Two Polyak-Type Step Sizes for Mirror Descent

Jun-Kai You[1,2] and Yen-Huan Li[1,3,4]

[1]Department of Computer Science and Information Engineering,
National Taiwan University
[2]MediaTek Research
[3]Department of Mathematics, National Taiwan University
[4]Center for Quantum Science and Engineering, National Taiwan
University

October 5, 2022

Abstract

We propose two Polyak-type step sizes for mirror descent and prove their convergences for minimizing convex locally Lipschitz functions. Both step sizes, unlike the original Polyak step size, do not need the optimal value of the objective function.

1   Introduction

Throughout this paper, consider the optimization problem

$$x^\star \in \underset{x \in \mathcal{X}}{\arg\min} f(x), \tag{P}$$

for some proper closed convex function f and non-empty closed convex set $\mathcal{X} \subseteq \mathbb{R}^d$. Denote by $g(x)$ a subgradient of f at x. When $\mathcal{X} = \mathbb{R}^d$, Polyak [21] proposed the subgradient method in Algorithm 1. The specific choice of η_k in Algorithm 1 is called the Polyak step size. The Polyak step size avoids evaluating the Lipschitz nor smoothness parameters of f, sometimes a difficult task. Nevertheless, its requires knowing $f^\star$, which limits its direct applications in practice.

Remarkably, the convergence rate of the Polyak step size is often satisfactory and can be even optimal. Polyak [21] proved that projected gradient descent with the Polyak step size converges at a linear rate when f is strongly convex and either smooth or Lipschitz. Later, Polyak [22, Chapter 5.3] showed that asymptotically, gradient descent with the Polyak step size converges at an $O(1/\sqrt{k})$ rate, where k denotes the number of iterations, when f is locally Lipschitz. Hazan and Kakade [10] showed that the Polyak step size yields the optimal iteration complexities achievable by gradient descent when f is Lipschitz, Lipschitz and strongly convex, smooth, or

1

Algorithm 1 Subgradient method with the Polyak step size

Require: $x_1 \in \mathbb{R}^d$.
 1: **for all** $k \in \mathbb{N}$ **do**
 2: $\quad \eta_k \leftarrow \frac{f(x_k) - f^\star}{\|g(x_k)\|_2^2}$, where $\|\cdot\|_2$ denotes the ℓ_2-norm.
 3: $\quad x_{k+1} = x_k - \eta_k g(x_k)$.
 4: **end for**
---

smooth and strongly convex. Loizou et al. [17] proposed a "stochastic Polyak step size" for stochastic gradient descent when f is a finite sum and studied the iteration complexity under standard smoothness and strong convexity conditions. The stochastic Polyak step size was later extended by D'Orazio et al. [8] for stochastic mirror descent. Ren et al. [23] studied the convergence of gradient descent with the Polyak step size under a generalized smoothness and generalized Łojasiewicz condition.

In this paper, we study the Polyak step size in the style of Polyak [22]. We do not aim to provide an iteration complexity bound under restrictive conditions on the objective function f, such as Lipschitzness and smoothness. Instead, we aim to guarantee asymptotic convergence for a very large class of f. In particular, we propose two algorithms based on mirror descent with Polyak-type step sizes and prove their convergences when f is locally Lipschitz. Our motivation is twofold.

1. There are several applications that violate the Lipschitz and smoothness assumptions, such as portfolio selection [18], Poisson inverse problems [5, 11], quantum state tomography [12, 20], and minimization of quantum Rényi divergences [26].

2. Though Polyak [22] has proved convergence of gradient descent with the Polyak step size, it is desirable to generalize the result for mirror descent. First, Polyak [22] does not consider the constrained optimization case. Second, projected gradient descent can generate infeasible iterates, causing the algorithm to "stall" before approaching the minimizer [15, 26].

The interested reader is referred to Appendix A for the details. Moreover, unlike the original Polyak step size studied by Polyak [22] and its mirror descent extension D'Orazio et al. [8], the two algorithms we propose do not need the optimal value $f^\star$.

There have been several variants of the Polyak stpe size that do not need $f^\star$ either; see, e.g., the discussions by Polyak [21], Brännlund [6] and a recent solution by Hazan and Kakade [10]. Among existing works, the most relevant to this paper are Nedić and Bertsekas [19] and Goffin and Kiwiel [9]. In particular, the two algorithms we propose are generalizations of the "first adjustment" considered by Nedić and Bertsekas [19] and the subgradient level method analyzed by Goffin and Kiwiel [9][1], respectively.

Our analyses differ significantly from those by Nedić and Bertsekas [19] and by Goffin and Kiwiel [9]. In particular, the non-Euclidean nature of mirror descent and

---

[1]According to Goffin and Kiwiel [9], the algorithm was proposed by Brännlund in his PhD thesis, but we cannot find the PhD thesis. Therefore, we cite Goffin and Kiwiel [9] instead of Brännlund's PhD thesis.

the fact that the mirror map may not be defined on the boundary of $\mathscr{X}$ render the proof strategies of Nedić and Bertsekas [19] and Goffin and Kiwiel [9] not directly applicable. The measure of the "traveling distance" in the subgradient level method is also slightly modified to fit in the mirror descent framework; see Section 4.1 for a discussion.

The first algorithm we propose had appeared in our recent paper [26]. That paper focuses on an application in quantum information theory and is targeted at information theorists. We isolate the optimization theory part and present it in a slightly more general form in Section 3.

2 Problem Formulation and Useful Facts

2.1 Problem Formulation

We consider solving the optimization problem (P) by mirror descent. Let h be a convex function and D_h the associated Bregman divergence, i.e.,

$$D_h(x, y) := h(x) - h(y) - \langle \nabla h(y), x - y \rangle, \quad \forall (x, y) \in \mathrm{dom}\, h \times \mathrm{dom}\, \nabla h.$$

Let $g(x)$ be a subgradient of f at x. Define $T(x; \eta)$ as a solution to the following minimization problem

$$\min_{y \in \mathscr{X}} \langle g(x), y - x \rangle + \frac{D_h(y, x)}{\eta} \tag{1}$$

for any *step size $\eta > 0$*. Mirror descent iterates by iteratively applying the mapping T with possibly different step sizes. The two "algorithms" we propose are indeed mirror descent with different step size selection rules. For convenience, we present the step size selection rule and mirror descent step (1) together and call them "algorithms."

2.2 Assumptions

We make the following assumptions in the rest of this paper. The first assumption is standard and ensures that f is continuous around the minimizer [2, Corollary 8.39].

Assumption 2.1. The optimal value $f^\star$ is finite.

The following two assumptions ensure that mirror descent with D_h is well-defined for solving (P). In particular, Assumption 2.2 resolves the "domain consistency" issue [1]; Assumption 2.3 ensures that $g(x)$ is well defined at any $x \in \mathscr{X} \cap \mathrm{int}\, \mathrm{dom}\, h$ [3, Corollary 16.18]. See, e.g., the discussion in Bauschke et al. [3].

Assumption 2.2. The function h is Legendre and the set $\mathscr{X} \cap \mathrm{dom}\, f$ is contained in the closure of $\mathrm{dom}\, h$.

Assumption 2.3. The relative interior of f contains $\mathscr{X} \cap \mathrm{int}\, \mathrm{dom}\, h$.

The following assumption is standard for analyzing mirror descent-type methods; see, e.g., Juditsky and Nemirovski [13, 14].

Assumption 2.4. The function h is strongly convex with respect to a norm $\|\cdot\|$ (not necessarily ℓ_2) on $\mathcal{X} \cap \operatorname{int} \operatorname{dom} h$; that is,

$$D_h(x,y) \geq \frac{1}{2}\|y-x\|^2, \quad \forall x, y \in \mathcal{X} \cap \operatorname{int} \operatorname{dom} h.$$

We will write $\|\cdot\|_*$ for the norm dual to $\|\cdot\|$. The last assumption, local boundedness of the subgradients, is key to our analyses. This assumption is also exploited by, e.g., Polyak [22] and Goffin and Kiwiel [9].

Assumption 2.5. The mapping $g(\cdot)$ is bounded on any compact subset of $\mathcal{X} \cap \operatorname{int} \operatorname{dom} h$.

One may equivalently assume that f is Lipschitz on any compact subset of $\mathcal{X} \cap \operatorname{int} \operatorname{dom} h$.

2.3 Useful Facts

The following results, which will be used in our analyses, are perhaps familiar to experts. We present them for the convenience of the reader.

The following theorem [1, Theorem 3.8(i)] helps verify that the iterates all lie in the interior of $\operatorname{dom} h$.

Theorem 2.1. *Suppose that h is Legendre. Let $x \in \operatorname{int} \operatorname{dom} h$ and $(x_k)_{k \in \mathbb{N}}$ be a sequence in $\operatorname{int} \operatorname{dom} h$. If $D_h(x, x_k) \to \infty$, then $(x_k)_{k \in \mathbb{N}}$ converges to a point on the boundary of $\operatorname{dom} h$.*

Corollary 2.2. *Suppose that h is Legendre. Let $x \in \operatorname{int} \operatorname{dom} h$ and $\{x_k\}_{k \in \mathbb{N}} \subset \operatorname{int} \operatorname{dom} h$. If $D_h(x, x_k) \leq c$ for some $c > 0$ for all $k \in \mathbb{N}$, then the closure of $\{x_k\}_{k \in \mathbb{N}}$ lies in $\operatorname{int} \operatorname{dom} h$.*

Let $x^\star$ be the minimizer. In our analyses, it is desirable to set x to be $x^\star$ in Theorem 2.1. Nevertheless, it can happen that $x^\star$ does not belong to[2] $\operatorname{int} \operatorname{dom} h$. The issue can be easily circumvented as f is closed.

Lemma 2.3. *Let $x^\star$ be a minimizer of f on $\mathcal{X}$. For any $\varepsilon > 0$, there exists some $x_\varepsilon \in \mathcal{X} \cap \operatorname{dom} h$ such that $f(x_\varepsilon) \leq f^\star + \varepsilon$.*

Proof. If $x^\star \in \mathcal{X} \cap \operatorname{int} \operatorname{dom} h$, then we can simply choose $x_\varepsilon = x^\star$. Otherwise, by Assumption 2.3, there is some y in the intersection of $\mathcal{X} \cap \operatorname{int} \operatorname{dom} h$ and the relative interior of $\operatorname{dom} f$. Define

$$\varphi(t) := f(x^\star + t(y - x^\star)), \quad \forall t \in \mathbb{R}.$$

It is easily checked that φ is also proper closed convex. Then, φ is continuous on the closure of its domain [3, Corollary 9.15]. Therefore, we can choose $x_\varepsilon = x^\star + t(y - x^\star)$ for some t small enough. $\qquad\square$

---

[2] Consider, for example, minimizing the function $f(x, y) := x$ on the the set

$$\{(x, y) \in \mathbb{R}^2 : x \geq 0, y \geq 0, x + y = 1\}$$

by entropic mirror descent. Then, h is the engative Shannon entropy and $\operatorname{int} \operatorname{dom} h$ is the interior of the positive orthant. Obviously, the minimizer is $(1, 0)$ and does not belong to $\operatorname{int} \operatorname{dom} h$

The following is an intermediate result in the standard analysis of mirror descent [14] and is seldom explicitly stated as a lemma. We provide its proof for completeness.

Lemma 2.4. *Let $x \in \mathcal{X} \cap \operatorname{int} \operatorname{dom} h$. Let $x_+ = T(x; \eta)$ for some $\eta > 0$. Then, $x_+ \in \mathcal{X} \cap \operatorname{int} \operatorname{dom} h$ and*

$$\eta\big(f(x) - f(y)\big) \le D_h(y, x) - D_h(y, x_+) + \frac{\eta^2 \|g(x)\|_*^2}{2}.$$

Proof. We write

$$\begin{aligned}
\eta_k\big(f(x) - f(y)\big) &\le \eta_k \langle g(x), x - y \rangle \\
&= \eta_k \langle g(x), x_+ - x \rangle + \eta_k \langle g(x), x - x_+ \rangle \\
&\le D_h(y, x) - D_h(y, x_+) - D_h(x_+, x) + \eta_k \langle g(x), x - x_+ \rangle \\
&\le D_h(y, x) - D_h(y, x_+) - \frac{\|x - x_+\|^2}{2} + \eta \langle g(x), x - x_+ \rangle \\
&\le D_h(y, x) - D_h(y, x_+) - \frac{\|x - x_+\|^2}{2} + \eta \|g(x)\|_* \|x - x_+\|.
\end{aligned}$$

The first inequality follows from the convexity of f; the second follows from the Bregman proximal inequality [24]; the third follows from Assumption 2.4. It remains to maximize the right-hand side with respect to $\|x - x_+\|$. $\qquad\square$

3   First Algorithm

3.1   Algorithm and Convergence Guarantee

Algorithm 2 presents the first algorithm we propose. The algorithm generalizes the "first adjustment" of Nedić and Bertsekas [19] by replacing the ℓ_2-norm with a general norm. Compared to Algorithm 1, the only difference lies in that $f^\star$ is replaced by a sequence of its estimates $(\hat{f}_k)_{k \in \mathbb{N}}$.

The following theorem guarantees that Algorithm 2 asymptotically converges to an approximate solution to (P).

Theorem 3.1. *If Assumptions 1–5 hold, then Algorithm 2 satisfies*

$$\inf_{k \in \mathbb{N}} f(x_k) \le f^\star + \delta.$$

Algorithm 2 requires deciding the error tolerance δ in advance and does not guarantee convergence to the exact minimum. Our second algorithm fixes the weaknesses.

3.2   Proof of Theorem 3.1

Suppose, for contradiction, that

$$\inf_{k \in \mathbb{N}} f(x_k) \ge f^\star + \delta + \varepsilon \tag{2}$$

for some $\varepsilon > 0$.

Algorithm 2 First Algorithm

Require: $\delta_1 \geq \delta > 0$, $\beta < 1$, $\gamma \geq 1$, $c > 1/2$, and $x_1 \in \mathcal{X} \cap \operatorname{int} \operatorname{dom} h$.

1: **for all** $k \in \mathbb{N}$ **do**
2:     **if** $g(x_k) = 0$ **then**
3:         Return x_k as a minimizer and terminate.
4:     **end if**
5:     $\hat{f}_k \leftarrow \min_{1 \leq \kappa \leq k} f(x_\kappa) - \delta_k$.
6:     $\eta_k \leftarrow \frac{f(x_k) - \hat{f}_k}{c \|g(x_k)\|_*^2}$.
7:     $x_{k+1} \leftarrow T(x_k; \eta_k)$.
8:     **if** $f(x_{k+1}) \leq \hat{f}_k$ **then**
9:         $\delta_{k+1} = \gamma \delta_k$.
10:    **else**
11:        $\delta_{k+1} = \max\{\beta \delta_k, \delta\}$.
12:    **end if**
13: **end for**

Lemma 3.2. *There is some $\tilde{k} \in \mathbb{N}$ such that $\delta_k = \delta$ for all $k \geq \tilde{k}$.*

Proof. Suppose that $f(x_{k+1}) \leq \hat{f}_k$ holds for infinitely many k's. Then, since $\delta_k \geq \delta$ for all $k \in \mathbb{N}$, we have $f(x_k) \to -\infty$, violating Assumption 2.1. Therefore, $f(x_{k+1}) \leq \hat{f}_k$ can only hold for finitely many k's. This implies that Line 11 of Algorithm 2 must be executed infinitely many times. The lemma follows. $\qquad\square$

By the definition of $\hat{f}_k$, (2), and Lemma 3.2, we have

$$\hat{f}_k \geq \inf_{k \in \mathbb{N}} f(x_k) - \delta_k \geq f^\star + \delta + \varepsilon - \delta = f^\star + \varepsilon, \quad \forall k \geq \tilde{k}.$$

By Lemma 2.3, there is some $\tilde{x} \in \mathcal{X} \cap \operatorname{int} \operatorname{dom} h$ such that

$$\hat{f}_k \geq f(\tilde{x}), \quad \forall k \geq \tilde{k}. \tag{3}$$

Then, we write

$$
\begin{aligned}
D_h(\tilde{x}, x_{k+1}) &\leq D_h(\tilde{x}, x_k) - \eta_k \left(f(x_k) - f(\tilde{x}) \right) + \frac{\eta_k^2 \|g(x_k)\|_*^2}{2} \\
&\leq D_h(\tilde{x}, x_k) - \eta_k \left(f(x_k) - \hat{f}_k \right) + \frac{\eta_k^2 \|g(x_k)\|_*^2}{2} \\
&= D_h(\tilde{x}, x_k) - \frac{\left(f(x_k) - \hat{f}_k \right)^2}{c \|g(x_k)\|_*^2} + \frac{\left(f(x_k) - \hat{f}_k \right)^2}{2c^2 \|g(x_k)\|_*^2} \\
&= D_h(\tilde{x}, x_k) - \frac{1}{c}\left(1 - \frac{1}{2c} \right) \left(\frac{f(x_k) - \hat{f}_k}{\|g(x_k)\|_*} \right)^2, \quad \forall k \geq \tilde{k}. \tag{4}
\end{aligned}
$$

In the above, the first inequality follows from Lemma 2.4; the second follows from (3); the third line follows from the definition of η_k.

By (4), the set $\{D_h(\tilde{x}, x_k)\}_{k \in \mathbb{N}}$ is bounded. By the strong convexity of h (Assumption 2.4), the set of iterates $\{x_k\}_{k \in \mathbb{N}}$ is also bounded in $\mathcal{X} \cap \operatorname{int} \operatorname{dom} h$. Hence, by Corollary 2.2, the closure of the set of iterates $\{x_k\}_{k \in \mathbb{N}}$ is bounded in $\mathcal{X} \cap \operatorname{int} \operatorname{dom} h$. Then, there is some $G > 0$ such that $\|g(x_k)\|_* \leq G$ for all $k \in \mathbb{N}$. The inequality (4) implies

$$D_h(\tilde{x}, x_{k+1}) \leq D_h(\tilde{x}, x_k) - \frac{1}{cG^2}\left(1 - \frac{1}{2c}\right)\left(f(x_k) - \hat{f}_k\right)^2, \quad \forall k \geq \tilde{k}$$

A telescopic sum gives

$$\frac{1}{cG^2}\left(1 - \frac{1}{2c}\right) \sum_{k=\tilde{k}}^{\infty} \left(f(x_k) - \hat{f}_k\right)^2 \leq D_h(\tilde{x}, x_{\tilde{k}}) < \infty,$$

showing that $f(x_k) - \hat{f}_k \to 0$. Nevertheless, Lemma 3.2 implies

$$f(x_k) - \hat{f}_k = f(x_k) - \min_{\kappa \leq k} f(\kappa) + \delta \geq \delta, \quad \forall k \geq \tilde{k},$$

a contradiction. The theorem follows.

4   Second Algorithm

4.1   Algorithm and Convergence Guarantee

Algorithm 3 presents the second algorithm we propose. The algorithm generalizes the subgradient level method analyzed by Goffin and Kiwiel [9].

Algorithm 3 is a direct adaptation of the subgradient level method to the mirror descent setup, except for Line 20. Line 1–17 are exactly the same as the corresponding parts in the subgradient level method. Line 18 generalizes the Polyak-type step size with a general norm, as in Algorithm 2. Line 19 replaces a gradient descent step with a mirror descent step. Line 20 is the major difference. In the original subgradient level method, σ_k is the sum of the Euclidean distances between consecutive iterates; in Line 20, σ_k becomes the sum of the magnitudes of the gradients scaled by $c\eta_k$. The two definitions of σ_k coincide, up to a scaling factor c, in the unconstrained optimization setup Goffin and Kiwiel [9] considered, but are obviously different in general.

Theorem 4.1. *If Assumptions 1–5 hold, then Algorithm 3 satisfies*

$$\inf_{k \in \mathbb{N}} f(x_k) = f^\star.$$

4.2   Proof of Theorem 4.1

We will prove by contradiction: If $\inf_{k \in \mathbb{N}} f(x_k) > f^\star$, then Lemma 4.2 and Lemma 4.3 below show that $\delta_l \to 0$ and δ_l bounded away from zero, respectively. Therefore, the desired convergence must hold.

Algorithm 3 Second Algorithm

Require: $\delta_1 > 0$, $\sigma_1 = 0$, $l = 1$, $k(1) = 1$, $B > 0$, $c > 1/2$, and $x_1 \in \mathscr{X} \cap \text{int dom} \, h$.

1: **for all** $k \in \mathbb{N}$ **do**
2:     **if** $g(x_k) = 0$ **then**
3:         Return x_k as a minimizer and terminate.
4:     **end if**
5:     $f_k^{\text{rec}} \leftarrow \min_{1 \le \kappa \le k} f(x_\kappa)$
6:     **if** $f(x_k) \le f_{k(l)}^{\text{rec}} - (1/2)\delta_l$ **then**
7:         $k(l+1) \leftarrow k$
8:         $\sigma_k \leftarrow 0$
9:         $\delta_{l+1} \leftarrow \delta_l$
10:        $l \leftarrow l+1$
11:     **else if** $\sigma_k > B$ **then**
12:        $k(l+1) \leftarrow k$
13:        $\sigma_k \leftarrow 0$
14:        $\delta_{l+1} \leftarrow (1/2)\delta_l$
15:        $l \leftarrow l+1$
16:     **end if**
17:     $\hat{f}_k \leftarrow f_{k(l)}^{\text{rec}} - \delta_l$
18:     $\eta_k \leftarrow \frac{f(x_k) - \hat{f}_k}{c \|g(x_k)\|_*^2}.$
19:     $x_{k+1} \leftarrow T(x_k; \eta_k).$
20:     $\sigma_{k+1} \leftarrow \sigma_k + c\eta_k \|g(x_k)\|_*.$
21: **end for**

Lemma 4.2. *If* $\inf_{k \in \mathbb{N}} f(x_k) > f^\star$, *then* $l \to \infty$ *and* $\delta_l \to 0$.

Proof. By Lemma 2.3, there is some point $\tilde{x} \in \mathscr{X} \cap \text{int dom} \, h$ near $x^\star$ such that

$$\inf_{k \in \mathbb{N}} f(x_k) \ge f(\tilde{x}). \tag{5}$$

Now we prove $l \to \infty$ by contradiction. Suppose that l is always bounded above by some $\bar{l} \in \mathbb{N}$. Define $s_k := c\eta_k \|g(x_k)\|_*$. By Line 12–15 in Algorithm 3, we have

$$\delta_l \ge \underline{\delta}, \quad \forall 1 \le l \le \bar{l}, \tag{6}$$

for some $\underline{\delta} > 0$, and there is some $\tilde{k} \in \mathbb{N}$ such that

$$\sum_{k=\tilde{k}+1}^{\infty} s_k \le B. \tag{7}$$

Iteratively applying Lemma 2.4, we write

$$D_h(\tilde{x}, x_{k+1}) \leq D_h(\tilde{x}, x_k) - \eta_k \left(f(x_k) - f(\tilde{x}) \right) + \frac{s_k^2}{2c^2}$$

$$\leq D_h(\tilde{x}, x_{\tilde{k}+1}) - \sum_{\kappa=\tilde{k}+1}^{k} \eta_\kappa \left(f(x_\kappa) - f(\tilde{x}) \right) + \frac{1}{2c^2} \sum_{\kappa=\tilde{k}+1}^{k} s_\kappa^2$$

$$\leq D_h(\tilde{x}, x_{\tilde{k}+1}) + \frac{B^2}{2c^2}, \quad \forall k > \tilde{k},$$

where the last inequality follows from (5) and the fact that $\sum_k s_k^2 \leq \left(\sum_k s_k \right)^2$. Therefore, $\{D_h(\tilde{x}, x_k)\}_{k \in \mathbb{N}}$ is bounded. By the strong convexity of h (Assumption 2.4), the set $\{x_k\}_{k \in \mathbb{N}}$ is also bounded. Corollary 2.2 then ensures that the closure of $\{x_k\}_{k \in \mathbb{N}}$ is bounded on the interior of $\operatorname{dom} h$. By Assumption 2.5, we have $\|g(x_k)\|_* \leq G$ for some $G > 0$ for all $k \in \mathbb{N}$.

The inequality (7) also implies that $\eta_k \|g(x_k)\|_* \to 0$. By the definition of η_k, we write

$$\eta_k \|g(x_k)\|_* = \frac{f(x_k) - \min_{\kappa \leq k} f(x_\kappa) + \delta_l}{c \|\nabla f(x_k)\|_*}$$

$$\geq \frac{f(x_k) - \min_{\kappa \leq k} f(x_\kappa) + \delta_l}{cG},$$

showing that

$$\lim_{k \to \infty} f(x_k) - \min_{\kappa \leq k} f(x_\kappa) + \delta_l = 0.$$

Then, by (6), for large enough k we have

$$f(x_k) - \min_{\kappa \leq k} f(x_\kappa) + \underline{\delta} \leq f(x_k) - \min_{\kappa \leq k} f(x_\kappa) + \delta_l \leq \frac{1}{2} \underline{\delta},$$

contradicting the fact that $f(x_k) \geq \min_{\kappa \leq k} f(x_\kappa)$.

Now we have proved that $l \to \infty$. By the monotone convergence theorem, $\delta_l \to \delta_\infty$ for some $\delta_\infty \geq 0$. Suppose that $\delta_\infty > 0$. Then, Line 11–15 cannot be executed infinitely many times; this together with $l \to \infty$ implies that Line 6–10 is executed infinitely many times. Then, $f(x_k) \to \infty$, violating Assumption 2.1. Therefore, we have $\delta_\infty = 0$. □

Remark. The lemma is inspired by the analysis of Goffin and Kiwiel [9, Lemma 2]. If we can choose $\tilde{x} = x^\star$ in the proof above, which is the case in the analysis of Goffin and Kiwiel [9], then there is no need to assume $\inf_{k \in \mathbb{N}} f(x_k) > f^\star$. Nevertheless, $x^\star$ may lie on the boundary of $\operatorname{dom} h$, which hinders us from using Theorem 2.2. See the footnote before Lemma 2.3 for an example.

Lemma 4.3. *If* $\inf_{k \in \mathbb{N}} f(x_k) > f^\star$, *then* $\delta_l \to \delta_\infty$ *for some* $\delta_\infty > 0$.

Proof. Notice that $\hat{f}_k$ and $f_{k(l)}^{\mathrm{rec}}$ are decreasing and bounded from below by $f^\star$. By the monotone convergence theorem, both $\lim_{k \to \infty} \hat{f}_k$ and $\lim_{k \to \infty} f_{k(l)}^{\mathrm{rec}}$ exist. Also

notice that $\lim_{k\to\infty} f^{\mathrm{rec}}_{k(l)} > f^{\star}$; otherwise, we have $\inf_{k\in\mathbb{N}} f(x_k) = f^{\star}$. Since $\delta_l \to 0$ by Lemma 4.2, we write

$$\lim_{k\to\infty} \hat{f}_k = \lim_{k\to\infty} \left(f^{\mathrm{rec}}_{k(l)} - \delta_l \right)$$
$$= \lim_{k\to\infty} f^{\mathrm{rec}}_{k(l)}$$
$$> f^{\star}.$$

Similarly as in the proof of Lemma 4.2, there exist some $\varepsilon > 0$ and $\tilde{x} \in \mathcal{X} \cap \mathrm{int}\,\mathrm{dom}\,h$ and $\tilde{k} \in \mathbb{N}$ such that[3]

$$f(x_k) > f(\tilde{x}) \text{ and } \hat{f}_k \geq f(\tilde{x}) + \varepsilon, \quad \forall k \geq \tilde{k}. \tag{8}$$

Then, we have

$$s_k := c\eta_k \|g(x_k)\|_* = \frac{f(x_k) - \hat{f}_k}{\|g(x_k)\|_*} \leq \frac{f(x_k) - f(\tilde{x}) - \varepsilon}{\|g(x_k)\|_*}, \quad \forall k \geq \tilde{k}. \tag{9}$$

We write

$$D_h(\tilde{x}, x_{k+1}) \leq D_h(\tilde{x}, x_k) - \eta_k \left(f(x_k) - f(\tilde{x}) \right) + \frac{s_k^2}{2c^2}$$
$$= D_h(\tilde{x}, x_k) - \frac{s_k \left(f(x_k) - f(\tilde{x}) \right)}{c \|g(x_k)\|_*} + \frac{s_k^2}{2c^2}$$
$$\leq D_h(\tilde{x}, x_k) - \frac{s_k \left(f(x_k) - f(\tilde{x}) \right)}{c \|g(x_k)\|_*} + \frac{s_k \left(f(x_k) - f(\tilde{x}) - \varepsilon \right)}{2c^2 \|g(x_k)\|_*}$$
$$= D_h(\tilde{x}, x_k) - \left(c - \frac{1}{2} \right) \frac{s_k \left(f(x_k) - f(\tilde{x}) \right)}{c^2 \|g(x_k)\|_*} - \frac{s_k \varepsilon}{2c^2 \|g(x_k)\|_*}$$
$$\leq D_h(\tilde{x}, x_k) - \frac{s_k \varepsilon}{2c^2 \|g(x_k)\|_*}, \quad \forall k \geq \tilde{k}. \tag{10}$$

In the above, the first inequality follows from Lemma 2.4; the second inequality follows from (9); the last follows from (8). Therefore, $D_h(\tilde{x}, x_k)$ is strictly decreasing. Similarly as in the proof of Lemma 4.2, we conclude that $\|g(x_k)\|_* \leq G$ for some $G > 0$. Summing the inequality (10) for $k \geq \tilde{k}$, we get

$$\frac{\varepsilon}{2c^2 G} \sum_{k=\tilde{k}}^{\infty} s_k < \infty.$$

This implies that $\sigma_k > B$ (Line 11 in Algorithm 3) can only happen a finite number of times, so δ_l does not converge to zero. $\qquad\square$

Lemma 4.2 and Lemma 4.3 together provide the desired contradiction. The theorem follows.

---

[3] Notice that there is a slight abuse of notations: the $\tilde{x}$ and $\tilde{k}$ here are different from those in the proof of Lemma 4.2.

Acknowledgements

References

[1] H. H. Bauschke and J. M. Borwein. Legendre functions and the method of random Bregman projections. *J. Convex Anal.*, 4(1):27–67, 1997.

[2] H. H. Bauschke and P. L. Combettes. *Convex analysis and monotone operator thoery in Hilbert spaces*. Springer, Cham, CH, second edition, 2017.

[3] H. H. Bauschke, J. Bolte, and M. Teboulle. A descent lemma beyond Lipschitz gradient continuity: first-order methods revisited and applications. *Math. Oper. Res.*, 42(2):330–348, 2017.

[4] A. Ben-Tal, T. Margalit, and A. Nemirovski. The ordered subsets mirror descent optimization method with applications to tomography. *SIAM J. Optim.*, 12(1): 79–108, 2001.

[5] M. Bertero, P. Boccacci, G. Desiderà, and G. Vicidomini. Image deblurring with Poisson data: from cells to galaxies. *Inverse Probl.*, 25, 2009.

[6] U. Brännlund. Nondifferentiable optimization: Subgradient optimization methods. In C. A. Floudas and P. M. Pardalos, editors, *Encyclopedia of Optimization*. Springer, New York, NY, 2001.

[7] L. Condat. Fast projection onto the simplex and the ℓ_1 ball. *Math. Program., Ser. A*, 158:575–585, 2016.

[8] R. D'Orazio, N. Loizou, I. Laradji, and I. Mitliagkas. On stochastic mirror descent: Convergence analysis and adaptive variants. In *Beyond First-Order Methods in ML Systems Workshop, Int. Conf. Machine Learning*, 2021.

[9] J.-L. Goffin and K. C. Kiwiel. Convergence of a simple subgradient level method. *Math. Program.*, 85:207–211, 1999.

[10] E. Hazan and S. M. Kakade. Revisiting the Polyak step size. 2019. arXiv:1905.00313v1.

[11] T. Hohage and F. Werner. Inverse problems with Poisson data: statistical regularization theory, applications and algoithms. *Inverse Probl.*, 32, 2016.

[12] Z. Hradil. Quantum-state estimation. *Phys. Rev. A*, 55(3), 1997.

[13] A. Juditsky and A. Nemirovski. First-order methods for nonsmooth convex large-scale optimization, I: General purpose methods. In S. Sra, S. Nowozin, and S. J. Wright, editors, *Optimization for Machine Learning*, chapter 5. MIT Press, Cambridge, MA, 2012.

[14] A. Juditsky and A. Nemirovski. First-order methods for nonsmooth convex large-scale optimization, II: Utilizing problem's structure. In S. Sra, S. Nowozin, and S. J. Wright, editors, *Optimization for Machine Learning*, chapter 5. MIT Press, Cambridge, MA, 2012.

[15] G. C. Knee, E. Bolduc, J. Leach, and E. M. Gauger. Maximum-likelihood quantum process tomography via projected gradient descent. 2018. arXiv:1803.10062v1.

[16] Y.-H. Li and V. Cevher. Convergence of the exponentiated gradient method with Armijo line search. *J. Optim. Theory Appl.*, 181(2):588–607, 2019.

[17] N. Loizou, S. Vaswani, I. Laradji, and S. Lacoste-Julien. Stochastic Polyak step-size for SGD: An adaptive learning rate for fast convergence. In *Proc. 24th Int. Conf. Artificial Intelligence and Statistics*, pages 1306–1314, 2021.

[18] L. C. MacLean, E. O. Thorp, and W. T. Ziemba, editors. *The Kelly Capital Growth Investment Criterion*. World Scientific, Singapore, 2012.

[19] A. Nedić and D. P. Bertsekas. Incremental subgradient methods for nondifferentiable optimization. *SIAM J. Optim.*, 12(1):109–138, 2001.

[20] M. Paris and J. Řeháček, editors. *Quantum State Estimation*. Springer, Berlin, 2004.

[21] B. T. Polyak. Minimization of unsmooth functionals. *USSR Comput. Math. Math. Phys.*, 9(3):14–29, 1969.

[22] B. T. Polyak. *Introduction to Optimization*. Optimization Softw., Inc., New York, NY, 1987.

[23] T. Ren, F. Cui, A. Atsidakou, S. Sanghavi, and N. Ho. Towards statistical and computational complexities of Polyak step size gradient descent. In *Proc. 25th Int. Conf. Artificial Intelligence and Statistics*, pages 3930–3961, 2022.

[24] M. Teboulle. A simplified view of first order methods for optimization. *Math. Program., Ser. B*, 170:67–96, 2018.

[25] Y. Vardi and D. Lee. From image deblurring to optimal investments: Maximum likelihood solutions for positive linear inverse problems. *J. R. Stat. Soc., Ser. B*, 55(3):569–612, 1993.

[26] J.-K. You, H.-C. Cheng, and Y.-H. Li. Minimizing quantum Rényi divergences via mirror descent with Polyak step size. In *2022 IEEE Int. Symp. Information Theory (ISIT)*, 2022.

A Optimization Problems Violating Lipschitz/Smoothness Conditions

Consider a stock market of d stocks. Denote the return rates of the n stocks by an entry-wise non-negative random variable a taking values in $\mathbb{R}^d$. Suppose a takes the value a_i with probability p_i[4]. The growth-optimal portfolio selection (aka the Kelly criterion) is given by

$$x^\star \in \underset{x \in \Delta}{\arg\min}\, f(x), \quad f(x) := \sum_{i=1}^n p_i \left[-\log \langle a_i, x \rangle \right], \tag{11}$$

where Δ denotes the probability simplex in $\mathbb{R}^n$. It is easily checked that the optimization problem (11) is convex and the function f is neither Lipschitz nor smooth. A simple calculation shows that f is not smooth relative to the negative Shannon entropy [16]. Although f is smooth relative to the Burg entropy [3], the resulting iteration rule lacks a closed-form and can be computationally expensive when the n is very large.

The maximum-likelihood estimator (MLE) for Poisson inverse problems can be transformed to (11) after a smart entry-wise scaling [25, 4]. The MLE for quantum state tomography takes the form [12]

$$\rho^\star \in \underset{\rho \in \mathscr{D}}{\arg\min}\, \frac{1}{n} \sum_{i=1}^n \left[-\log \operatorname{tr}\left(A_i \rho\right)\right], \tag{12}$$

where $\mathscr{D}$ denotes the set of Hermitian positive semi-definite matrices of unit traces and A_i are Hermitian positive semi-definite matrices. When all matrices involved share the same eigenbasis, (12) becomes a vector optimization problem of the form (11); that is, (12) is exactly the non-commutative counterpart of (11). The expressions of quantum Rényi divergences are complicated, so we omit them; the interested reader is referred to You et al. [26].

Consider solving (11) by projected gradient descent. Even if we ignore the step size selection issue, the so-called "stalling issue" [15] arises. Projection onto the probability simplex typically results in sparse vectors [7]. Since a_i can also be sparse and the product of sparse vectors can be exactly zero, the function f and its gradient can be undefined at a sparse iterate; then, projected gradient descent "stalls", though the iterate may be far from the minimizer. Knee et al. [15] and You et al. [26] provide related discussions for quantum process tomography and minimizing quantum Rényi divergences, respectively.

---

[4]In general, the support of a does not need to be finite. We focus on the finite support case to ease the discussion.