

A MORE ON FRANK-WOLFE-TYPE ALGORITHMS

A.1 STANDARD FRANK-WOLFE ALGORITHM

Algorithm 2 Frank-Wolfe algorithm (Frank et al., 1956)

Input: obj. $f : \mathcal{Y} \mapsto \mathbb{R}$, oracle $O(\cdot)$, init. $w_0 \in \mathcal{Y}$

- 1: **for** $t=1, 2, 3 \dots, T$ **do**
- 2: $v_t \leftarrow \text{Oracle}(\nabla f(w_{t-1})) = \arg \min_{v \in \mathcal{Y}} v^T \nabla f(w_{t-1})$
- 3: $w_t \leftarrow (1 - \eta_t w_{t-1}) + \eta_t v_t$, for $\eta := \frac{2}{t+2}$
- 4: **end for**
- 5: **return** w_T

For a convex function $f : \mathcal{X} \mapsto \mathbb{R}$ the Frank-Wolfe algorithm (FW) solves the constrained optimization problem over a compact and convex set \mathcal{X} . The standard FW is known to have sublinear convergence rate, and various methods are proposed to improve the convergence rate. For example, when the underlying feasible set is a polytope, and the objective function is strongly convex, multiple variants, such as away-step FW (Wolfe, 1970; Jaggi, 2013), pairwise FW (Mitchell et al., 1974), and Wolfe’s method (Wolfe, 1976) are shown to enjoy linear convergence rate.

A.2 WOLFE’S METHOD FOR MINIMUM NORM POINT

Algorithm 3 Wolfe’s Method for Minimum Norm Point

Initialize $x \in \mathcal{P}$, active set $\mathcal{S} = [x]$ and weight $\lambda = [1]$.

Output: $x \in \mathcal{P}$ that has the minimum Euclidean norm.

- 1: **while** true **do** // Major cycle
- 2: $s \leftarrow \text{Oracle}(x)$ // Potential improving point
- 3: **if** $\|x\|^2 \leq x^T s + \epsilon$ **then break**
- 4: $\mathcal{S} \leftarrow \mathcal{S} \cup \{s\}$
- 5: **while** true **do** // Minor cycle
- 6: $y, \alpha \leftarrow \text{AffineMinimizer}(\mathcal{S})$ // $y = \arg \min_{s \in \text{aff}(\mathcal{S})} \|s\|_2$
- 7: **if** $\alpha_s > 0$ for all s **then break** // $y \in \text{conv}(\mathcal{S})$
- 8: // If $y \notin \text{conv}(\mathcal{S})$, then update y to the intersection of $\text{conv}(\mathcal{S})$ and segment joining x and y . Then remove points in \mathcal{S} unnecessary for describing y .
- 9: $\theta \leftarrow \min_{i: \alpha_i \leq 0} \frac{\lambda_i}{\lambda_i - \alpha_i}$ // Recall λ satisfies $x = \sum_{s \in \mathcal{S}} \lambda_s s$
- 10: $y \leftarrow \theta y + (1 - \theta)x, \lambda_i = \theta \alpha_i + (1 - \theta)\lambda_i$
- 11: $\mathcal{S} \leftarrow \{s_i | s_i \in \mathcal{S} \text{ and } \lambda_i > 0\}$
- 12: **end while**
- 13: Update $x = y$ and $\lambda = \alpha$.
- 14: **end while**
- 15: **return** x

Wolfe’s method is an iterative algorithm for finding the point with minimum Euclidean norm in a polytope, which is defined as the convex hull of a set of finitely many points.

The Wolfe’s method consists of a finite number of major cycles, each of which consists of a finite number of minor cycles. At the start of each major cycle, let $H(x) := \{y^T x = x^x\}$ be the hyperplane defined by x . If $H(x)$ separates the polytope from the origin, then the major cycle is terminated. Otherwise, we invoke an oracle to find any point that is on the near side of the hyperplane. The point is then added into the active set \mathcal{S} , and minor cycle starts.

In a minor cycle, let y be the point of smallest norm in of the affine hull $\text{aff}(\mathcal{S})$. If y is in the relative interior of the convex hull $\text{conv}(\mathcal{S})$, the x is updated to y and the minor cycle is terminated. Otherwise, y is updated to the nearest point to y on the line segment $\text{conv}(\mathcal{S}) \cap [x, y]$. Thus y is updated to a boundary point of $\text{conv}(\mathcal{S})$, and any point that is not on the face of $\text{conv}(\mathcal{S})$ in which y lies are deleted. The minor cycles are executed repeatedly until the \mathcal{S} becomes a *corral*, that is, a set whose affine minimizer lies inside its convex hull. Since a set of one point is always a corral, the minor cycles are terminated after a finite number of runs.

B PROOF OF THEOREM 4.1

Theorem 4.1 (Approximation Error Strictly Decreases). *For any non-terminal step t , we have $\text{err}(\mu^{t+1}) < \text{err}(\mu^t)$. That is, the measurement vector of μ^t found by Algorithm 1 gets strictly closer to the convex set Ω after major cycle step.*

Proof. If the current step is a major cycle with no minor cycle, then \mathbf{x}^{t+1} is the affine minimizer of $\text{aff}(\mathcal{S} \cup \{\mathbf{s}^t\})$ with respect to ω^t . Then the affine minimizer property implies $(\mathbf{s}^t - \mathbf{x}^{t+1})(\mathbf{x}^{t+1} - \omega^t) = 0$. Since iteration does not terminate at step t , we have $(\mathbf{x}^t - \omega^t)^T(\mathbf{x}^t - \mathbf{s}^t) > 0$, and therefore \mathbf{x}^{t+1} not equal to \mathbf{x}^t . Then \mathbf{x}^{t+1} is the unique affine minimizer implies $f_\Omega(\mathbf{x}^{t+1}) = \min_{\omega \in \Omega} \|\mathbf{x}^{t+1} - \omega\|^2 \leq \|\mathbf{x}^{t+1} - \omega^t\|^2 < \|\mathbf{x}^t - \omega^t\|^2 = f_\Omega(\mathbf{x}^t)$.

Otherwise the current step contains one or more minor cycles. In this case, we show that the first minor cycle strictly reduces the approximation error, and the (possibly) following minor cycles cannot increase it. For the first minor cycle, the affine minimizer \mathbf{y}^0 of $\text{aff}(\mathcal{S} \cup \{\mathbf{s}^t\})$ with respect to ω^t is outside $\text{conv}(\mathcal{S} \cup \{\mathbf{s}^t\})$. Let $\mathbf{z} = \theta \mathbf{y}^0 + (1 - \theta) \mathbf{x}^t$ be the intersection of $\text{conv}(\mathcal{S} \cup \{\mathbf{s}^t\})$ and segment joining \mathbf{x} and \mathbf{y} . Let $\mathcal{V}^0 := \mathcal{S}^t$ and \mathcal{V}^i denotes the active set after the i -th minor cycle. Then since \mathbf{y}^1 is the affine minimizer of \mathcal{V}^1 with respect to ω^t , we have

$$\|\mathbf{z} - \omega^t\| = \|\theta \mathbf{y}^0 + (1 - \theta) \mathbf{x}^t - \omega^t\| \leq \theta \|\mathbf{y}^0 - \omega^t\| + (1 - \theta) \|\mathbf{x}^t - \omega^t\| < \|\mathbf{x}^t - \omega^t\|, \quad (14)$$

where the second step uses the triangle inequality and the last step follows since the segment $\mathbf{x}^t \mathbf{y}^0$ intersects the interior of $\text{conv}(\mathcal{S} \cup \{\mathbf{s}^t\})$, and the distance to ω^t strictly decreases along this segment. Therefore the point \mathbf{z} found by first minor cycle satisfies

$$f_\Omega(\mathbf{z}) = \min_{\omega \in \Omega} \|\mathbf{z} - \omega\|^2 \leq \|\mathbf{z} - \omega^t\|^2 < \|\mathbf{x}^t - \omega^t\|^2 = f_\Omega(\mathbf{x}^t). \quad (15)$$

Hence $h(\mathbf{y}^1) < h(\mathbf{x}^t)$, and the first minor cycle strictly decreases the approximation error. By a similar argument, in subsequent minor cycles the approximation error cannot be increased. However, after the first minor cycle, the iterating point may already at the intersection point and the strict inequality in last step of Eq. 14 need to be replaced by non-strict inequality.

Therefore any major cycle either finds an improving point and continue, or enter minor cycles where the first minor cycle finds an improving point, and the subsequent minor cycles does not increase the distance. Adding both side of $f_\Omega(\mathbf{x}^{t+1}) < f_\Omega(\mathbf{x}^t)$ by $f_\Omega(\mathbf{x}^*)$ and we have the approximation error $h(\mathbf{x}^{t+1}) < h(\mathbf{x}^t)$ strictly decreases. \square

C PROOF OF THEOREM 4.2

We first prove the Theorem 4.2 using Lemma 4.3 and Lemma 4.4. Then present the proof of the lemmas.

Theorem 4.2 (Convergence in Approximation Error). *For $t \geq 1$, the mixed policy μ^t found by Algorithm 1 satisfies*

$$\text{err}(\mu^t) \leq 16Q^2/(t+2). \quad (16)$$

where $Q := \max_{\mu \in \Delta(\mathcal{U})} \|\mathbf{c}(\mu)\|$ is the maximum norm of a measurement vector.

Proof. Since Lemma 4.4 shows that drop steps are no more than half of total major cycle steps, and Theorem 4.1 guarantees these drop steps reducing the approximation error, we can safely skip these step, and re-index the step numbers to include non-drop steps only using k .

For these non-drop steps, we claim that $\text{err}(\mu^k) \leq 8Q^2/(k+1)$. Using Lemma 4.3, we prove the convergence rate using induction. We first bound the error of any $\text{err}(\mu^k)$. For any $k \geq 1$

$$\text{err}(\mu^k) = \text{dist}(\mathbf{c}(\mu^k), \Omega) - \text{dist}(\mathbf{c}(\mu^*), \Omega) \quad (17)$$

$$= 1/2 \|\mathbf{c}(\mu^k) - \text{Proj}_\Omega(\mathbf{c}(\mu^k))\|^2 - 1/2 \|\mathbf{c}(\mu^*) - \text{Proj}_\Omega(\mathbf{c}(\mu^*))\|^2 \quad (18)$$

$$\leq 1/2 (\|\mathbf{c}(\mu^k)\|^2 + \|\text{Proj}_\Omega(\mathbf{c}(\mu^k))\|^2 - \|\mathbf{c}(\mu^*)\|^2 - \|\text{Proj}_\Omega(\mathbf{c}(\mu^*))\|^2) \quad (19)$$

$$\leq \|\mathbf{c}(\mu^k)\|^2 - \|\mathbf{c}(\mu^*)\|^2 \quad (20)$$

$$\leq \|\mathbf{c}(\mu^k)\|^2 \quad (21)$$

$$\leq Q^2, \quad (22)$$

where Eq. 18 uses the definition of our squared Euclidean distance function. Eq. 19 follows from triangle inequality, and Eq. 20 is by the contractive property of the Euclidean distance.

When $k = 1$, the Eq. 22 established the based case. Now for $k \geq 1$, assume that $\text{err}(\mu^k) \leq 8Q^2/(k+1)$ for $k \geq 1$, then Lemma 4.3 gives $\text{err}(\mu^{k+1}) \leq \text{err}(\mu^k) - \text{err}^2(\mu^k)/8Q^2$. Since the quadratic function of the right hand side is monotonically increasing on $(-\infty, 4Q^2]$, using the inductive hypothesis

$$\text{err}(\mu^{k+1}) \leq \text{err}(\mu^k) - \text{err}^2(\mu^k)/8Q^2 \leq 8Q^2/(k+1) - 8Q^2/(k+1)^2 \leq Q^2/(k+2) \quad (23)$$

Then since for t steps of major cycle steps, the number of non-drop steps $k > t/2$, we conclude that $\text{err}(\mu^t) \leq 16Q^2/(t+2)$. □

Then we prove the lemmas.

Lemma 4.3. *For a non-drop step, we have $\text{err}(\mu^t) - \text{err}(\mu^{t+1}) \geq \text{err}^2(\mu^t)/8Q^2$.*

Proof. The non-drop step contains either no minor cycle or one minor cycle. We first consider the no minor cycle case.

If a major cycle contains no minor cycle, then \mathbf{x}^{t+1} is the affine minimizer of the $\mathcal{S} \cup \{\mathbf{s}^t\}$.

$$\text{err}(\mu^t) - \text{err}(\mu^{t+1}) = \text{dist}(\mathbf{x}^t, \Omega) - \text{dist}(\mathbf{x}^{t+1}, \Omega) \quad (24)$$

$$= 1/2(\|\mathbf{x}^t - \boldsymbol{\omega}^t\|^2 - \min_{\boldsymbol{\omega} \in \Omega} \|\mathbf{x}^{t+1} - \boldsymbol{\omega}\|^2) \quad (25)$$

$$\geq 1/2(\|\mathbf{x}^t - \boldsymbol{\omega}^t\|^2 - \|\mathbf{x}^{t+1} - \boldsymbol{\omega}^t\|^2) \quad (26)$$

$$= 1/2(\|\mathbf{x}^t - \boldsymbol{\omega}^t\|^2 + \|\mathbf{x}^{t+1} - \boldsymbol{\omega}^t\|^2 - 2\|\mathbf{x}^{t+1} - \boldsymbol{\omega}^t\|^2) \quad (27)$$

$$= 1/2(\|\mathbf{x}^t - \boldsymbol{\omega}^t\|^2 + \|\mathbf{x}^{t+1} - \boldsymbol{\omega}^t\|^2 - 2(\mathbf{x}^t - \boldsymbol{\omega}^t)^T(\mathbf{x}^{t+1} - \boldsymbol{\omega}^t)) \quad (28)$$

$$= 1/2(\|\mathbf{x}^t - \mathbf{x}^{t+1}\|^2), \quad (29)$$

where the equation (28) follows from the affine minimizer property Eq. (9). For $\|\mathbf{x}^t - \mathbf{x}^{t+1}\|$ in the last equation, and $\forall \mathbf{q} \in \text{aff}(\mathcal{S} \cup \{\mathbf{s}^t\})$, we have

$$\|\mathbf{x}^t - \mathbf{x}^{t+1}\| \geq \|\mathbf{x}^t - \mathbf{x}^{t+1}\| \frac{\|\mathbf{x}^t\| + \|\mathbf{q}\|}{2Q} \quad (\text{Definition of } Q) \quad (30)$$

$$\geq \|\mathbf{x}^t - \mathbf{x}^{t+1}\| \frac{\|\mathbf{x}^t - \mathbf{q}\|}{2Q} \quad (\text{Triangle inequality}) \quad (31)$$

$$\geq \frac{1}{2Q}(\mathbf{x}^t - \mathbf{x}^{t+1})(\mathbf{x}^t - \mathbf{q}) \quad (\text{Cauchy-Schwarz inequality}) \quad (32)$$

$$= \frac{1}{2Q}(\mathbf{x}^t - \boldsymbol{\omega}^t)(\mathbf{x}^t - \mathbf{q}) \quad (\text{Affine minimizer property}). \quad (33)$$

Then it suffices to show that $(\mathbf{x}^t - \boldsymbol{\omega}^t)(\mathbf{x}^t - \mathbf{q}) \geq \text{err}(\mu^t)$.

Since Ω is a convex set, the squared Euclidean distance function $\text{dist}(\mathbf{x}, \Omega)$ is convex for \mathbf{x} , which implies

$$\text{dist}(\mathbf{x}^t, \Omega) + (\mathbf{q} - \mathbf{x}^t) \nabla \text{dist}(\mathbf{x}^t, \Omega) \leq \text{dist}(\mathbf{q}, \Omega). \quad (34)$$

Putting in $\nabla \text{dist}(\mathbf{x}^t, \Omega) = (\mathbf{x}^t - \text{Proj}_{\Omega}(\mathbf{x}^t)) = (\mathbf{x}^t - \boldsymbol{\omega}^t)$, we get $(\mathbf{x}^t - \boldsymbol{\omega}^t)(\mathbf{x}^t - \mathbf{q}) \geq \text{err}(\mu^t)$, which together with Eq. 29 and Eq. 33 concludes that for non-drop step with no minor cycles, we have $\text{err}(\mu^t) - \text{err}(\mu^{t+1}) \geq \text{err}^2(\mu^t)/8Q^2$.

For non-drop step with one minor cycle, we use the Theorem 6 of (Chakrabarty et al., 2014). By a linear translation of adding all points with $-\boldsymbol{\omega}^t$, it gives

$$\|\mathbf{x}^t - \boldsymbol{\omega}^t\|^2 - \|\mathbf{x}^{t+1} - \boldsymbol{\omega}^t\|^2 \geq ((\mathbf{x}^t - \boldsymbol{\omega}^t)(\mathbf{x}^t - \mathbf{q}))^2/8Q^2. \quad (35)$$

Then applying the same argument as Eq. 34, and we finished our proof. □

Lemma 4.4. *After t major cycle steps of Algorithm 1 the number of drop steps is less than $t/2$.*

Proof. Recall that at the termination of a minor cycle, the size of the active set $|\mathcal{S}_c| \in [1, m]$. Since in each major cycle steps, the size of active set \mathcal{S}_t increases by one, and each drop step reduces the size of \mathcal{S}_t by at least one, the number of drop steps is always less than half of total number of the major cycle steps. \square

D PROOF OF THEOREM 4.5

Theorem 4.5 (Memory Complexity Bound). *For an constrained RL problem with m -dimensional measurement vector, in the worst case, a mixed policy needs to randomize among $m + 1$ individual policies to ensure convergence of RL oracles that search for deterministic policies.*

Proof. We give a constructive proof. Consider a m -dimensional vector-valued MDP with a single state, $m + 1$ actions, and $\mathbf{c}(a_i) := \mathbf{e}_i$ is the unit vector of i -th dimension for $i \in [1, m]$, and $\mathbf{c}(a_{m+1}) := \mathbf{0}$, and the episode terminates after 1 steps. The constrained RL problem is to find a policy whose measurement vector lies in the convex set of a single point $\{\mathbf{1}/2m\}$. By linear programming, it is clear that the only feasible mixed deterministic policy is to select a_{m+1} with $1/2$ probability, and the rest m actions with $1/2m$ probability; i.e. the unique feasible policy to this problem has an active set containing $m + 1$ deterministic policies. Therefore any method randomize among less than $m + 1$ individual policies does not ensure convergence when used with RL algorithms searching for deterministic policies. \square