1}}^\top B^k a_{i_{k+1}} - \frac{(a_{i_{k+1}}^\top B^k a_{i_k})^2}{\beta_k^{-1} + Q_{i_k, i_k}^k} \right).
\end{aligned}$$

Therefore, given B^k and Q_{i_k, i_k}^k , $Q_{i_{k+1}, i_{k+1}}^k$ can be computed in $O(n^2)$ time, for all $k \geq 0$.

B Proof of Proposition 2.1

Note that $(\omega^*)'(s) = s/(1+s) > 0$ for any $s > 0$, so ω^* is strictly increasing on $[0, +\infty)$. In addition, $0 \leq s \leq 1/2$ implies that $(\omega^*)'(s) \geq \frac{2}{3}s$, whereby

$$\omega^*(s) = \omega^*(0) + \int_0^s (\omega^*)'(t) dt \geq \int_0^s \frac{2}{3}t dt = s^2/3 \quad \text{for } s \in [0, 1/2].$$

Define the linear function $l(s) := (1 - 2 \ln(1.5))s$, and notice that $l(s) = \omega^*(s)$ at $s = 0$ and $s = 1/2$. Therefore from the convexity of ω^* it follows for $s \geq 1/2$ that

$$\omega^*(s) \geq l(s) = (1 - 2 \ln(1.5))s \geq s/5.3.$$

□

C Proof of Proposition 2.2

For any $s \in (-1, 1)$, we have

$$\begin{aligned}
\ln(1+s) &= s - \frac{s^2}{2} + \frac{s^3}{3} - \frac{s^4}{4} + \dots \\
&\geq s - \frac{|s|^2}{2} - \frac{|s|^3}{2} - \frac{|s|^4}{2} - \dots \\
&= s - \frac{|s|^2}{2} \left(1 + |s| + |s|^2 + \dots \right) \\
&= s - \frac{|s|^2}{2(1-|s|)}.
\end{aligned}$$

□

D Proof of inequality (2.18)

As shown in Khachiyan [35, Lemma 3], with $p_0 = (1/m)e$ we have

$$\delta_0 \leq n \ln(1 + \varepsilon_0) \quad \text{where } \varepsilon_0 := \frac{m}{n} \max_{i \in [m]} a_i^\top \left(\sum_{j=1}^m a_j a_j^\top \right)^{-1} a_i - 1. \quad (\text{D.1})$$

Define $C := [a_1 \dots a_m] \in \mathbb{R}^{n \times m}$ and $P := C^\top (CC^\top)^{-1} C$, and it holds for all $i \in [m]$ that:

$$a_i^\top \left(\sum_{j=1}^m a_j a_j^\top \right)^{-1} a_i = (e^i)^\top C^\top (CC^\top)^{-1} C e^i = \|P e^i\|_2^2 \leq \|e^i\|_2^2 = 1, \quad (\text{D.2})$$

where e^i denotes the i -th standard coordinate vector, and the second equality and the inequality above follow since P is a projection matrix. As a result, $\varepsilon_0 \leq m/n - 1$ and the proof follows using the left inequality in (D.1). \square

E Proof of Proposition 2.4

Note that (2.51) follows easily from

$$\frac{1}{d_{j+1}} \geq \frac{1}{d_j} + \frac{1}{M} \quad \forall j \geq 0, \quad (\text{E.1})$$

and so let us prove (E.1). It follows from the hypotheses of Proposition 2.4 that:

$$d_{j+1} \leq d_j - \frac{g_j^2}{M} \leq d_j - \frac{d_j^2}{M}, \quad (\text{E.2})$$

and dividing both sides of (E.2) by $d_j d_{j+1}$ and rearranging yields

$$\frac{1}{d_{j+1}} \geq \frac{1}{d_j} + \frac{d_j}{M d_j d_{j+1}} \geq \frac{1}{d_j} \left(1 + \frac{d_j}{M} \right) = \frac{1}{d_j} + \frac{1}{M},$$

where the second inequality uses $d_j \geq d_{j+1}$.

We next prove (2.52). The result is trivially true for $j = 0$, so let us consider $j \geq 1$. For convenience define $\bar{g}_j := \min\{g_0, \dots, g_j\}$, and consider some l satisfying some $0 < l \leq j$. We then have

$$\frac{M}{l} > d_l = d_{j+1} + \sum_{i=l}^j d_i - d_{i+1} \geq d_{j+1} + \sum_{i=l}^j \frac{g_i^2}{M} \geq (j-l+1) \frac{\bar{g}_j^2}{M},$$

where the first inequality uses (2.51) and the second inequality is from the first hypothesis of the proposition. If we take $l = \lceil j/2 \rceil$, then $j/2 \leq l \leq j/2 + 1$. Therefore

$$\bar{g}_j^2 < \frac{M^2}{l(j-l+1)} \leq \frac{M^2}{(j/2)^2} = \left(\frac{2M}{j} \right)^2,$$

which then implies (2.52). \square

F Exact line-search procedure to compute α_k

Let us define

$$\bar{F}(x) := -\sum_{i \in [m]} y_i \ln(a_i^\top x) + \xi^\top x, \quad (\text{F.1})$$

where $y_i \in \{1, 2, \dots\}$ and $a_i \in \mathbb{R}^n$ for $i \in [m]$ and $\xi \in \mathbb{R}^n$, and we wish to solve $\min_{x \in \mathcal{X}} \bar{F}(x)$. Note that the optimization problems (4.7) and (1.10) both conveniently fit into the format (F.1). Given $x \in \mathcal{X} \cap \text{dom } \bar{F}$, let $d \in \mathbb{R}^n$ be a descent direction at x (namely $\langle \nabla \bar{F}(x), d \rangle < 0$) such that $x + d \in \mathcal{X}$. The exact line-search problem involves finding

$$\alpha^* := \arg \min_{\alpha \in [0, 1]} [\zeta(\alpha) := \bar{F}(x + \alpha d)]. \quad (\text{F.2})$$

Define

$$\mathcal{I}_+ := \{i \in [m] : a_i^\top d > 0\} \quad \text{and} \quad \mathcal{I}_- := \{i \in [m] : a_i^\top d < 0\}, \quad (\text{F.3})$$

so that

$$\zeta(\alpha) \stackrel{c}{=} -\sum_{i \in \mathcal{I}_+} y_i \ln(a_i^\top x + \alpha a_i^\top d) - \sum_{i \in \mathcal{I}_-} y_i \ln(a_i^\top x + \alpha a_i^\top d) + \alpha \xi^\top d,$$

where $\stackrel{c}{=}$ denotes equality up to a constant independent of α . Consequently, we have

$$\begin{aligned} \zeta'(\alpha) &= -\sum_{i \in \mathcal{I}_+} \frac{y_i a_i^\top d}{a_i^\top x + \alpha a_i^\top d} - \sum_{i \in \mathcal{I}_-} \frac{y_i a_i^\top d}{a_i^\top x + \alpha a_i^\top d} + \xi^\top d \\ &= -\sum_{i \in \mathcal{I}_+} \frac{y_i}{\gamma_i + \alpha} - \sum_{i \in \mathcal{I}_-} \frac{y_i}{\gamma_i + \alpha} + \xi^\top d, \end{aligned}$$

where $\gamma_i = a_i^\top x / a_i^\top d$ for $i \in \mathcal{I}_+ \cup \mathcal{I}_-$. Clearly, we have $\gamma_i > 0$ for $i \in \mathcal{I}_+$ and $\gamma_i < 0$ for $i \in \mathcal{I}_-$. We have the following observations about ζ and ζ' :

- (i) Domain: $\text{dom } \zeta = \text{dom } \zeta' = (\gamma_-, \gamma_+)$, where $\gamma_- := -\min_{i \in \mathcal{I}_+} \gamma_i < 0$ and $\gamma_+ := -\max_{i \in \mathcal{I}_-} \gamma_i > 0$.
- (ii) Monotonicity: Since ζ is strictly convex, ζ' is strictly increasing on (γ_-, γ_+) . Also, $\zeta'(0) = \langle \nabla \bar{F}(x), d \rangle < 0$.
- (iii) Behavior as $\alpha \uparrow +\infty$: If $\mathcal{I}_- \neq \emptyset$, then $\gamma_+ < +\infty$ and $\zeta'(\alpha) \uparrow +\infty$ as $\alpha \uparrow \gamma_+$; if $\mathcal{I}_- = \emptyset$, then $\gamma_+ = +\infty$ and $\zeta'(\alpha) \uparrow \xi^\top d$ as $\alpha \uparrow +\infty$.

Based on these observations, our exact line search procedure can be described as follows:

- $\mathcal{I}_- = \emptyset$: We consider two cases, namely $\xi^\top d \leq 0$ and $\xi^\top d > 0$. If $\xi^\top d \leq 0$, then $\zeta'(\alpha) < 0$ for all $\alpha \geq 0$, and hence $\alpha^* = 1$. Otherwise, we first find α' as the unique solution of $\zeta'(\alpha) = 0$, namely $\alpha' = (\zeta')^{-1}(0)$. If $\alpha' > 1$, then $\zeta'(\alpha) < 0$ for all $\alpha \in [0, 1]$ and hence $\alpha^* = 1$; otherwise $\alpha^* = \alpha'$. In other words, $\alpha^* = \min\{\alpha', 1\}$.
- $\mathcal{I}_- \neq \emptyset$: In this case, by observation (iii), we see that $\zeta'(\alpha) = 0$ must have a unique solution α' on $(0, \gamma_+)$. To find such a solution, we can first find j as the largest integer such that $\zeta(\gamma_+ - 10^j) > 0$, and then use Newton's method starting from $\gamma_+ - 10^j$ or bisection over $[0, \gamma_+ - 10^j]$ to find α' . Once we find α' , we can let $\alpha^* = \min\{\alpha', 1\}$.

References

- [1] Atwood, C.L.: Optimal and efficient designs of experiments. *The Annals of Mathematical Statistics* pp. 1570–1602 (1969)
- [2] Auslender, A., Teboulle, M., Ben-Tiba, S.: Interior proximal and multiplier methods based on second order homogeneous kernels. *Math. Oper. Res.* **24**(3), 645–668 (1999)
- [3] Bach, F.: Duality between subgradient and conditional gradient methods. *SIAM J. Optim.* **25**(1), 115–129 (2015)
- [4] Bach, F.: Duality between subgradient and conditional gradient methods. *SIAM Journal on Optimization* **25**(1), 115—129 (2015)
- [5] Bauschke, H.H., Bolte, J., Teboulle, M.: A descent lemma beyond lipschitz gradient continuity: First-order methods revisited and applications. *Math. Oper. Res.* **42**(2), 330–348 (2017)
- [6] Ben-Tal, A., Margalit, T., Nemirovski, A.: The ordered subsets mirror descent optimization method with applications to tomography. *SIAM J. Optim.* **12**(1), 79–108 (2001)
- [7] Boyd, S., Vandenberghe, L.: *Convex Optimization*. Cambridge University Press (2004)
- [8] Canon, M.D., Cullum, C.D.: A tight upper bound on the rate of convergence of frank-wolfe algorithm. *SIAM J. Control* **6**(4), 509–516 (1968)
- [9] Censor, Y., Zenios, S.: Proximal minimization algorithm with d-functions. *J Optim. Theory Appl.* **73**, 451—464 (1992)
- [10] Chambolle, A., Ehrhardt, M.J., Richtárik, P., Schönlieb, C.B.: Stochastic primal-dual hybrid gradient algorithm with arbitrary sampling and imaging applications. *SIAM J. Optim.* **28**(4), 2783–2808 (2018)
- [11] Clarkson, K.: Coresets, sparse greedy approximation, and the Frank-Wolfe algorithm. 19th ACM-SIAM Symposium on Discrete Algorithms pp. 922–931 (2008)
- [12] Cover, T.: An algorithm for maximizing expected log investment return. *IEEE Transactions on Information Theory* **30**(2), 369–373 (1984)
- [13] Croux, C., Haesbroeck, G., Rousseeuw, P.J.: Location adjustment for the minimum volume ellipsoid estimator. *Statistics and Computing* **12**(3), 191–200 (2002)
- [14] Demyanov, V., Rubinov, A.: The minimization of a smooth convex functional on a convex set. *SIAM Journal on Control* (1967)
- [15] Dey, N., Blanc-Féraud, L., Zimmer, C., Roux, P., Kam, Z., Olivo-Marin, J., Zerubia, J.: Richardson-lucy algorithm with total variation regularization for 3d confocal microscope deconvolution. *Microsc. Res. Tech.* **69** (2006)
- [16] Diamond, S., Boyd, S.: CVXPY: A Python-embedded modeling language for convex optimization. *J. Mach. Learn. Res.* **17**(83), 1–5 (2016)

- [17] Dunn, J.: Rates of convergence for conditional gradient algorithms near singular and nonsingular extremals. *SIAM Journal on Control and Optimization* **17**(2), 187–211 (1979)
- [18] Dunn, J.: Convergence rates for conditional gradient sequences generated by implicit step length rules. *SIAM Journal on Control and Optimization* **18**(5), 473–487 (1980)
- [19] Dunn, J., Harshbarger, S.: Conditional gradient algorithms with open loop step size rules. *Journal of Mathematical Analysis and Applications* **62**, 432–444 (1978)
- [20] Dvurechensky, P., Ostroukhov, P., Safin, K., Shtern, S., Staudigl, M.: Self-concordant analysis of frank-wolfe algorithms. In: *Proc. ICML*, pp. 2814–2824 (2020)
- [21] Eckstein, J.: Nonlinear proximal point algorithms using bregman functions, with applications to convex programming. *Math. Oper. Res.* **18**(1), 202–226 (1993)
- [22] Fedorov, V.V.: *Theory of Optimal Experiments*. Academic Press (1972)
- [23] Frank, M., Wolfe, P.: An algorithm for quadratic programming. *Nav. Res. Logist. Q.* **3**(1-2), 95–110 (1956)
- [24] Freund, R.M., Grigas, P.: New analysis and results for the frank-wolfe method. *Math. Program.* **155**, 199—230 (2016)
- [25] Freund, R.M., Grigas, P., Mazumder, R.: An extended frank-wolfe method with “in-face” directions, and its application to low-rank matrix completion. *SIAM J. Optim.* **27**(1), 319–346 (2017)
- [26] Ghadimi, S.: Conditional gradient type methods for composite nonlinear and stochastic optimization. *Math. Program.* **173**, 431—464 (2019)
- [27] Grigas, P.: *Methods for convex optimization and statistical learning*. Ph.D. thesis, Massachusetts Institute of Technology (2016)
- [28] Gurobi Optimization, L.: *Gurobi optimizer reference manual* (2021). URL <http://www.gurobi.com>
- [29] Hager, W.: Updating the inverse of a matrix. *SIAM Review* **31**, 221–239 (1989)
- [30] Harchaoui, Z., Juditsky, A., Nemirovski, A.: Conditional gradient algorithms for norm-regularized smooth convex optimization. *Math. Program.* **152**, 75–112 (2015)
- [31] Hariharan, B.: All about convolution (2018). URL <https://www.cs.cornell.edu/courses/cs4670/2018sp/lec04-filtering2.pdf>
- [32] Harmany, Z.T., Marcia, R.F., Willett, R.M.: This is spiral-tap: Sparse poisson intensity reconstruction algorithms—theory and practice. *IEEE Transactions on Image Processing* **21**(3), 1084–1096 (2012)
- [33] Jaggi, M.: Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In: *Proc. ICML*, pp. 427–435 (2013)

- [34] John, F.: Extremum problems with inequalities as subsidiary conditions. in *Studies and Essays, Presented to R. Courant on His 60th Birthday*, Interscience, New York **30**, 187–204 (1948)
- [35] Khachiyan, L.G.: Rounding of polytopes in the real number model of computation. *Mathematics of Operations Research* **21**(2), 307–320 (1996)
- [36] Kiefer, J., Wolfowitz, J.: The equivalence of two extremum problems. *Canadian Journal of Mathematics* **12**(5), 363–365 (1960)
- [37] Knorr, E.M., Ng, R.T., Zamar, R.H.: Robust space transformations for distance-based operations. In: *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 126–135. ACM (2001)
- [38] Levitin, E., Polyak, B.: Constrained minimization methods. *USSR Comput. Math. Math. Phys.* **6**(5), 1–50 (1966)
- [39] Lu, H., Freund, R.: Generalized stochastic frank-wolfe algorithm with stochastic “substitute” gradient for structured convex optimization. *Mathematical Programming* (2021). To appear
- [40] Lu, H., Freund, R.M., Nesterov, Y.: Relatively smooth convex optimization by first-order methods, and applications. *SIAM J. Optim.* **28**(1), 333–354 (2018)
- [41] Nemirovskii, A., Yudin, D.: Efficient methods for large-scale convex problems. *Ekonomika i Matematicheskie Metody* (in Russian) **15**, 135–152 (1979)
- [42] Nesterov, Y.: *Introductory Lectures on Convex Optimization: A Basic Course*. Springer (2004)
- [43] Nesterov, Y.: Gradient methods for minimizing composite functions. *Math. Program.* **140**(1), 125–161 (2013)
- [44] Nesterov, Y.: Complexity bounds for primal-dual methods minimizing the model of objective function. *Math. Program.* **171**, 311–330 (2018)
- [45] Nesterov, Y., Nemirovskii, A.: *Interior-Point Polynomial Algorithms in Convex Programming*. SIAM (1994)
- [46] Peypouquet, J.: *Convex optimization in normed spaces : theory, methods and examples*. Springer (2015)
- [47] de Pierro, A.R.: On the convergence of an em-type algorithm for penalized likelihood estimation in emission tomography. *IEEE Trans. Med. Imaging* **14**(4), 762–765 (1995)
- [48] Renegar, J.: *A Mathematical View of Interior-point Methods in Convex Optimization*. SIAM, Philadelphia, PA, USA (2001)
- [49] Rudin, L.I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. *Phys. D: Nonlinear Phenom.* **60**(1), 259–268 (1992)
- [50] Shepp, L.A., Logan, B.F.: The fourier reconstruction of a head section. *IEEE Trans. Nucl. Sci.* **21**(3), 21–43 (1974)

- [51] Shepp, L.A., Vardi, Y.: Maximum likelihood reconstruction for emission tomography. *IEEE Trans. Med. Imaging* **1**(2), 113–122 (1982)
- [52] Stonyakin, F., Tyurin, A., Gasnikov, A., Dvurechensky, P., Agafonov, A., Dvinskikh, D., Alkousa, M., Pasechnyuk, D., Artamonov, S., Piskunova, V.: Inexact relative smoothness and strong convexity for optimization and variational inequalities by inexact model. [arXiv:2001.09013](https://arxiv.org/abs/2001.09013) (2020)
- [53] Sun, P., Freund, R.M.: Computation of minimum-volume covering ellipsoids. *Operations Research* **52**(5), 690–706 (2004)
- [54] Todd, M.J.: *Minimum-Volume Ellipsoids: Theory and Algorithms*. SIAM (2016)
- [55] Vardi, Y., Lee, D.: From image deblurring to optimal investments: Maximum likelihood solutions for positive linear inverse problems. *J. R. Stat. Soc. Ser. B* **55**(3), 569–598 (1993)
- [56] Yang, C., Fan, J., Wu, Z., Udell, M.: Efficient automl pipeline search with matrix and tensor factorization (2020)
- [57] Yildirim, E.A.: On the minimum volume covering ellipsoid of ellipsoids. *SIAM Journal on Optimization* **17**(3), 621–641 (2006)