

Complexity of Linear Minimization and Projection on Some Sets

Cyrille W. Combettes^{1 3}

cyrille@gatech.edu

Sebastian Pokutta^{2 3}

pokutta@zib.de

¹ *School of Industrial and Systems Engineering, Georgia Institute of Technology, USA*

² *Institute of Mathematics, Technische Universität Berlin, Germany*

³ *Department for AI in Society, Science, and Technology, Zuse Institute Berlin, Germany*

Abstract

The Frank-Wolfe algorithm is a method for constrained optimization that relies on linear minimizations, as opposed to projections. Therefore, a motivation put forward in a large body of work on the Frank-Wolfe algorithm is the computational advantage of solving linear minimizations instead of projections. However, the discussions supporting this advantage are often too succinct or incomplete. In this paper, we review the complexity bounds for both tasks on several sets commonly used in optimization. Projection methods onto the ℓ_p -ball, $p \in]1, 2[\cup]2, +\infty[$, and the Birkhoff polytope are also proposed.

1 Introduction

We consider the constrained optimization problem

$$\min_{x \in \mathcal{C}} f(x), \tag{1}$$

where $\mathcal{C} \subset \mathbb{R}^n$ is a compact convex set and $f: \mathbb{R}^n \rightarrow \mathbb{R}$ is a smooth function. Among all general purpose methods addressing problem (1), the Frank-Wolfe algorithm [17], a.k.a. conditional gradient algorithm [41], has the particularity of never requiring projections onto \mathcal{C} . It uses linear minimizations over \mathcal{C} instead and is therefore often referred to as a *projection-free* algorithm in the literature, in the sense that it does not call for solutions to quadratic optimization subproblems.

Thus, a motivation put forward in a large body of work on the Frank-Wolfe algorithm is the computational advantage of solving linear minimizations instead of projections. However, only a few works actually provide examples. On the other hand, the complexities of linear minimizations over several sets are available in [27, 29, 19], but they do not always (accurately) discuss the complexities of the respective projections. Therefore, while it is intuitive that a linear minimization is simpler to solve than a projection in general, a complete quantitative assessment is necessary to properly motivate the projection-free property of the Frank-Wolfe algorithm.

Contributions. We review the complexity bounds of linear minimizations and projections on several sets commonly used in optimization: the standard simplex, the ℓ_p -balls for $p \in [1, +\infty]$, the nuclear norm-ball, the flow polytope, the Birkhoff polytope, and the permutahedron. These sets are selected because linear minimizations or projections can be solved very efficiently, rather than by resorting to a general purpose method, in which case the analysis is less interesting. We also propose two methods for projecting onto the ℓ_p -ball and the Birkhoff polytope respectively,

and we analyze their complexity. Computational experiments for the ℓ_1 -ball and the nuclear norm-ball are presented.

Remark 1.1. *We would like to stress that, while it is possible that a projection-based algorithm requires less iterations than the Frank-Wolfe algorithm to find a solution to problem (1), our goal here is to demonstrate its advantage in terms of per-iteration complexity. We discuss the Frank-Wolfe algorithm and some successful applications in Section 2.2.*

2 Preliminaries

2.1 Notation and definitions

We work in the Euclidean space \mathbb{R}^n or $\mathbb{R}^{m \times n}$ equipped with the standard scalar product $\langle x, y \rangle = x^\top y$ or $\langle X, Y \rangle = \text{tr}(X^\top Y)$. We denote by $\|\cdot\|$ the norm induced by the scalar product, i.e., the ℓ_2 -norm $\|\cdot\|_2$ or the Frobenius norm $\|\cdot\|_F$ respectively. For any closed convex set $\mathcal{C} \subset \mathbb{R}^n$, the projection operator onto \mathcal{C} , the distance function to \mathcal{C} , and the diameter of \mathcal{C} , all with respect to $\|\cdot\|$, are denoted by $\text{proj}(\cdot, \mathcal{C})$, $\text{dist}(\cdot, \mathcal{C})$, and $\text{diam}(\mathcal{C})$ respectively.

For every $i, j \in \mathbb{N}$ such that $i \leq j$, the brackets $\llbracket i, j \rrbracket$ denote the set of integers between (and including) i and j . For all $x \in \mathbb{R}^n$ and $i, j \in \llbracket 1, n \rrbracket$ such that $i \leq j$, $[x]_i$ denotes the i -th entry of x and $[x]_{i:j} = ([x]_i, \dots, [x]_j)^\top \in \mathbb{R}^{j-i+1}$. The signum function is $\text{sign}: \lambda \in \mathbb{R} \mapsto 1$ if $\lambda > 0$, -1 if $\lambda < 0$, and 0 if $\lambda = 0$. The characteristic function of an event E is $\mathbb{1}_E = 1$ if E is true, else 0 . The indicator function of a set $\mathcal{C} \subset \mathbb{R}^n$ is $\iota_{\mathcal{C}}: x \in \mathbb{R}^n \mapsto 0$ if $x \in \mathcal{C}$, else $+\infty$. Operations on vectors in \mathbb{R}^n , such as $\text{sign}(x)$, $|x|$, x^p , $\max\{x, y\}$, xy , that are conventionally applied to scalars, are carried out entrywise and return a vector in \mathbb{R}^n . The shape of 0 and 1 will be clear from context, i.e., a scalar or a vector. The identity matrix in $\mathbb{R}^{n \times n}$ is denoted by I_n . The matrix with all ones in $\mathbb{R}^{m \times n}$ is denoted by $J_{m,n}$, and by J_n if $m = n$.

We adopt the real-number infinite-precision model of computation. The complexity of a computational task is the number of arithmetic operations necessary to execute it. We ran the experiments on a laptop under Linux Ubuntu 20.04 with Intel Core i7-10750H. The code is available at <https://github.com/cyrillewcombettes/complexity>.

2.2 The Frank-Wolfe algorithm

The Frank-Wolfe algorithm (FW) [17], a.k.a. conditional gradient algorithm [41], is a first-order projection-free algorithm for solving constrained optimization problems (1). It is presented in Algorithm 1.

Algorithm 1 Frank-Wolfe (FW)

Input: Start point $x_0 \in \mathcal{C}$, step-size strategy $(\gamma_t)_{t \in \mathbb{N}} \subset [0, 1]$.

- 1: **for** $t = 0$ **to** $T - 1$ **do**
 - 2: $v_t \leftarrow \arg \min_{v \in \mathcal{C}} \langle v, \nabla f(x_t) \rangle$
 - 3: $x_{t+1} \leftarrow x_t + \gamma_t(v_t - x_t)$
 - 4: **end for**
-

At each iteration, FW minimizes the linear approximation of f at x_t over \mathcal{C} (Line 2), i.e.,

$$\min_{v \in \mathcal{C}} f(x_t) + \langle v - x_t, \nabla f(x_t) \rangle,$$

and then moves in the direction of a solution $v_t \in \mathcal{C}$ with a step-size $\gamma_t \in [0, 1]$ (Line 3). This ensures that the new iterate $x_{t+1} = (1 - \gamma_t)x_t + \gamma_tv_t \in \mathcal{C}$ is feasible by convexity, and there is no need for a projection back onto \mathcal{C} . For this *projection-free* property, FW has encountered numerous applications, including solving traffic assignment problems [39], performing video co-localization [30], or, e.g., developing adversarial attacks [6].

When f is convex, FW converges at a rate $f(x_t) - \min_{\mathcal{C}} f = \mathcal{O}(1/t)$ for different step-size strategies [17, 15, 29], which is optimal in general [5, 29]. Faster rates can be established under additional assumptions on the properties of f or the geometry of \mathcal{C} [41, 23, 20, 31]. Recently, several variants have also been developed to improve its performance [36, 38, 21, 18, 4, 9]. When f is nonconvex, [35] showed that FW converges to a stationary point at a rate $\mathcal{O}(1/\sqrt{t})$ in the gap $\max_{v \in \mathcal{C}} \langle x_t - v, \nabla f(x_t) \rangle$ [28], which has inspired a line of work in stochastic optimization [48, 52, 51, 10, 53].

Lastly, note that FW is also popular for the natural sparsity of its iterates with respect to the vertices of \mathcal{C} , as $x_t \in \text{conv}\{x_0, v_0, \dots, v_{t-1}\}$ [7, 26, 37, 43, 8].

3 Projections versus linear minimizations

The Frank-Wolfe algorithm avoids projections by computing linear minimizations instead. In Table 1, we summarize the complexities of a linear minimization and a (Euclidean) projection on several sets commonly used in optimization. That is, we compare the complexities of solving

$$\min_{x \in \mathcal{C}} \langle x, y \rangle \quad \text{and} \quad \min_{x \in \mathcal{C}} \|x - y\|. \quad (2)$$

When an exact solution cannot be computed directly, we compare to the complexity of finding an ε -approximate solution; note that the two objectives in (2) are homogeneous. For the projection problem, it means to solve $\min_{x \in \mathcal{C}} \|x - y\|^2$ using any method but the Frank-Wolfe algorithm, since it would go against the purpose of this paper. Note however that the Frank-Wolfe algorithm can generate a solution with complexity $\mathcal{O}(\text{iter}(\mathcal{C}) \text{diam}(\mathcal{C})^2 / \varepsilon^2)$, where $\text{iter}(\mathcal{C})$ denotes the complexity of an iteration, which amounts to that of a linear minimization over \mathcal{C} . When addressing problem (1), solving projection subproblems via the Frank-Wolfe algorithm is known as conditional gradient sliding [38].

Table 1: Complexities of linear minimizations and (Euclidean) projections on some sets commonly used in optimization. We denote by x^* a solution, $\rho = p \sup_{t \in \mathbb{N}} \|x_t\|_{2(p-1)}^{p-1} \|x_t\|_2 < +\infty$ where $(x_t)_{t \in \mathbb{N}}$ is the sequence generated by Algorithm 2, by ν and σ_1 the number of nonzero entries and the top singular value of $-Y$ respectively, and by $\varepsilon > 0$ the additive error in the objective of (2) when an approximate solution is computed. The constant d_z is defined in (12) and $\tilde{\mathcal{O}}$ hides polylogarithmic factors in m and n .

Set \mathcal{C}	Linear minimization	Projection	Reference
ℓ_p -ball, $p \in \{1, 2, +\infty\}$	$\mathcal{O}(n)$	$\mathcal{O}(n)$	Sections 3.1–3.2
ℓ_p -ball, $p \in]1, 2[\cup]2, +\infty[$	$\mathcal{O}(n)$	$\mathcal{O}(n\rho^2 \ y - x^*\ _2^2 / \varepsilon^2)$	Section 3.3
Nuclear norm-ball	$\mathcal{O}(\nu \ln(m+n) \sqrt{\sigma_1} / \sqrt{\varepsilon})$	$\mathcal{O}(mn \min\{m, n\})$	Section 3.4
Flow polytope	$\mathcal{O}(m+n)$	$\tilde{\mathcal{O}}(m^3n + n^2)$	Section 3.5
Birkhoff polytope	$\mathcal{O}(n^3)$	$\mathcal{O}(n^2 d_z^2 / \varepsilon^2)$	Section 3.6
Permutahedron	$\mathcal{O}(n \ln(n))$	$\mathcal{O}(n \ln(n) + n)$	Section 3.7

We now provide details for the complexities reported in Table 1. Slightly abusing notation though we may have $\text{card}(\arg \min_{x \in \mathcal{C}} \langle x, y \rangle) > 1$, we write $\arg \min_{x \in \mathcal{C}} \langle x, y \rangle = x^*$ instead of $\arg \min_{x \in \mathcal{C}} \langle x, y \rangle \ni x^*$.

3.1 The ℓ_1 -ball and the standard simplex

Let $\{e_1, \dots, e_n\}$ denote the standard basis in \mathbb{R}^n . The ℓ_1 -ball is

$$\{x \in \mathbb{R}^n \mid \|x\|_1 \leq 1\} = \left\{ x \in \mathbb{R}^n \mid \sum_{i=1}^n |[x]_i| \leq 1 \right\} = \text{conv}\{\pm e_1, \dots, \pm e_n\},$$

and the standard simplex is

$$\Delta_n = \{x \in \mathbb{R}^n \mid \langle x, 1 \rangle = 1, x \geq 0\} = \text{conv}\{e_1, \dots, e_n\}.$$

A projection onto the ℓ_1 -ball amounts to computing a projection onto the standard simplex, for which the most efficient algorithms have a complexity $\mathcal{O}(n)$; see [12] for a review. On the other hand, linear minimizations are available in closed form: for all $y \in \mathbb{R}^n$,

$$\arg \min_{x \in \Delta_n} \langle x, y \rangle = e_{i_{\min}} \quad \text{and} \quad \arg \min_{\|x\|_1 \leq 1} \langle x, y \rangle = -\text{sign}([y]_{i_{\max}}) e_{i_{\max}},$$

where $i_{\min} \in \arg \min_{i \in [1, n]} [y]_i$ and $i_{\max} \in \arg \max_{i \in [1, n]} |[y]_i|$. Thus, while their complexities can both be written $\mathcal{O}(n)$, in practice linear minimizations are much simpler to solve than projections. Figure 1 presents a computational comparison. The results are averaged over 5 runs and the shaded areas represent ± 1 standard deviation.




Figure 1: Solving a linear minimization and a projection on the ℓ_1 -ball. The input vector $y \in \mathbb{R}^n$ is generated by sampling entries from the standard normal distribution and the projection method is [12, Fig. 2], which is state-of-the-art in practice [12, Tab. 3]. In this situation, the plots suggest that linear minimizations are about $100\times$ faster to solve than projections when n is large enough.

3.2 The ℓ_2 -ball and the ℓ_∞ -ball

For all $x \in \mathbb{R}^n$, $\|x\|_2 = \sqrt{\sum_{i=1}^n [x]_i^2}$ and $\|x\|_\infty = \max_{i \in [1, n]} |[x]_i|$. Here, linear minimizations have no significant advantage over projections as they are all available in closed form. For all $y \in \mathbb{R}^n$,

$$\arg \min_{\|x\|_2 \leq 1} \langle x, y \rangle = \frac{y}{\|y\|_2} \quad \text{and} \quad \arg \min_{\|x\|_2 \leq 1} \|x - y\|_2 = \frac{y}{\max\{\|y\|_2, 1\}},$$

and

$$\arg \min_{\|x\|_\infty \leq 1} \langle x, y \rangle = -\text{sign}(y) \quad \text{and} \quad \arg \min_{\|x\|_\infty \leq 1} \|x - y\|_2 = \text{sign}(y) \min\{|y|, 1\}.$$

3.3 The ℓ_p -balls for $p \in]1, +\infty[$

Let $p \in]1, +\infty[$. The ℓ_p -ball is

$$\{x \in \mathbb{R}^n \mid \|x\|_p \leq 1\} = \left\{ x \in \mathbb{R}^n \mid \left(\sum_{i=1}^n |[x]_i|^p \right)^{1/p} \leq 1 \right\}.$$

Linear minimizations are available in closed form: by duality, for all $y \in \mathbb{R}^n$,

$$\arg \min_{\|x\|_p \leq 1} \langle x, y \rangle = -\nabla \|\cdot\|_q(y) = -\frac{\text{sign}(y)|y|^{q-1}}{\|y\|_q^{q-1}},$$

where $q = p/(p-1) \in]1, +\infty[$. To the best of our knowledge, there is no projection method specific to the ℓ_p -ball when $p \in]1, 2[\cup]2, +\infty[$. We use [11, Alg. 6.5], a Haugazeau-like algorithm [25] for projecting onto the intersection of sublevel sets of convex functions. For a single sublevel set, the problem reads

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \quad & \|x - y\|_2^2 \\ \text{s.t.} \quad & g(x) \leq 0, \end{aligned} \tag{3}$$

and we assume that $g: \mathbb{R}^n \rightarrow \mathbb{R}$ is convex and differentiable for ease of exposition. In our case, $g = \|\cdot\|_p^p - 1$. Alternatively, one could use a Lagrange multiplier to formulate (3) as a strongly convex unconstrained problem, but finding the corresponding multiplier may require a considerable effort of tuning; also note that information is usually given in the form $g(x) \leq 0$ rather than in the form of a Lagrange multiplier. The method is presented in Algorithm 2, where for all $a, b \in \mathbb{R}^n$,

$$H(a, b) = \{x \in \mathbb{R}^n \mid \langle x - b, a - b \rangle \leq 0\},$$

and z_t is the projection of x_t onto $\{x \in \mathbb{R}^n \mid g(x_t) + \langle x - x_t, \nabla g(x_t) \rangle \leq 0\} = H(x_t, z_t)$. If $g(x_t) \leq 0$, then $x_{t+s} = x^*$ for all $s \in \mathbb{N}$ [11, Prop. 3.1], where x^* is the solution to problem (3).

Algorithm 2 Haugazeau-like for problem (3)

Input: Point to project $y \in \mathbb{R}^n$.

- 1: $x_0 \leftarrow y$
 - 2: **for** $t = 0$ **to** $T - 1$ **do**
 - 3: $z_t \leftarrow x_t - \frac{g(x_t)}{\|\nabla g(x_t)\|_2^2} \nabla g(x_t) \mathbb{1}_{\{g(x_t) > 0\}}$
 - 4: $x_{t+1} \leftarrow \text{proj}(x_0, H(x_0, x_t) \cap H(x_t, z_t))$
 - 5: **end for**
-

The projection in Line 4 is available in closed form [25, Thm. 3-1]; see also [2, Cor. 29.25]. The complexity of an iteration of Algorithm 2 is $\mathcal{O}(n)$. We propose in Theorem 3.1 the convergence rate of Algorithm 2, based on a key result from [1, Thm. 7.12]. A convergence rate is also proposed in [46], however it uses a stronger assumption and has a minor error in the exponent of the constant.

Theorem 3.1. *Let $g: \mathbb{R}^n \rightarrow \mathbb{R}$ be a differentiable convex function and $\mathcal{C} = \{x \in \mathbb{R}^n \mid g(x) \leq 0\}$, and suppose that there exists $\hat{x} \in \mathcal{C}$ such that $g(\hat{x}) < 0$. Consider Algorithm 2 and let $x^* = \text{proj}(x_0, \mathcal{C})$. Then, for all $t \in \mathbb{N}$,*

$$\|x^* - x_0\|_2^2 - \|x_t - x_0\|_2^2 \leq \frac{\max\{8\rho^2, 2\} \|x_0 - x^*\|_2^2}{t + 2}, \tag{4}$$

and

$$\|x_t - x^*\|_2 \leq \frac{\max\{2\sqrt{2}\rho, \sqrt{2}\}\|x_0 - x^*\|_2}{\sqrt{t+2}}, \quad (5)$$

where $\rho = (-1/g(\hat{x})) \sup_{t \in \mathbb{N}} \|\nabla g(x_t)\|_2 \|x_t - \hat{x}\|_2 < +\infty$.

Proof. First, note that by [11, Prop. 3.1], for all $t \in \mathbb{N}$,

$$\|x_t - x_0\|_2 \leq \|x_{t+1} - x_0\|_2 \leq \|x^* - x_0\|_2. \quad (6)$$

We prove by induction that (4) holds for all $t \in \mathbb{N}$. The base case $t = 0$ is trivial. Suppose that (4) holds at iteration $t \in \mathbb{N}$. Since $x_{t+1} \in H(x_0, x_t)$ and $x_{t+1} \in H(x_t, z_t)$, we have

$$\begin{aligned} \|x_{t+1} - x_0\|_2^2 - \|x_t - x_0\|_2^2 &= \|x_{t+1} - x_t\|_2^2 + 2\langle x_{t+1} - x_t, x_t - x_0 \rangle \\ &\geq \|x_{t+1} - x_t\|_2^2 \\ &\geq \text{dist}(x_t, H(x_t, z_t))^2. \end{aligned} \quad (7)$$

By [1, Thm. 7.12],

$$\text{dist}(x_t, \mathcal{C}) \leq \rho \text{dist}(x_t, H(x_t, z_t)), \quad (8)$$

where $\rho = (-1/g(\hat{x})) \sup_{t \in \mathbb{N}} \|\nabla g(x_t)\|_2 \|x_t - \hat{x}\|_2$, and $\rho < +\infty$ because $(x_t)_{t \in \mathbb{N}}$ converges [2, Cor. 30.9]. Now, $x^* = \text{proj}(x_0, \mathcal{C})$ so $\mathcal{C} \subset H(x_0, x^*)$. We can assume that $x_t \neq x^*$, so $x_t \notin H(x_0, x^*)$ by (6). By [2, Ex. 29.20],

$$\text{dist}(x_t, H(x_0, x^*)) = \frac{\langle x_t - x^*, x_0 - x^* \rangle}{\|x_0 - x^*\|_2}.$$

Thus,

$$\begin{aligned} \text{dist}(x_t, \mathcal{C}) &\geq \text{dist}(x_t, H(x_0, x^*)) \\ &= \frac{\langle x_t - x^*, x_0 - x^* \rangle}{\|x_0 - x^*\|_2} \\ &= \frac{\langle x_t - x_0, x_0 - x^* \rangle + \|x_0 - x^*\|_2^2}{\|x_0 - x^*\|_2} \\ &\geq \|x_0 - x^*\|_2 - \|x_t - x_0\|_2 \\ &= \|x_0 - x^*\|_2 - \sqrt{\|x^* - x_0\|_2^2 - (\|x^* - x_0\|_2^2 - \|x_t - x_0\|_2^2)} \\ &\geq 0, \end{aligned} \quad (9)$$

where we used the Cauchy-Schwarz inequality in the second inequality and (6) in the last inequality. Let $\varepsilon_s = \|x^* - x_0\|_2^2 - \|x_s - x_0\|_2^2$ for $s \in \{t, t+1\}$. Combining (7)–(9),

$$\varepsilon_{t+1} \leq \varepsilon_t - \frac{1}{\rho^2} \left(\|x_0 - x^*\|_2 - \sqrt{\|x^* - x_0\|_2^2 - \varepsilon_t} \right)^2.$$

Since $\sqrt{\alpha - \beta} \leq \sqrt{\alpha} - \beta/(2\sqrt{\alpha})$ for all $\alpha \geq \beta > 0$, we obtain

$$\begin{aligned} \varepsilon_{t+1} &\leq \varepsilon_t - \frac{1}{\rho^2} \left(\|x_0 - x^*\|_2 - \left(\|x^* - x_0\|_2 - \frac{\varepsilon_t}{2\|x^* - x_0\|_2} \right) \right)^2 \\ &= \left(1 - \frac{\varepsilon_t}{4\rho^2\|x^* - x_0\|_2^2} \right) \varepsilon_t. \end{aligned}$$

Let $\kappa = \max\{8\rho^2, 2\}\|x_0 - x^*\|_2^2$. We have $\varepsilon_t \leq \kappa/(t+2)$ by the induction hypothesis, so

$$\begin{aligned} \varepsilon_{t+1} &\leq \left(1 - \frac{\varepsilon_t}{\kappa/2}\right) \varepsilon_t \\ &\leq \begin{cases} \frac{\kappa/2}{t+2} & \text{if } \varepsilon_t \leq \frac{\kappa/2}{t+2} \\ \left(1 - \frac{1}{t+2}\right) \frac{\kappa}{t+2} & \text{if } \varepsilon_t \geq \frac{\kappa/2}{t+2} \end{cases} \\ &\leq \frac{\kappa}{t+3}. \end{aligned}$$

We conclude that (4) holds for all $t \in \mathbb{N}$. Then, for all $t \in \mathbb{N}$,

$$\begin{aligned} \|x_t - x^*\|_2^2 &= \|x^* - x_0\|_2^2 - \|x_t - x_0\|_2^2 + 2\langle x^* - x_t, x_0 - x_t \rangle \\ &\leq \|x^* - x_0\|_2^2 - \|x_t - x_0\|_2^2, \end{aligned}$$

because $x^* \in H(x_0, x_t)$, since $\mathcal{C} \subset H(x_0, x_t)$ [11, Prop. 5.2]. This proves (5). \square

In our case, $g = \|\cdot\|_p^p - 1$ and $g(0) = -1 < 0$, so Theorem 3.1 holds with

$$\rho = \sup_{t \in \mathbb{N}} \|\nabla g(x_t)\|_2 \|x_t\|_2 = p \sup_{t \in \mathbb{N}} \|x_t\|_{2(p-1)}^{p-1} \|x_t\|_2 < +\infty.$$

Therefore, the complexity of an ε -approximate projection onto the ℓ_p -ball is $\mathcal{O}(n\rho^2\|y - x^*\|_2^2/\varepsilon^2)$.

Remark 3.2. If $p \in [1, +\infty[\cap \mathbb{Q}$, another option, although probably less practical, is to formulate the projection problem as a conic quadratic program and to obtain an ε -approximate solution using an interior-point algorithm, with complexity $\mathcal{O}(\text{poly}(n) \ln(1/\varepsilon))$ [3].

3.4 The nuclear norm-ball

This is probably the most popular example of the computational advantage of linear minimizations over projections in the literature. The nuclear norm, a.k.a. trace norm, of a matrix is the sum of its singular values and serves as a convex surrogate for the rank constraint [16]. The nuclear norm-ball is the convex hull of rank-1 matrices:

$$\{X \in \mathbb{R}^{m \times n} \mid \|X\|_{\text{nuc}} \leq 1\} = \text{conv}\{uv^\top \mid u \in \mathbb{R}^m, v \in \mathbb{R}^n, \|u\|_2 = \|v\|_2 = 1\}.$$

For all $Y \in \mathbb{R}^{m \times n}$,

$$\arg \min_{\|X\|_{\text{nuc}} \leq 1} \|X - Y\|_F = U \text{diag}(\hat{\sigma}) V^\top,$$

where $Y = U \text{diag}(\sigma) V^\top$ is the singular value decomposition (SVD) of Y , $(U, \sigma, V) \in \mathbb{R}^{m \times k} \times \mathbb{R}^k \times \mathbb{R}^{n \times k}$, $k = \min\{m, n\}$, and $\hat{\sigma}$ is the projection of σ onto the standard simplex Δ_k . The SVD can be computed with complexity $\mathcal{O}(mn \min\{m, n\} + \min\{m^3, n^3\})$ using the Golub-Reinsch algorithm or the R -SVD algorithm [22, Fig. 8.6.1]. On the other hand, a linear minimization requires only a truncated SVD:

$$\arg \min_{\|X\|_{\text{nuc}} \leq 1} \langle X, Y \rangle = \arg \max_{\|u\|_2 = \|v\|_2 = 1} \text{tr}((uv^\top)^\top (-Y)) = \arg \max_{\|u\|_2 = \|v\|_2 = 1} u^\top (-Y)v = uv^\top,$$

where u and v are the top left and right singular vectors of $-Y$. A pair of unit vectors $(u, v) \in \mathbb{R}^m \times \mathbb{R}^n$ satisfying $\sigma_1 - u^\top (-Y)v \leq \varepsilon$ with high probability can be obtained using the Lanczos

algorithm with complexity $\mathcal{O}(\nu \ln(m+n) \sqrt{\sigma_1}/\sqrt{\varepsilon})$, where σ_1 and ν denote the top singular value and the number of nonzero entries in $-Y$ respectively [29, 33]. Note that $\nu \leq mn$ and that in many applications of interest, e.g., in recommender systems, $\nu \ll mn$.

In practice, the package ARPACK [40] is often used to compute the top pair of singular vectors. Furthermore, if the input matrix Y is symmetric, then the package LOBPCG [32] can be particularly efficient. Figure 2 illustrates both cases, where linear minimizations are solved to machine precision. The results are averaged over 5 runs and the shaded areas represent ± 1 standard deviation.




Figure 2: Solving a linear minimization and a projection on the nuclear norm-ball. A matrix $Y \in \mathbb{R}^{n \times n}$ is generated by sampling entries from the standard normal distribution. The full and truncated SVDs are computed using the functions `svd` and `svds` from the Python packages `numpy.linalg` [24] and `scipy.sparse.linalg` [50] respectively. The function `svds` is used with `tol=0`. *Left:* The input is Y and the function `svds` is used with `solver='arpack'`. *Right:* The input is the symmetric matrix $(Y + Y^\top)/2$ and the function `svds` is used with `solver='lobpcg'`. We see that the ratio of CPU times increases as n increases.

3.5 The flow polytope

Let G be a single-source single-sink directed acyclic graph (DAG) with m vertices and n edges. Index by $\llbracket 1, m \rrbracket$ the set of vertices such that the edges are directed from a smaller to a larger vertex index; this can be achieved with complexity $\mathcal{O}(m+n)$ via topological sort [13, Sec. 22.4]. Let $A_G \in \mathbb{R}^{m \times n}$ be the incidence matrix of G . The flow polytope induced by G is

$$\mathcal{F}_G = \{x \in \mathbb{R}^n \mid A_G x = (-1, 0, \dots, 0, 1)^\top, x \geq 0\},$$

i.e., \mathcal{F}_G is the set of unit flows $x \in \mathbb{R}^n$ on G , where $[x]_i \geq 0$ denotes the flow going through edge i . Thus, for all $y \in \mathbb{R}^n$, $\arg \min_{x \in \mathcal{F}_G} \langle x, y \rangle$ is a flow $x \in \{0, 1\}^n$ identifying a shortest path on G weighted by y . Its computation has complexity $\mathcal{O}(m+n)$ [13, Sec. 24.2]. This is significantly cheaper than the complexity $\tilde{\mathcal{O}}(m^3 n + n^2)$ of a projection [49, Thm. 20], where $\tilde{\mathcal{O}}$ hides polylogarithmic factors.

3.6 The Birkhoff polytope

The Birkhoff polytope, a.k.a. assignment polytope, is the set of doubly stochastic matrices

$$\mathcal{B}_n = \{X \in \mathbb{R}^{n \times n} \mid X \mathbf{1} = \mathbf{1}, X^\top \mathbf{1} = \mathbf{1}, X \geq 0\}.$$

It is the convex hull of the permutation matrices and arises in matching, ranking, and seriation problems. Linear minimizations can be solved with complexity $\mathcal{O}(n^3)$ using the Hungarian algorithm [34]. To the best of our knowledge, there is no projection method specific to the Birkhoff

polytope so we propose one here. Let $Y \in \mathbb{R}^{n \times n}$. By reshaping it into a vector $y \in \mathbb{R}^{n^2}$, projecting Y onto the Birkhoff polytope is equivalent to solving

$$\begin{aligned} \min_{x \in \mathbb{R}^{n^2}} \|x - y\|_2^2 \\ \text{s.t. } Ax = 1 \\ x \geq 0, \end{aligned} \quad \text{where } A = \begin{pmatrix} 1^\top & 0^\top & \cdots & 0^\top \\ 0^\top & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0^\top \\ 0^\top & \cdots & 0^\top & 1^\top \\ I_n & \cdots & \cdots & I_n \end{pmatrix} \in \mathbb{R}^{2n \times n^2}. \quad (10)$$

This can again be reformulated as

$$\min_{x \in \mathbb{R}^{n^2}} \left(\iota_{\mathcal{K}}(x) + \frac{1}{2} \|x - y\|_2^2 \right) + \left(\iota_{\mathcal{A}}(x) + \frac{1}{2} \|x - y\|_2^2 \right), \quad (11)$$

where $\mathcal{K} = \{x \in \mathbb{R}^{n^2} \mid x \geq 0\}$ and $\mathcal{A} = \{x \in \mathbb{R}^{n^2} \mid Ax = 1\}$. That is, we split the constraints into two sets enjoying efficient projections. We can now apply the Douglas-Rachford algorithm [42] to problem (11). The method is presented in Algorithm 3; see Appendix A for details. Line 3 computes the projection of u_t onto the affine subspace \mathcal{A} and in Line 4 is computed a projection onto the nonnegative orthant \mathcal{K} . We can set $u = 1/n \in \mathcal{A}$ and we denote by $A^\dagger \in \mathbb{R}^{n^2 \times 2n}$ the Moore-Penrose inverse of A .

Algorithm 3 Douglas-Rachford for problem (11)

Input: Point to project $y \in \mathbb{R}^{n^2}$, start point $z_0 \in \mathbb{R}^{n^2}$, offset point $u \in \mathcal{A}$.

- 1: **for** $t = 0$ **to** $T - 1$ **do**
 - 2: $u_t \leftarrow \frac{z_t + y}{2}$
 - 3: $x_t \leftarrow u_t - A^\dagger A(u_t - u)$
 - 4: $z_{t+1} \leftarrow \max \left\{ \frac{2x_t - z_t + y}{2}, 0 \right\} + z_t - x_t$
 - 5: **end for**
-

The complexity of an iteration of Algorithm 3 is dominated by the matrix-vector multiplication in Line 3. We can assume that $A^\dagger A \in \mathbb{R}^{n^2 \times n^2}$ is precomputed. In fact,

$$A^\dagger A = \frac{1}{n^2} \begin{pmatrix} B_1 & B_2 & \cdots & B_2 \\ B_2 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & B_2 \\ B_2 & \cdots & B_2 & B_1 \end{pmatrix} \in \mathbb{R}^{n^2 \times n^2}, \quad \text{where } \begin{cases} B_1 = nI_n + (n-1)J_n \in \mathbb{R}^{n \times n} \\ B_2 = nI_n - J_n \in \mathbb{R}^{n \times n}, \end{cases}$$

so $A^\dagger A$ is block circulant with circulant blocks (BCCB) and has only three distinct entries: $2n - 1$, $n - 1$, and -1 . The expression of $A^\dagger A$ can be shown by checking that $A^\dagger = A^\top / n - J_{n^2, 2n} / (2n^2) \in \mathbb{R}^{n^2 \times 2n}$ using the necessary and sufficient Moore-Penrose conditions [47, Thm. 1]. Thus, the multiplication of $A^\dagger A$ and any vector $x \in \mathbb{R}^{n^2}$ can be performed with complexity $\mathcal{O}(n^2)$. Indeed, it amounts to computing $(nI_n + (n-1)J_n)[x]_{i;j}$ and $(nI_n - J_n)[x]_{i;j}$ for every $(i, j) \in \{(kn+1, (k+1)n) \mid k \in \llbracket 0, n-1 \rrbracket\}$, each of which has complexity $\mathcal{O}(n)$.

It remains to bound the number of iterations required to achieve ε -convergence. Let $x^* = \text{proj}(y, \mathcal{A} \cap \mathcal{K})$ be the solution to problem (11), i.e., the projection of Y onto the Birkhoff polytope after reshaping, and let $\bar{x}_t = (\sum_{s=0}^t x_s) / (t+1) \in \mathcal{A}$ for all $t \in \mathbb{N}$. By [14, Thm. 1],

$$\|\bar{x}_t - x^*\|_2 \leq \frac{\|z_0 - z^*\|_2}{\sqrt{2(t+1)}},$$

where z^* is a fixed point of $\text{rprox}_{\iota_{\mathcal{K}}+(1/2)\|\cdot-y\|_2^2} \circ \text{rprox}_{\iota_{\mathcal{A}}+(1/2)\|\cdot-y\|_2^2}$, $\text{rprox}_{\varphi} = 2\text{prox}_{\varphi} - \text{id}$, and prox_{φ} is the proximity operator of φ [44]. Therefore, the complexity of an ε -approximate projection onto the Birkhoff polytope is $\mathcal{O}(n^2 d_z^2 / \varepsilon^2)$, where

$$d_z = \|z_0 - z^*\|_2. \quad (12)$$

3.7 The permutahedron

Let \mathfrak{S}_n be the set of permutations on $\llbracket 1, n \rrbracket$ and $w \in \mathbb{R}^n$. The permutahedron induced by w is the convex hull of all permutations of the entries in w , i.e.,

$$\mathcal{P}_w = \text{conv}\{w_{\sigma} \in \mathbb{R}^n \mid w_{\sigma} = ([w]_{\sigma_1}, \dots, [w]_{\sigma_n})^{\top}, \sigma \in \mathfrak{S}_n\}.$$

It is related to the Birkhoff polytope via $\mathcal{P}_w = \{Xw \mid X \in \mathcal{B}_n\}$. With no loss of generality, we can assume that the weights are already sorted in ascending order: $[w]_1 \leq \dots \leq [w]_n$. Thus, for all $y \in \mathbb{R}^n$,

$$\arg \min_{x \in \mathcal{P}_w} \langle x, y \rangle = w_{\sigma^{-1}},$$

where σ satisfies $[y]_{\sigma_1} \geq \dots \geq [y]_{\sigma_n}$. Sorting the entries of y has complexity $\mathcal{O}(n \ln(n))$ [13]. A projection can be obtained with a slightly higher complexity $\mathcal{O}(n \ln(n) + n)$, by sorting the entries of y and solving an isotonic regression problem [45].

Acknowledgment

Research reported in this paper was partially supported by the Research Campus MODAL funded by the German Federal Ministry of Education and Research under grant 05M14ZAM.

A An application of the Douglas-Rachford algorithm

Let \mathcal{H} be a Euclidean space with norm $\|\cdot\|$ and denote by $\Gamma_0(\mathcal{H})$ the set of proper lower semicontinuous convex functions $\mathcal{H} \rightarrow \mathbb{R} \cup \{+\infty\}$. The Douglas-Rachford algorithm [42] can be used to solve

$$\min_{x \in \mathcal{H}} f(x) + g(x),$$

when $f, g \in \Gamma_0(\mathcal{H})$ satisfy $\arg \min_{\mathcal{H}}(f + g) \neq \emptyset$ and $(\text{ri dom } f) \cap (\text{ri dom } g) \neq \emptyset$ [2], where ri and dom denote the relative interior of a set and the domain of a function respectively. It is presented in Algorithm 4. For every function $\varphi \in \Gamma_0(\mathcal{H})$, the proximity operator is $\text{prox}_{\varphi} = \arg \min_{x \in \mathcal{H}} \varphi(x) + (1/2)\|\cdot - x\|^2$ [44].

Algorithm 4 Douglas-Rachford

Input: Start point $z_0 \in \mathcal{H}$.

- 1: **for** $t = 0$ **to** $T - 1$ **do**
 - 2: $x_t \leftarrow \text{prox}_g(z_t)$
 - 3: $z_{t+1} \leftarrow \text{prox}_f(2x_t - z_t) + z_t - x_t$
 - 4: **end for**
-

We are interested in an application to problem (11), where $\mathcal{H} = \mathbb{R}^{n^2}$, $f = \iota_{\mathcal{K}} + (1/2)\|\cdot - y\|_2^2$, $g = \iota_{\mathcal{A}} + (1/2)\|\cdot - y\|_2^2$, $y \in \mathcal{H}$, $\mathcal{K} = \{x \in \mathbb{R}^{n^2} \mid x \geq 0\}$, $\mathcal{A} = \{x \in \mathbb{R}^{n^2} \mid Ax = 1\}$, and $A \in \mathbb{R}^{2n \times n^2}$ is

defined in (10). Problem (11) admits a (unique) solution since it is a projection problem onto the intersection of the closed convex sets \mathcal{K} and \mathcal{A} , and $1/n \in (\text{ri dom } f) \cap (\text{ri dom } g) = (\text{ri } \mathcal{K}) \cap \mathcal{A}$ so the Douglas-Rachford algorithm is well defined here. We now show that it reduces to Algorithm 3. For all $t \in \mathbb{N}$,

$$\begin{aligned}
x_t &= \text{prox}_g(z_t) \\
&= \arg \min_{x \in \mathcal{H}} g(x) + \frac{1}{2} \|z_t - x\|_2^2 \\
&= \arg \min_{x \in \mathcal{A}} \frac{1}{2} \|x - y\|_2^2 + \frac{1}{2} \|x - z_t\|_2^2 \\
&= \arg \min_{x \in \mathcal{A}} \left\| x - \frac{z_t + y}{2} \right\|_2^2 \\
&= \text{proj} \left(\frac{z_t + y}{2}, \mathcal{A} \right) \\
&= \frac{z_t + y}{2} - A^\dagger A \left(\frac{z_t + y}{2} - u \right),
\end{aligned}$$

since $\text{proj}(\cdot, \mathcal{A}) = \cdot - A^\dagger A(\cdot - u)$, given any $u \in \mathcal{A}$ [2, Ex. 29.17]. Thus, Lines 2–3 in Algorithm 3 are equivalent to Line 2 in Algorithm 4. Similarly, Line 4 in Algorithm 3 is equivalent to Line 3 in Algorithm 4:

$$\begin{aligned}
z_{t+1} &= \text{prox}_f(2x_t - z_t) + z_t - x_t \\
&= \arg \min_{z \in \mathcal{H}} \left(f(z) + \frac{1}{2} \|2x_t - z_t - z\|_2^2 \right) + z_t - x_t \\
&= \arg \min_{z \in \mathcal{K}} \left(\frac{1}{2} \|z - y\|_2^2 + \frac{1}{2} \|z - (2x_t - z_t)\|_2^2 \right) + z_t - x_t \\
&= \arg \min_{z \in \mathcal{K}} \left\| z - \frac{2x_t - z_t + y}{2} \right\|_2^2 + z_t - x_t \\
&= \text{proj} \left(\frac{2x_t - z_t + y}{2}, \mathcal{K} \right) + z_t - x_t \\
&= \max \left\{ \frac{2x_t - z_t + y}{2}, 0 \right\} + z_t - x_t,
\end{aligned}$$

since $\text{proj}(\cdot, \mathcal{K}) = \max\{\cdot, 0\}$.

References

- [1] H. H. Bauschke and J. M. Borwein. On projection algorithms for solving convex feasibility problems. *SIAM Review*, 38(3):367–426, 1996.
- [2] H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer, 2nd edition, 2017.
- [3] A. Ben-Tal and A. S. Nemirovski. *Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications*. Society for Industrial and Applied Mathematics, 2001.
- [4] G. Braun, S. Pokutta, and D. Zink. Lazifying conditional gradient algorithms. *Journal of Machine Learning Research*, 20(71):1–42, 2019.

- [5] M. D. Canon and C. D. Cullum. A tight upper bound on the rate of convergence of Frank-Wolfe algorithm. *SIAM Journal on Control*, 6(4):509–516, 1968.
- [6] J. Chen, D. Zhou, J. Yi, and Q. Gu. A Frank-Wolfe framework for efficient and effective adversarial attacks. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*, pages 3486–3494, 2020.
- [7] K. L. Clarkson. Coresets, sparse greedy approximation, and the Frank-Wolfe algorithm. *ACM Transactions on Algorithms*, 6(4):1–30, 2010.
- [8] C. W. Combettes and S. Pokutta. Revisiting the approximate Carathéodory problem via the Frank-Wolfe algorithm. *arXiv preprint arXiv:1911.04415*, 2019.
- [9] C. W. Combettes and S. Pokutta. Boosting Frank-Wolfe by chasing gradients. In *Proceedings of the 37th International Conference on Machine Learning*, pages 2111–2121, 2020.
- [10] C. W. Combettes, C. Spiegel, and S. Pokutta. Projection-free adaptive gradients for large-scale optimization. *arXiv preprint arXiv:2009.14114*, 2020.
- [11] P. L. Combettes. Strong convergence of block-iterative outer approximation methods for convex optimization. *SIAM Journal on Control and Optimization*, 38(2):538–565, 2000.
- [12] L. Condat. Fast projection onto the simplex and the ℓ_1 ball. *Mathematical Programming*, 158(1):575–585, 2016.
- [13] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, 3rd edition, 2009.
- [14] D. Davis and W. Yin. Faster convergence rates of relaxed Peaceman-Rachford and ADMM under regularity assumptions. *Mathematics of Operations Research*, 42(3):783–805, 2017.
- [15] J. C. Dunn and S. Harshbarger. Conditional gradient algorithms with open loop step size rules. *Journal of Mathematical Analysis and Applications*, 62(2):432–444, 1978.
- [16] M. Fazel, H. Hindi, and S. P. Boyd. A rank minimization heuristic with application to minimum order system approximation. In *Proceedings of the 2001 American Control Conference*, pages 4734–4739, 2001.
- [17] M. Frank and P. Wolfe. An algorithm for quadratic programming. *Naval Research Logistics Quarterly*, 3(1-2):95–110, 1956.
- [18] R. M. Freund, P. Grigas, and R. Mazumder. An extended Frank-Wolfe method with “in-face” directions, and its application to low-rank matrix completion. *SIAM Journal on Optimization*, 27(1):319–346, 2017.
- [19] D. Garber. *Projection-free Algorithms for Convex Optimization and Online Learning*. Ph.D. thesis, Technion, 2016.
- [20] D. Garber and E. Hazan. Faster rates for the Frank-Wolfe method over strongly-convex sets. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 541–549, 2015.
- [21] D. Garber and O. Meshi. Linear-memory and decomposition-invariant linearly convergent conditional gradient algorithm for structured polytopes. In *Advances in Neural Information Processing Systems*, volume 29, pages 1001–1009, 2016.
- [22] G. H. Golub and C. F. van Loan. *Matrix Computations*. Johns Hopkins University Press, 4th edition, 2013.
- [23] J. Guélat and P. Marcotte. Some comments on Wolfe’s ‘away step’. *Mathematical Programming*, 35(1):110–119, 1986.
- [24] C. R. Harris et al. Array programming with NumPy. *Nature*, 585(7825):357–362, 2020.

- [25] Y. Haugazeau. *Sur les Inéquations Variationnelles et la Minimisation de Fonctionnelles Convexes*. Thèse de doctorat, Université de Paris, 1968.
- [26] E. Hazan. Sparse approximate solutions to semidefinite programs. In *Proceedings of the 8th Latin American Symposium on Theoretical Informatics*, pages 306–316, 2008.
- [27] E. Hazan and S. Kale. Projection-free online learning. In *Proceedings of the 29th International Conference on Machine Learning*, 2012.
- [28] D. W. Hearn. The gap function of a convex program. *Operations Research Letters*, 1(2):67–71, 1982.
- [29] M. Jaggi. Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In *Proceedings of the 30th International Conference on Machine Learning*, pages 427–435, 2013.
- [30] A. Joulin, K. Tang, and L. Fei-Fei. Efficient image and video co-localization with Frank-Wolfe algorithm. In *European Conference on Computer Vision*, pages 253–268, 2014.
- [31] T. Kerdreux, A. d’Aspremont, and S. Pokutta. Projection-free optimization on uniformly convex sets. In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics*, pages 19–27, 2021.
- [32] A. V. Knyazev, M. E. Argentati, I. Lashuk, and E. E. Ovtchinnikov. Block locally optimal preconditioned eigenvalue solvers (BLOPEX) in hypre and PETSc. *SIAM Journal on Scientific Computing*, 29(5):2224–2239, 2007.
- [33] J. Kuczyński and H. Woźniakowski. Estimating the largest eigenvalue by the power and Lanczos algorithms with a random start. *SIAM Journal on Matrix Analysis and Applications*, 13(4):1094–1122, 1992.
- [34] H. W. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1–2):83–97, 1955.
- [35] S. Lacoste-Julien. Convergence rate of Frank-Wolfe for non-convex objectives. *arXiv preprint arXiv:1607.00345*, 2016.
- [36] S. Lacoste-Julien and M. Jaggi. On the global linear convergence of Frank-Wolfe optimization variants. In *Advances in Neural Information Processing Systems*, volume 28, pages 496–504, 2015.
- [37] S. Lacoste-Julien, M. Jaggi, M. Schmidt, and P. Pletscher. Block-coordinate Frank-Wolfe optimization for structural SVMs. In *Proceedings of the 30th International Conference on Machine Learning*, pages 53–61, 2013.
- [38] G. Lan and Y. Zhou. Conditional gradient sliding for convex optimization. *SIAM Journal on Optimization*, 26(2):1379–1409, 2016.
- [39] L. J. LeBlanc, E. K. Morlok, and W. P. Pierskalla. An efficient approach to solving the road network equilibrium traffic assignment problem. *Transportation Research*, 9(5):309–318, 1975.
- [40] R. B. Lehoucq, D. C. Sorensen, and C. Yang. *ARPACK Users’ Guide: Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods*. Society for Industrial and Applied Mathematics, 1998.
- [41] E. S. Levitin and B. T. Polyak. Constrained minimization methods. *USSR Computational Mathematics and Mathematical Physics*, 6(5):1–50, 1966.
- [42] P.-L. Lions and B. Mercier. Splitting algorithms for the sum of two nonlinear operators. *SIAM Journal on Numerical Analysis*, 16(6):964–979, 1979.
- [43] G. Luise, S. Salzo, M. Pontil, and C. Ciliberto. Sinkhorn barycenters with free support via Frank-Wolfe algorithm. In *Advances in Neural Information Processing Systems*, volume 32, pages 9322–9333, 2019.

- [44] J. J. Moreau. Fonctions convexes duales et points proximaux dans un espace hilbertien. *Comptes Rendus Hebdomadaires des Séances de l'Académie des Sciences*, 255:2897–2899, 1962.
- [45] R. Negrinho and A. F. T. Martins. Orbit regularization. In *Advances in Neural Information Processing Systems*, volume 27, pages 3221–3229, 2014.
- [46] C. H. J. Pang. First order constrained optimization algorithms with feasibility updates. *arXiv preprint arXiv:1506.08247*, 2015.
- [47] R. Penrose. A generalized inverse for matrices. *Mathematical Proceedings of the Cambridge Philosophical Society*, 51(3):406–413, 1955.
- [48] S. J. Reddi, S. Sra, B. Póczos, and A. Smola. Stochastic Frank-Wolfe methods for nonconvex optimization. In *54th Annual Allerton Conference on Communication, Control, and Computing*, pages 1244–1251, 2016.
- [49] L. A. Végh. A strongly polynomial algorithm for a class of minimum-cost flow problems with separable convex objectives. *SIAM Journal on Computing*, 45(5):1729–1761, 2016.
- [50] P. Virtanen et al. Scipy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, 17(3):261–272, 2020.
- [51] J. Xie, Z. Shen, C. Zhang, H. Qian, and B. Wang. Efficient projection-free online methods with stochastic recursive gradient. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*, pages 6446–6453, 2020.
- [52] A. Yurtsever, S. Sra, and V. Cevher. Conditional gradient methods via stochastic path-integrated differential estimator. In *Proceedings of the 36th International Conference on Machine Learning*, pages 7282–7291, 2019.
- [53] M. Zhang, Z. Shen, A. Mokhtari, H. Hassani, and A. Karbasi. One sample stochastic Frank-Wolfe. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, pages 4012–4023, 2020.