ON THE LOWER BOUND OF MINIMIZING POLYAK-ŁOJASIEWICZ FUNCTIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

Polyak-Łojasiewicz (PL) (Polyak, 1963) condition is a weaker condition than the strong convexity but suffices to ensure a global convergence for the Gradient Descent algorithm. In this paper, we study the lower bound of algorithms using first-order oracles to find an approximate optimal solution. We show that any first-order algorithm requires at least $\Omega\left(\frac{L}{\mu}\log\frac{1}{\varepsilon}\right)$ gradient costs to find an ε-approximate optimal solution for a general L-smooth function that has an μ-PL constant. This result demonstrates the *optimality* of the Gradient Descent algorithm to minimize smooth PL functions in the sense that there exists a "hard" PL function such that no first-order algorithm can be faster than Gradient Descent when ignoring a numerical constant. In contrast, it is well-known that the momentum technique, e.g. (Nesterov, 2003, chap. 2) can provably accelerate Gradient Descent to $O\left(\sqrt{\frac{L}{\hat{\mu}}}\log\frac{1}{\varepsilon}\right)$ gradient costs for functions that are L-smooth and $\hat{\mu}$-strongly convex. Therefore, our result distinguishes the hardness of minimizing a smooth PL function and a smooth strongly convex function as the complexity of the former cannot be improved by any polynomial order in general.

1 INTRODUCTION

We consider the problem

$$\min_{\mathbf{x}\in\mathbb{R}^d} f(\mathbf{x}), \tag{1}$$

where the function f is L-smooth and satisfies the Polyak-Łojasiewicz condition. A function f is said to satisfy the Polyak-Łojasiewicz condition if (2) holds for some $\mu > 0$:

$$\|\nabla f(\mathbf{x})\|^2 \geq 2\mu\left(f(\mathbf{x}) - \inf_{\mathbf{y}\in\mathbb{R}^d} f(\mathbf{y})\right), \quad \forall \mathbf{x}\in\mathbb{R}^d. \tag{2}$$

We refer to (2) as the μ-PL condition and simply denote $\inf_{\mathbf{y}\in\mathbb{R}^d} f(\mathbf{y})$ by f^*. The PL condition may be originally introduced by Polyak (Polyak, 1963) and Łojasiewicz (Lojasiewicz, 1963) independently. The PL condition is strictly weaker than strong convexity as one can show that any $\hat{\mu}$-strongly convex function which by definition satisfies:

$$f(\mathbf{x}) \geq f(\mathbf{y}) + \langle\nabla f(\mathbf{y}), \mathbf{x}-\mathbf{y}\rangle + \frac{\hat{\mu}}{2}\|\mathbf{x}-\mathbf{y}\|^2$$

is also $\hat{\mu}$-PL by minimizing both sides with respect to $\mathbf{x}$ (Karimi et al., 2016). However, the PL condition does not even imply convexity. From a geometric view, the PL condition suggests that the sum of the squares of the gradient dominates the optimal function value gap, which implies that any local stationary point is a global minimizer. Because it is relatively easy to obtain an approximate local stationary point by first-order algorithms, the PL condition serves as an ideal and weaker alternative to strong convexity.

In machine learning, the PL condition has received wide attention recently. Lots of models are found to satisfy this condition under different regimes. Examples include, but are not limited to, matrix decomposition and linear neural networks under a specific initialization (Hardt & Ma, 2016; Li et al., 2018), nonlinear neural networks in the so-called neural tangent kernel regime (Liu et al.,

2022), reinforcement learning with linear quadratic regulator (Fazel et al., 2018). Compared with strong convexity, the PL condition is much easier to hold since the reference point in the latter only is a minimum point such that $\mathbf{x}^* = \operatorname{argmin}_{\mathbf{y}} f(\mathbf{y})$, instead of any $\mathbf{y}$ in the domain.

Turning to the theoretic side, it is known (Karimi et al., 2016) that the standard Gradient Descent algorithm admits a linear converge to minimize a L-smooth and μ-PL function. To be specific, in order to find an ε-approximate optimal solution $\hat{\mathbf{x}}$ such that $f(\hat{\mathbf{x}}) - f^* \leq \varepsilon$, Gradient Decent needs $O(\frac{L}{\mu}\log\frac{1}{\varepsilon})$ gradient computations. However, it is still not clear whether there exist algorithms that can achieve a provably faster convergence rate. In the optimization community, it is perhaps well-known that the momentum technique, e.g. Nesterov (2003, chap. 2), can provably accelerate Gradient Descent from $O(\frac{L}{\hat{\mu}}\log\frac{1}{\varepsilon})$ to $O\left(\sqrt{\frac{L}{\hat{\mu}}}\log\frac{1}{\varepsilon}\right)$ for functions that are L-smooth and $\hat{\mu}$-strongly convex. Even though some works (J Reddi et al., 2016; Lei et al., 2017) have considered accelerations under different settings, probably faster convergence of first-order algorithms for PL functions is still not obtained up to now.

In this paper, we study the first-order complexities to minimize a generic smooth PL function and ask the question:

"Is the Gradient Decent algorithm (nearly) optimal or can we design a much faster algorithm?"

We answer the question in the language of min-max lower bound complexity for minimizing the L-smooth and μ-PL function class. We analyze the worst complexity of minimizing any function that belongs to the class using first-order algorithms. Excitingly, we construct a hard instance function showing that any first-order algorithm requires at least $\Omega\left(\frac{L}{\mu}\log\frac{1}{\varepsilon}\right)$ gradient costs to find an ε-approximate optimal solution. This answers the aforementioned question in an explicit way: the Gradient Descent algorithm is already *optimal* in the sense that no first-order algorithm can achieve a provably faster convergence rate in general ignoring a numerical constant. For the first time, we distinguish the hardness of minimizing a PL function and a strongly convex function in terms of first-order complexities, as the momentum technique for smooth and strongly convex functions provably accelerates Gradient Descent by a certain polynomial order.

It is worth mentioning that the optimization problem under our consideration is high-dimensional and the goal is to obtain the complexity bounds that do not have an explicit dependency on the dimension.

Our technique to establish the lower bound follows from the previous lower bounds in convex (Nesterov, 2003) and non-convex optimization (Carmon et al., 2021). The main idea is to construct a so-called "zero-chain" function ensuring that any first-order algorithm per-iteratively can only solve one coordinate of the optimization variable. Then for a "zero-chain" function that has a sufficiently high dimension, some number of entries will never reach their optimal values after the execution of any first-order algorithm in certain iterations. To obtain the desired $\Omega\left(\frac{L}{\mu}\log\frac{1}{\varepsilon}\right)$ lower bound, we propose a "zero-chain" function similar to Carmon et al. (2020), which is composed of the worst convex function designed by Nesterov (2003) and a separable function in the form as $\sigma\sum_{i=1}^{T} v_{T,c}(\mathbf{x}_i)$ or $\sum_{i=1}^{T} v_{\mathbf{y}_i}(\mathbf{x}_i)$ to destroy the convexity. Different from their separable function, the one that we introduce has a large Lipshictz constant. This property helps us to estimate the PL constant in a convenient way. This new idea gives new insights into the constructions and analyses of instance functions, which might be potentially generalized to establish the lower bounds for other non-convex problems.

NOTATION

We use bold letters, such as $\mathbf{x}$, to denote vectors in the Euclidean space $\mathbb{R}^d$, and bold capital letters, such as $\mathbf{A}$, to denote matrices. $\mathbf{I}_d$ denotes the identity matrix of size $d \times d$. We omit the subscript and simply denote $\mathbf{I}$ as the identity matrix when the dimension is clear from context. For $\mathbf{x} \in \mathbb{R}^d$, we use $\mathbf{x}_i$ to denote its ith coordinate. We use $\operatorname{supp}(\mathbf{x})$ to denote the subscripts of non-zero entries of $\mathbf{x}$, i.e. $\operatorname{supp}(\mathbf{x}) = \{i : \mathbf{x}_i \neq 0\}$. We use $\operatorname{span}\left\{\mathbf{x}^{(1)}, \cdots, \mathbf{x}^{(n)}\right\}$ to denote the linear subspace spanned by $\mathbf{x}^{(1)}, \cdots, \mathbf{x}^{(n)}$, i.e. $\left\{\mathbf{y} : \mathbf{y} = \sum_{i=1}^{n} a_i\mathbf{x}^{(i)}, a_i \in \mathbb{R}\right\}$. We call a function f L-smooth if ∇f is L-Lipschitz continuous, i.e. $\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\| \leq L\|\mathbf{x} - \mathbf{y}\|$. We denote $f^* = \inf_{\mathbf{x}} f(\mathbf{x})$.

We let $\mathbf{x}^*$ be any minimizer of f, i.e., $\mathbf{x}^* = \operatorname{argmin} f$. We always assume the existence of $\mathbf{x}^*$. We say that $\mathbf{x}$ is an ε-approximate optimal point of f when $f(\mathbf{x}) - f^* \leq \varepsilon$.

2 Related Work

Lower Bounds There has been a line of research concerning the lower bounds of algorithms on certain function classes. To the best of our knowledge, (Nemirovskij & Yudin, 1983) defines the oracle model to measure the complexity of algorithms, and most existing research on lower bounds follow this formulation of complexity. For convex functions and first-order oracles, the lower bound is studied in Nesterov (2003), where well-known optimal lower bound $\Omega(\varepsilon^{-\frac{1}{2}})$ and $\Omega(\kappa \log \frac{1}{\varepsilon})$ are obtained. For convex functions and nth-order oracles, lower bounds $\Omega\left(\varepsilon^{-\frac{2}{3n+1}}\right)$ have been proposed in Arjevani et al. (2019b). When the function is non-convex, it is generally NP-hard to find its global minima, or to test whether a point is a local minimum or a saddle point (Murty & Kabadi, 1985). Instead of finding ε-approximate optimal points, an alternative measure is finding ε-stationary points where $\|\nabla f(\mathbf{x})\| \leq \varepsilon$. Sometimes, additional constraints on the Hessian matrices of second-order stationary points are needed. Results of this kind include Carmon et al. (2020; 2021); Fang et al. (2018); Zhou & Gu (2019); Arjevani et al. (2019a; 2020). Though a PL function may be non-convex, it is tractable to find an ε-approximate optimal point, as local minima of a PL function must be global minima. In this paper, we give the lower complexity bound for finding ε-approximate optimal points.

PL Condition The PL condition was introduced by Polyak (Polyak, 1963) and Łojasiewicz (Lojasiewicz, 1963) independently. Besides the PL condition, there are other relaxations of the strong convexity, including error bounds (Luo & Tseng, 1993), essential strong convexity (Liu et al., 2014), weak strong convexity (Necoara et al., 2019), restricted secant inequality (Zhang & Yin, 2013), and quadratic growth (Anitescu, 2000). Karimi et al. (2016) discussed the relationships between these conditions. All these relaxations implies the PL condition except for the quadratic growth, which implies that the PL condition is quite general. There are many other papers that study designing practical algorithms to optimize a PL objective function under different scenarios, for example, Bassily et al. (2018); Nouiehed et al. (2019); Hardt & Ma (2016); Fazel et al. (2018); J Reddi et al. (2016); Lei et al. (2017).

3 Preliminaries

3.1 Upper bound on PL functions

Although the PL condition is a weaker condition than strong convexity, it guarantees linear convergence for Gradient Descent. The result can be found in Polyak (1963) and Karimi et al. (2016). We present it here for completeness.

Theorem 1. *If f is L-smooth and satisfies μ-PL condition, then the Gradient Descent algorithm with a constant step-size $\frac{1}{L}$:*

$$\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} - \frac{1}{L}\nabla f(\mathbf{x}^{(k)}), \tag{3}$$

has a linear convergence rate. We have:

$$f(\mathbf{x}^{(k)}) - f^* \leq \left(1 - \frac{\mu}{L}\right)^k (f(\mathbf{x}^{(0)}) - f^*). \tag{4}$$

Theorem 1 shows that the Gradient Descent algorithm finds the ε-approximate optimal point of f in $O\left(\frac{L}{\mu}\log\frac{1}{\varepsilon}\right)$ gradient computations. This gives an upper complexity bound for first-order algorithms. However, it remains open to us whether there are faster algorithms for smooth PL functions. We will establish a lower complexity bound on first-order algorithms, which nearly matches the upper bound.

3.2 DEFINITIONS OF ALGORITHM CLASSES AND FUNCTION CLASSES

An algorithm is a mapping from real-valued functions to sequences. For algorithm A and $f : \mathbb{R}^d \to \mathbb{R}$, we define $\mathsf{A}[f] = \{\mathbf{x}^{(i)}\}_{i\in\mathbb{N}}$ to be the sequence of algorithm A acting on f, where $\mathbf{x}^{(i)} \in \mathbb{R}^d$.

Note here, the algorithm under our consideration works on function defined on any Euclidean space. We call it the dimension-free property of the algorithm.

The definition of algorithms abstracts away from the the optimization process of a function. We consider algorithms which only make use of the first-order information of the iteration sequence. We call them first-order algorithms. If an algorithm is a first-order algorithm, then

$$\mathbf{x}^{(i)} = \mathsf{A}^{(i)}\left(\mathbf{x}^{(0)}, \nabla f(\mathbf{x}^{(0)}), \cdots, \mathbf{x}^{(i-1)}, \nabla f(\mathbf{x}^{(i-1)})\right), \tag{5}$$

where $\mathsf{A}^{(i)}$ is a function depending on A. Perhaps the simplest example of first-order algorithms is Gradient Descent.

We are interested in finding an ε-approximate point of a function f. Given a function $f : \mathbb{R}^d \to \mathbb{R}$ and an algorithm A, the complexity of A on f is the number of queries to the first-order oracle needed to find an ε-approximate point. We denote $T_\varepsilon(\mathsf{A}, f)$ to be the gradient complexity of A on f, then

$$T_\varepsilon(\mathsf{A}, f) = \min_t \left\{t : f(\mathbf{x}^{(t)}) - f^* \le \varepsilon\right\}. \tag{6}$$

In practice, we do not have the full information of the function f. We only know that f is in a particular function class $\mathcal{F}$, such as L-smooth functions. Given an algorithm A. We denote $\mathsf{T}_\varepsilon(\mathsf{A}, \mathcal{F})$ to be the complexity of A on $\mathcal{F}$, and define $\mathsf{T}_\varepsilon(\mathsf{A}, \mathcal{F})$ as follows:

$$\mathsf{T}_\varepsilon(\mathsf{A}, \mathcal{F}) = \sup_{f\in\mathcal{F}} T_\varepsilon(\mathsf{A}, f). \tag{7}$$

Thus, $\mathsf{T}_\varepsilon(\mathsf{A}, \mathcal{F})$ is the worst-case complexity of functions $f \in \mathcal{F}$.

For searching an ε-approximate optimal point of a function in $\mathcal{F}$, we need to find an algorithm which have a low complexity on $\mathcal{F}$. Denote an algorithm class by $\mathcal{A}$. The lower bound of an algorithm class on $\mathcal{F}$ describes the efficiency of algorithm class $\mathcal{A}$ on function class $\mathcal{F}$, which is defined to be

$$\mathcal{T}_\varepsilon(\mathcal{A}, \mathcal{F}) = \inf_{\mathsf{A}\in\mathcal{A}} \mathsf{T}_\varepsilon(\mathsf{A}, \mathcal{F}) = \inf_{\mathsf{A}\in\mathcal{A}} \sup_{f\in\mathcal{F}} T_\varepsilon(\mathsf{A}, f). \tag{8}$$

3.3 ZERO-RESPECTING ALGORITHM

Among all the algorithms, a special algorithm class is called zero-respecting algorithms. If A is a zero-respecting algorithm and $\mathsf{A}[f] = \left\{\mathbf{x}^{(t)}\right\}_{t\in\mathbb{N}}$, then the following condition holds for all $f : \mathbb{R}^d \to \mathbb{R}$:

$$\operatorname{supp}\{\mathbf{x}^{(n)} - \mathbf{x}^{(0)}\} \in \bigcup_{i=1}^{n-1} \operatorname{supp}\{\nabla f(\mathbf{x}^{(i)})\}. \tag{9}$$

That is to say, $\mathbf{x}^{(n)} - \mathbf{x}^{(0)}$ lies in the linear subspace spanned by $\nabla f(\mathbf{x}^{(0)}), \cdots, \nabla f(\mathbf{x}^{(n-1)})$. We denote the collection of first-order zero-respecting algorithms with $\mathbf{x}^{(0)} = \mathbf{0}$ by $\mathcal{A}_{\mathrm{zr}}$. It is shown by Nemirovskij & Yudin (1983) that a lower complexity bound on first-order zero-respecting algorithms are also a lower complexity bound on all the first-order algorithm when the function class satisfies the orthogonal invariance property.

3.4 ZERO-CHAIN

A zero-chain f is a function that safisfies the following condition:

$$\operatorname{supp}(\mathbf{x}) \subseteq \{1, 2, \cdots, k\} \implies \operatorname{supp}(\nabla f(\mathbf{x})) \subseteq \{1, 2, \cdots, k+1\}, \quad \forall \mathbf{x}. \tag{10}$$

In other words, the support of $\nabla f(\mathbf{x})$ lies in a restricted linear subspace depending on the support of $\mathbf{x}$.

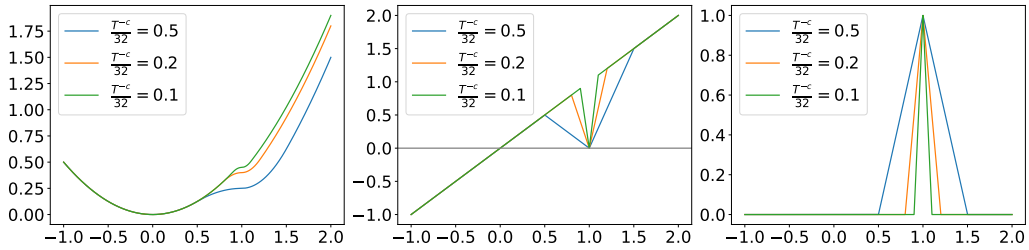

(a) $v_{T,c}$ (b) $v'_{T,c}$ (c) $b_{T,c}$

Figure 1: Illustration of $v_{T,c}$, $v'_{T,c}$ and $b_{T,c}$.

The "worst function in the (convex) world" in Nesterov (2003) defined as

$$f_k(\mathbf{x}) = \frac{1}{2}(\mathbf{x}_1 - 1)^2 + \sum_{i=1}^{k-1}(\mathbf{x}_{i+1} - \mathbf{x}_i)^2 \tag{11}$$

is a zero-chain, because if $\mathbf{x}_i = 0$ for $i > n$, then $(\nabla f_k(\mathbf{x}))_{i+1} = 0$ for $i > n$. A zero-chain is difficult to optimize for zero-respecting algorithms, because zero-respecting algorithms only discover one coordinate by one gradient computation.

4 Main Results

According to Theorem 1, we already have an upper complexity bound $O\left(\frac{L}{\mu}\log\frac{1}{\varepsilon}\right)$ by applying Gradient Descent to all the PL functions. In this section, we establish the lower complexity bound of first-order algorithms on PL functions. Let $\mathcal{P}(\Delta, \mu, L)$ be the collection of all L-smooth and μ-PL functions f with $f(\mathbf{x}^{(0)}) - f^* \leq \Delta$. We establish a lower bound of $\mathcal{T}_\varepsilon\left(\mathcal{A}_{\mathrm{zr}}, \mathcal{P}(\Delta, \mu, L)\right)$ by constructing a function which is hard to optimize for zero-respecting algorithms, and extend the result to first-order algorithms. We first propose a relatively simple hard instance which leads to a $\Omega\left(\kappa^{1-a}\right)$ lower bound in Section 4.1, where a can be any real number that belongs to $(0, 1)$. This helps in understanding our intuitions. Then we present a more complicated hard instance that can achieve the desired $\Omega\left(\kappa\log\frac{1}{\varepsilon}\right)$ lower bound in Section 4.2.

4.1 $\Omega\left(\kappa^{1-a}\right)$ Lower Bound

We introduce a hard function $f_{T,c,\sigma} : \mathbb{R}^T \to \mathbb{R}$ for first-order algorithms:

$$f_{T,c,\sigma}(\mathbf{x}) = q(\mathbf{x}) + \sigma\sum_{i=1}^{T} v_{T,c}(\mathbf{x}_i), \tag{12}$$

where

$$q(\mathbf{x}) = \frac{1}{2}\mathbf{x}_1^2 + \frac{1}{2}\sum_{i=1}^{T-1}\left(\mathbf{x}_{i+1} - \mathbf{x}_i\right)^2 \tag{13}$$

is a quadratic function, and we define $v_{T,c}$ as follows:

$$v_{T,c}(x) = \begin{cases} \frac{1}{2}x^2, & x \leq 1 - \frac{1}{32}T^{-c}, \\ \frac{1}{2}x^2 - 16T^c\left(x - 1 + T^{-c}\right)^2, & 1 - \frac{1}{32}T^{-c} < x \leq 1, \\ \frac{1}{2}x^2 - \frac{1}{32}T^{-c} + 16T^c\left(x - 1 - T^{-c}\right)^2, & 1 < x \leq 1 + \frac{1}{32}T^{-c}, \\ \frac{1}{2}x^2 - \frac{1}{32}T^{-c}, & x > 1 + \frac{1}{32}T^{-c} \end{cases} \tag{14}$$

where T is a positive integer, and c, σ are positive real numbers.

The function q can be rewritten as

$$q(\mathbf{x}) = \frac{1}{2}\mathbf{x}^T\mathbf{A}\mathbf{x}, \tag{15}$$

where

$$\mathbf{A} = \begin{bmatrix} 2 & -1 & & & \\ -1 & 2 & -1 & & \\ & \ddots & \ddots & \ddots & \\ & & -1 & 2 & -1 \\ & & & -1 & 1 \end{bmatrix} \tag{16}$$

is a positive-definite symmetric matrix.

From the definition of $v_{T,c}$ in (14), we obtain the derivative of $v_{T,c}$, which is:

$$v'_{T,c}(x) = \begin{cases} x, & x \le 1 - \frac{1}{32}T^{-c}, \\ x - 32T^c\left(x - 1 + \frac{1}{32}T^{-c}\right), & 1 - \frac{1}{32}T^{-c} < x \le 1, \\ x + 32T^c\left(x - 1 - \frac{1}{32}T^{-c}\right), & 1 < x \le 1 + \frac{1}{32}T^{-c}, \\ x, & x > 1 + \frac{1}{32}T^{-c}. \end{cases} \tag{17}$$

As we can see in (17), $v'_{T,c}$ is a piecewise linear function. To further simplify notations, we define $b_{T,c}(x)$ in (18):

$$b_{T,c}(x) = \begin{cases} 1 - 32T^c|x-1|, & 1 - \frac{1}{32}T^{-c} \le x \le 1 + \frac{1}{32}T^{-c}, \\ 0, & \text{otherwise}. \end{cases} \tag{18}$$

Then $v'_{T,c}(x) = x - b_{T,c}(x)$. Figure 1 provides geometric view of $v_{T,c}$ and $b_{T,c}$.

The quadratic part q is a translation of "the worst function in the (convex) world" in Nesterov (2003), and the definition of $v_{T,c}$ is inspired by the hard instance in Carmon et al. (2021). Our hard instance differs from previous ones mainly in the large Lipschitz constant of its gradient. We note that the controlled degree of nonsmoothness is crucial for our estimate of PL constant. In Lemma 1 we list some important properties of $f_{T,c,\sigma}$, which we prove in Section B.3.

Lemma 1. *If $\sigma < 1$ and $T^c \ge \frac{1}{2}\sigma^{-1}$, then $f_{T,c,\sigma}$ satisfies the following.*

1. *$g_{T,c,\sigma}(\mathbf{x}) = f_{T,c,\sigma}(\mathbf{1} - \mathbf{x})$ is a zero-chain.*
2. *$\mathbf{x}^* = \mathbf{0}$, $f^*_{T,c,\sigma} = 0$, $f_{T,c,\sigma}(x) \le \frac{1}{2}\mathbf{x}^T(\mathbf{A} + \sigma\mathbf{I})\mathbf{x}$.*
3. *$f_{T,c,\sigma}$ is $(4 + \sigma + 32\sigma T^c)$-smooth.*
4. *$f_{T,c,\sigma}$ satisfies the $\frac{1}{C_1 T^{1+5c}}$-PL condition, where C_1 is a universal constant.*

Now we study the lower bound of zero-respecting algorithms first. Let $\tilde{f}$ denote the following function:

$$\tilde{f}(\mathbf{x}) = \frac{LT^{-1}D^2}{42\sigma T^c} f_{T,c,\sigma}\left(\mathbf{1} - T^{1/2}D^{-1}\mathbf{x}\right), \tag{19}$$

where D is the distance between $\mathbf{0}$ and $\mathbf{x}^*$, and T, σ are parameters to be specified later. A change in D affects $\tilde{f}(\mathbf{0}) - \tilde{f}(\mathbf{x}^*)$ and $\|\mathbf{x}^* - \mathbf{0}\|$, but does not affect the condition number of $\tilde{f}$.

In Lemma 2, we show that $\tilde{f}$ is difficult to optimize for all first-order zero-respecting algorithms.

Lemma 2. *If $\sigma < 1$ and $T^c > \frac{1}{2}\sigma^{-1}$, a first-order zero-respecting algorithm with $\mathbf{x}^{(0)} = \mathbf{0}$ needs at least $\frac{T}{2}$ gradient computations to find a point $\mathbf{x}$ satisfying $\tilde{f}(\mathbf{x}) - \tilde{f}^* \le \frac{1}{16}(\tilde{f}(\mathbf{x}^{(0)}) - \tilde{f}^*)$.*

We give the full proof of Lemma 2 in Appendix **??** . The key point is that each gradient access of a zero-respecting algorithm only reveals one coordinate of a zero-chain. For $i > \frac{T}{2}$, $\mathbf{x}_i^{(k)}$ remains unchanged when $k \le \frac{T}{2}$, which gives a lower bound of the function value after the first $\frac{T}{2}$ gradient computations.

With Lemma 1 and Lemma 2 in hand, we are ready to give our lower complexity bound for zero-respecting first-order algorithms.

Theorem 2. *Given $L \geq \mu > 0$ and $c < 0.01$. When $\kappa = \frac{L}{\mu} > C_2$ where C_2 is a universal constant, we let*

$$T \in \left(\left(\frac{100}{C_2} \kappa \right)^{1/(1+6c)}, \left(\frac{200}{C_2} \kappa \right)^{1/(1+5c)} \right) \cap \mathbb{Z}, \tag{20}$$

and

$$\sigma = \frac{100}{C_2} \kappa T^{-(1+6c)}. \tag{21}$$

Then, $\tilde{f}$ is L-smooth and μ-PL. Moreover, any first-order zero-respecting algorithm with $\mathbf{x}^{(0)} = \mathbf{0}$ needs at least $\Omega\left(\kappa^{1/(1+6c)}\right)$ gradient computations to find a point $\mathbf{x}$ satisfying $\tilde{f}(\mathbf{x}) - \tilde{f}^* \leq \frac{1}{16}(\tilde{f}(\mathbf{x}^{(0)}) - \tilde{f}^*)$.

The proof of Theorem 2 is provided in Section B.5. For a satisfying $\frac{a}{6(1-a)} < 0.01$, let $c = \frac{a}{6(1-a)}$, then by Theorem 2, any zero-respecting algorithm needs at least $\Omega(\kappa^{1-a})$ gradient computations to find an ε-approximate optimal point of the function $\tilde{f}$.

Using the technique of Nemirovskij & Yudin (1983), for specific function classes such as PL functions, a lower complexity bound on first-order zero-respecting algorithms is also a lower complexity bound on all the first-order algorithms. Denoting the set of all first-order algorithms by $\mathcal{A}^{(1)}$, we have the following lemma:

Lemma 3.

$$\mathcal{T}_\varepsilon \left(\mathcal{A}^{(1)}, \mathcal{P}(\Delta, \mu, L) \right) \geq \mathcal{T}_\varepsilon \left(\mathcal{A}_{\text{zr}}, \mathcal{P}(\Delta, \mu, L) \right). \tag{22}$$

Note that given L and μ, the value $\tilde{f}(\mathbf{0})$ can be controlled by choosing an appropriate D. Lemma 2 and 3 leads to the main result of the paper:

Theorem 3. *For any $0 < a < 1$, when $\varepsilon \leq \frac{1}{16}\Delta$,*

$$\mathcal{T}_\varepsilon \left(\mathcal{A}^{(1)}, \mathcal{P}(\Delta, \mu, L) \right) \geq \Omega(\kappa^{1-a}) \tag{23}$$

Discussion about parameter setting In the definition of $\tilde{f}$, $\tilde{f}$ has hyper-parameters T, σ, c. From our construction, to achieve a lower bound of $\Omega(\kappa^{1-a})$ when a tends to zero, c also tends to zero. Parameters T and σ are chosen according to (20) and (21). When c tends to zero, $T = \Theta(\kappa)$ and σ tends to 1.

4.2 $\Omega\left(\kappa \log \frac{1}{\varepsilon}\right)$ LOWER BOUND

We show that the $\Omega(\kappa)$ lower bound can be further improved to $\Omega\left(\kappa \log \frac{1}{\varepsilon}\right)$ with a new hard instance based on $f_{T,c,\sigma}$ and a similar technique to estimate the PL constant. Detailed proof of Lemmas and Theorems in this subsection is provided in Appendix C .

We first introduce several components of the new hard instance. We define

$$v_y(x) = \begin{cases} \frac{1}{2}x^2, & x \leq \frac{31}{32}y, \\ \frac{1}{2}x^2 - 16\left(x - \frac{31}{32}y\right)^2, & \frac{31}{32}y < x \leq y, \\ \frac{1}{2}x^2 - \frac{y^2}{32} + 16\left(x - \frac{33}{32}y\right)^2, & y < x \leq \frac{33}{32}y, \\ \frac{1}{2}x^2 - \frac{y^2}{32}, & x > \frac{33}{32}y, \end{cases} \tag{24}$$

where $y > 0$ is a constant. By the definition of v_y, we have

$$v_y'(x) = \begin{cases} x, & x \leq \frac{31}{32}y, \\ x - 32\left(x - \frac{31}{32}y\right), & \frac{31}{32}y < x \leq y, \\ x + 32\left(x - \frac{33}{32}y\right), & y < x \leq \frac{33}{32}y, \\ x, & x > \frac{33}{32}y. \end{cases} \tag{25}$$

For the non-convex part, we define

$$b_y(x) = \begin{cases} y - 32|x-y|, & \frac{31}{32}y \le x \le \frac{33}{32}y, \\ 0, & \text{otherwise.} \end{cases} \tag{26}$$

Then we have $v_y'(x) = x - b_y(x)$.

For the convex part, we define $q_{T,t}(\mathbf{x})$ as follows (for the convenience of notation, we define $\mathbf{x}_0 = 0$):

$$q_{T,t}(\mathbf{x}) = \frac{1}{2}\sum_{i=0}^{t-1}\left[\left(\frac{7}{8}\mathbf{x}_{iT} - \mathbf{x}_{iT+1}\right)^2 + \sum_{j=1}^{T-1}\left(\mathbf{x}_{iT+j+1} - \mathbf{x}_{iT+j}\right)^2\right], \tag{27}$$

where $\mathbf{x} \in \mathbb{R}^{Tt}$. $q_{T,t}$ is a quadratic function of $\mathbf{x}$, thus can be written as

$$q_{T,t}(\mathbf{x}) = \frac{1}{2}\mathbf{x}^T\mathbf{B}\mathbf{x}, \tag{28}$$

where $\mathbf{B}$ is a positive semi-definite symmetric matrix. Like $\mathbf{A}$, $\mathbf{B}$ also satisfies $0 \preceq \mathbf{B} \preceq 4\mathbf{I}$, because the sum of absolute value of non-zero entries of each row of $\mathbf{B}$ is smaller or equal to 4.

Let $\mathbf{y} \in \mathbb{R}^{Tt}$ be a vector satisfying $\mathbf{y}_{qT+b} = \left(\frac{7}{8}\right)^q$, where $q \in \mathbb{N}$, $b \in \{1, 2, \cdots, T\}$. We define the hard instance $g_{T,t} : \mathbb{R}^{Tt} \to \mathbb{R}$ as follows:

$$g_{T,t}(\mathbf{x}) = q_{T,t}(\mathbf{x}) + \sum_{i=1}^{Tt} v_{\mathbf{y}_i}(\mathbf{x}_i). \tag{29}$$

Now we list some properties of $g_{T,t}$ in Lemma 4.

Lemma 4. *$f_{T,c,\sigma}$ satisfies the following.*

1. *$g_{T,t}(\mathbf{y} - \mathbf{x})$ is a zero-chain.*
2. *$\mathbf{x}^* = \mathbf{0}$, $g_{T,t}^* = 0$, $g_{T,t}(\mathbf{x}) \le \frac{1}{2}\mathbf{x}^T(\mathbf{B} + \mathbf{I})\mathbf{x}$.*
3. *$g_{T,t}$ is 37-smooth.*
4. *$g_{T,t}$ satisfies the $\frac{1}{C_3 T}$-PL condition, where C_3 is a universal constant.*

Define $\tilde{g}$ to be the following function, which is hard for first-order algorithms:

$$\tilde{g}(\mathbf{x}) = \frac{LT^{-1}D^2}{37} g_{T,t}\left(\mathbf{y} - T^{1/2}D^{-1}\mathbf{x}\right), \tag{30}$$

Similar to Lemma 2, we show that $\tilde{g}$ is hard for first-order zero-respecting algorithms:

Lemma 5. *Assume that $\varepsilon < 0.01$ and let $t = 2\left\lfloor \log_{\frac{8}{7}} \frac{3}{2\varepsilon} \right\rfloor$. A first-order zero-respecting algorithm with $\mathbf{x}^{(0)} = \mathbf{0}$ needs at least $\frac{1}{2}Tt$ gradient computations to find a point $\mathbf{x}$ satisfying $\tilde{g}(\mathbf{x}) - \tilde{g}^* \le \varepsilon(\tilde{g}(\mathbf{x}^{(0)}) - \tilde{g}^*)$.*

With Lemma 4 and 5, we obtain a lower bound for zero-respeting algorithms:

Theorem 4. *Given $L \ge \mu > 0$. When $\kappa = \frac{L}{\mu} > C_4$ where C_4 is a universal constant, there exists T and t such that $\tilde{g}$ is L-smooth and μ-PL. Moreover, any first-order zero-respecting algorithm with $\mathbf{x}^{(0)} = \mathbf{0}$ needs at least $\Omega\left(\kappa \log \frac{1}{\varepsilon}\right)$ gradient computations to find a point $\mathbf{x}$ satisfying $\tilde{g}(\mathbf{x}) - \tilde{g}^* \le \varepsilon(\tilde{g}(\mathbf{x}^{(0)}) - \tilde{g}^*)$.*

Finally, we arrive at a lower bound for first-order algorithms using Lemma 3:

Theorem 5. *For any $0 < a < 1$, when $\varepsilon \le \frac{1}{16}\Delta$,*

$$\mathcal{T}_\varepsilon\left(\mathcal{A}^{(1)}, \mathcal{P}(\Delta, \mu, L)\right) \ge \Omega\left(\kappa \log \frac{1}{\varepsilon}\right) \tag{31}$$

This bound matches the convergence rate of Gradient Descent up to a constant.

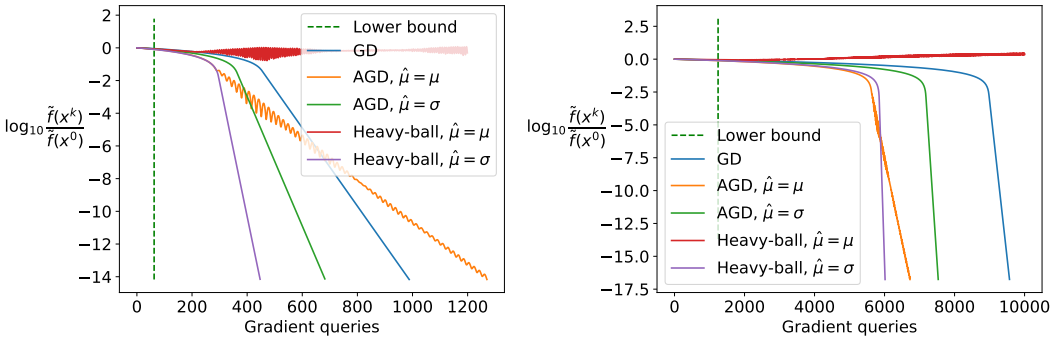

(a) $\kappa = 4.8384 \times 10^8$ (b) $\kappa = 4.8384 \times 10^{10}$

Figure 2: Convergence rate under Gradient Descent, Nesterov's Accelerated Gradient Descent and Polyak's Heavy-ball Method

5 Numerical Experiments

We conduct numerical experiments on our hard instance. We consider the κ relatively large, which can reduce the factors from the numerical constants. We $\kappa = 4.8384 \times 10^8$, $c = 0.001$ in Figure 2a and $\kappa = 4.8384 \times 10^{10}$, $c = 0.001$ in Figure 2b, then compute corresponding T and σ according to (20) and (21). We use Gradient Descent, Nesterov's Accelerated Gradient Descent (AGD) and Polyak's Heavy-ball Method to optimize the hard instance. As AGD and the Heavy-ball Method are designed for convex functions, we need to choose appropriate parameter $\hat{\mu}$ in both algorithms, because our hard instance is non-convex. The first choice is μ, which is the PL constant of our hard instance. The second choice is σ, because when $\mathbf{x}_i$ is far from 1, our hard instance can be treated as a σ-strongly convex quadratic function.

From Figure 2,we observe that the convergence can be roughly divided into two phases. In phase one, the optimization methods tries to pull every coordinate of $\tilde{\mathbf{x}}$ away from 1, and the function value does not decrease much. In phase two they converge linearly. This may be due to the fact that when each $\tilde{\mathbf{x}}_i$ is far from 1, our hard instance is exactly a strongly convex quadratic function. The observation is consistent with our theoretical results.

6 Conclusion

We construct a lower complexity bound on optimizing smooth PL functions with first-order methods. A first-order algorithm needs at least $\Omega\left(\frac{L}{\mu}\log\frac{1}{\varepsilon}\right)$ gradient access to find an ε-approximate optimal point of an L-smooth μ-PL function. Our lower bound matches the convergence rate of Gradient Descent up to constants.

We only focus on deterministic algorithms in this paper. We conjecture that our results can be extended to randomized algorithms, using the same technique in Nemirovskij & Yudin (1983) and explicit construction in Woodworth & Srebro (2016) and Woodworth & Srebro (2017). We leave its formal derivation to the future work.

References

Mihai Anitescu. Degenerate nonlinear programming with a quadratic growth condition. *SIAM Journal on Optimization*, 10(4):1116–1135, 2000.

Yossi Arjevani, Yair Carmon, John C Duchi, Dylan J Foster, Nathan Srebro, and Blake Woodworth. Lower bounds for non-convex stochastic optimization. *arXiv preprint arXiv:1912.02365*, 2019a.

Yossi Arjevani, Ohad Shamir, and Ron Shiff. Oracle complexity of second-order methods for smooth convex optimization. *Mathematical Programming*, 178(1-2):327–360, November 2019b. ISSN

0025-5610, 1436-4646. doi: 10.1007/s10107-018-1293-1. URL http://link.springer.com/10.1007/s10107-018-1293-1.

Yossi Arjevani, Yair Carmon, John C Duchi, Dylan J Foster, Ayush Sekhari, and Karthik Sridharan. Second-order information in non-convex stochastic optimization: Power and limitations. In *Conference on Learning Theory*, pp. 242–299. PMLR, 2020.

Raef Bassily, Mikhail Belkin, and Siyuan Ma. On exponential convergence of SGD in non-convex over-parametrized learning. *arXiv preprint arXiv:1811.02564*, 2018.

Yair Carmon, John C Duchi, Oliver Hinder, and Aaron Sidford. Lower bounds for finding stationary points I. *Mathematical Programming*, 184(1):71–120, 2020.

Yair Carmon, John C Duchi, Oliver Hinder, and Aaron Sidford. Lower bounds for finding stationary points II: first-order methods. *Mathematical Programming*, 185(1):315–355, 2021.

Cong Fang, Chris Junchi Li, Zhouchen Lin, and Tong Zhang. SPIDER: Near-optimal non-convex optimization via stochastic path-integrated differential estimator. *Advances in Neural Information Processing Systems*, 31, 2018.

Maryam Fazel, Rong Ge, Sham Kakade, and Mehran Mesbahi. Global convergence of policy gradient methods for the linear quadratic regulator. In *International Conference on Machine Learning*, pp. 1467–1476. PMLR, 2018.

Moritz Hardt and Tengyu Ma. Identity matters in deep learning. *arXiv preprint arXiv:1611.04231*, 2016.

Sashank J Reddi, Suvrit Sra, Barnabas Poczos, and Alexander J Smola. Proximal stochastic methods for nonsmooth nonconvex finite-sum optimization. *Advances in neural information processing systems*, 29, 2016.

Hamed Karimi, Julie Nutini, and Mark Schmidt. Linear convergence of gradient and proximal-gradient methods under the polyak-łojasiewicz condition. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 795–811. Springer, 2016.

Lihua Lei, Cheng Ju, Jianbo Chen, and Michael I Jordan. Non-convex finite-sum optimization via scsg methods. *Advances in Neural Information Processing Systems*, 30, 2017.

Yuanzhi Li, Tengyu Ma, and Hongyang Zhang. Algorithmic regularization in over-parameterized matrix sensing and neural networks with quadratic activations. In *Conference On Learning Theory*, pp. 2–47. PMLR, 2018.

Chaoyue Liu, Libin Zhu, and Mikhail Belkin. Loss landscapes and optimization in over-parameterized non-linear systems and neural networks. *Applied and Computational Harmonic Analysis*, 2022.

Ji Liu, Steve Wright, Christopher Ré, Victor Bittorf, and Srikrishna Sridhar. An asynchronous parallel stochastic coordinate descent algorithm. In *International Conference on Machine Learning*, pp. 469–477. PMLR, 2014.

Stanislaw Lojasiewicz. A topological property of real analytic subsets. *Coll. du CNRS, Les équations aux dérivées partielles*, 117(87-89):2, 1963.

Zhi-Quan Luo and Paul Tseng. Error bounds and convergence analysis of feasible descent methods: a general approach. *Annals of Operations Research*, 46(1):157–178, 1993.

Katta G Murty and Santosh N Kabadi. Some NP-complete problems in quadratic and nonlinear programming. Technical report, 1985.

Ion Necoara, Yu Nesterov, and Francois Glineur. Linear convergence of first order methods for non-strongly convex optimization. *Mathematical Programming*, 175(1):69–107, 2019.

Arkadij Semenovič Nemirovskij and David Borisovich Yudin. Problem complexity and method efficiency in optimization. 1983.

Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2003.

Maher Nouiehed, Maziar Sanjabi, Tianjian Huang, Jason D Lee, and Meisam Razaviyayn. Solving a class of non-convex min-max games using iterative first order methods. *Advances in Neural Information Processing Systems*, 32, 2019.

B.T. Polyak. Gradient methods for the minimisation of functionals. *USSR Computational Mathematics and Mathematical Physics*, 3(4):864–878, 1963.

Blake Woodworth and Nathan Srebro. Lower bound for randomized first order convex optimization. *arXiv preprint arXiv:1709.03594*, 2017.

Blake E Woodworth and Nati Srebro. Tight complexity bounds for optimizing composite objectives. *Advances in neural information processing systems*, 29, 2016.

Hui Zhang and Wotao Yin. Gradient methods for convex minimization: better rates under weaker conditions. *arXiv preprint arXiv:1303.4645*, 2013.

Dongruo Zhou and Quanquan Gu. Lower bounds for smooth nonconvex finite-sum optimization. In *International Conference on Machine Learning*, pp. 7574–7583. PMLR, 2019.

A Proof of Theorem 1

Theorem 1. *If f is L-smooth and satisfies μ-PL condition, then the Gradient Descent algorithm with a constant step-size $\frac{1}{L}$:*

$$\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} - \frac{1}{L}\nabla f(\mathbf{x}^{(k)}), \tag{32}$$

has a linear convergence rate. We have:

$$f(\mathbf{x}^{(k)}) - f^* \leq \left(1 - \frac{\mu}{L}\right)^k (f(\mathbf{x}^{(0)}) - f^*). \tag{33}$$

Proof. The proof is taken from Karimi et al. (2016). By the L-smoothness of f,

$$f(\mathbf{x}^{(k+1)}) \leq f(\mathbf{x}^{(k)}) + \left\langle \nabla f(\mathbf{x}^{(k)}), \mathbf{x}^{(k+1)} - \mathbf{x}^{(k)} \right\rangle + \frac{L}{2}\left\|\mathbf{x}^{(k+1)} - \mathbf{x}^{(k)}\right\|^2. \tag{34}$$

Applying (3) and (2) in (34), we have

$$f(\mathbf{x}^{(k+1)}) - f(\mathbf{x}^{(k)}) \leq \frac{1}{2L} - \left\|\nabla f(\mathbf{x}^{(k)})\right\|^2 \leq -\frac{\mu}{L}\left(f(\mathbf{x}^{(k)}) - f^*\right). \tag{35}$$

Rearranging terms in (35), we have

$$f(\mathbf{x}^{(k+1)}) - f^* \leq \left(1 - \frac{\mu}{L}\right)\left(f(\mathbf{x}^{(k)}) - f^*\right). \tag{36}$$

Applying (36) recursively gives the result. □

B Omitted proof in Section 4.1

B.1 Proof of Lemma 2

??

Lemma 2. *If $\sigma < 1$ and $T^c > \frac{1}{2}\sigma^{-1}$, a first-order zero-respecting algorithm with $\mathbf{x}^{(0)} = \mathbf{0}$ needs at least $\frac{T}{2}$ gradient computations to find a point $\mathbf{x}$ satisfying $\tilde{f}(\mathbf{x}) - \tilde{f}^* \leq \frac{1}{16}(\tilde{f}(\mathbf{x}^{(0)}) - \tilde{f}^*)$.*

Proof. By induction, we show that $\mathrm{supp}(\mathbf{x}^{(k)}) \subseteq \{1, 2, \cdots, k\}$.

- When $k = 0$, $\mathrm{supp}(\mathbf{x}^{(k)}) = \emptyset$.

- If $\mathrm{supp}(\mathbf{x}^{(k)}) \subseteq \{1, 2, \cdots, k\}$ holds for $k \leq n$, by the property of zero-chain, we have

$$\mathrm{supp}\left(\nabla \tilde{f}(\mathbf{x}^{(k)})\right) \in \{1, \cdots, k+1\}. \tag{37}$$

By the definition of zero-respecting algorithms,

$$\mathrm{supp}(\mathbf{x}^{(n+1)}) \subseteq \bigcup_{k=1}^{n} \mathrm{supp}\left(\nabla \tilde{f}(\mathbf{x}^{(k)})\right) \subseteq \{1, 2 \cdots, n+1\}. \tag{38}$$

Specifically, we have $\mathrm{supp}(\mathbf{x}^{(k)}) \in \{1, \cdots, T/2\}$ when $k \leq T/2$. Therefore,

$$\begin{aligned} \tilde{f}(\mathbf{x}^{(k)}) &\geq \frac{LT^{-1}D^2}{42\sigma T^c} \cdot \sigma \sum_{i=T/2+1}^{T} v_{T,c}(1) \\ &\geq \frac{LD^2}{336T^c}. \end{aligned} \tag{39}$$

According to the definition of $\tilde{f}$ in (19), $\tilde{f}^* = 0$, and

$$\begin{aligned} \tilde{f}(\mathbf{0}) &= \frac{LT^{-1}D^2}{42\sigma T^c} f_{T,c,\sigma}(\mathbf{1}) \\ &\leq \frac{LT^{-1}D^2}{42\sigma T^c}\left(\frac{1}{2} + \frac{1}{2}T\sigma\right) \quad \leq \frac{LD^2}{42T^c}. \end{aligned} \tag{40}$$

Therefore, for $k \leq \frac{T}{2}$, $\tilde{f}(\mathbf{x}^{(k)}) - \tilde{f}^* \geq \frac{1}{16}(\tilde{f}(\mathbf{x}^{(0)}) - \tilde{f}^*)$.

□

B.2 Proof of Lemma 3

Lemma 3.

$$\mathcal{T}_\varepsilon\left(\mathcal{A}^{(1)}, \mathcal{P}(\Delta, \mu, L)\right) \geq \mathcal{T}_\varepsilon\left(\mathcal{A}_{\mathrm{zr}}, \mathcal{P}(\Delta, \mu, L)\right). \tag{41}$$

Proof. The function class $\mathcal{P}(\Delta, \mu, L)$ enjoys an orthogonally invariance property, that is, if $f \in \mathcal{P}(\Delta, \mu, L)$, for any matrix $\mathbf{U}$ such that $\mathbf{U}^T\mathbf{U} = \mathbf{I}$, $f(\mathbf{U}\mathbf{x}) \in \mathcal{P}(\Delta, \mu, L)$ if the dimensions are compatible. This is a special case of Claim 1 in Bassily et al. (2018). Lemma 3 then follows from Proposition 1 in Carmon et al. (2020). □

B.3 Proof of Lemma 1

Before we prove Lemma 1, we present three auxiliary lemmas below.

Lemma 6. *If $\mathbf{x}_k = \max|\mathbf{x}_i| := m$ and $m \geq 1 + \frac{1}{32}T^{-c}$, then*

$$\|\nabla f_{T,c,\sigma}(\mathbf{x})\| \geq \left|\left(\nabla f_{T,c,\sigma}(\mathbf{x})\right)_k\right| \geq \frac{1}{240}T^{-(1+2c)/2}\sigma\|(\mathbf{A} + \sigma\mathbf{I})\mathbf{x}\|. \tag{42}$$

Lemma 7. *If $\sigma < 1$, $T^c \geq \frac{1}{2}\sigma^{-1}$, $1 - \frac{1}{32}T^{-c} \leq \max|\mathbf{x}_i| \leq 1 + \frac{1}{32}T^{-c}$, and*

$$\|\nabla f_{T,c,\sigma}(\mathbf{x})\| < \frac{1}{240}T^{-(1+2c)/2}\sigma\|(\mathbf{A} + \sigma\mathbf{I})\mathbf{x}\|, \tag{43}$$

then there exists a consecutive sequence of integers $k, k+1, \cdots, l$ ($k = l$ is allowed), such that

$$\mathbf{x}_{l+1} \geq 1 - \frac{1}{32}T^{-c}, \tag{44}$$

and

$$\mathbf{x}_{i+1} \geq \left(1 + \frac{\sigma}{2}\right)\mathbf{x}_i \tag{45}$$

for $i = k, \cdots, l$.

Lemma 8. *If $\sigma < 1$, $T \geq \frac{1}{2}\sigma^{-1/c}$, $\mathbf{x}_{l+1} \geq 1 - \frac{1}{32}T^{-c}$ and $\mathbf{x}_{l+1} \geq \left(1 + \frac{\sigma}{2}\right)\mathbf{x}_l$, then*

$$\|\nabla f_{T,c,\sigma}(\mathbf{x})\| \geq (\nabla f_{T,c,\sigma}(\mathbf{x}))_{l+1} \geq \frac{\sigma}{16}. \tag{46}$$

The proof of Lemma 6, 7 and Lemma 8 is provided in Appendix B.4. For $\mathbf{x} \in \mathbb{R}^T$, we may use one of the lemmas above to compute the PL constant at $\mathbf{x}$, and give a universal lower bound of PL constant which leads to Lemma 1. Now we present the proof of Lemma 1.

Lemma 1. *If $\sigma < 1$ and $T^c \geq \frac{1}{2}\sigma^{-1}$, then $f_{T,c,\sigma}$ satisfies the following.*

1. *$g_{T,c,\sigma}(\mathbf{x}) = f_{T,c,\sigma}(\mathbf{1} - \mathbf{x})$ is a zero-chain.*
2. *$\mathbf{x}^* = \mathbf{0}$, $f^*_{T,c,\sigma} = 0$, $f_{T,c,\sigma}(x) \leq \frac{1}{2}\mathbf{x}^T(\mathbf{A} + \sigma\mathbf{I})\mathbf{x}$.*
3. *$f_{T,c,\sigma}$ is $(4 + \sigma + 32\sigma T^c)$-smooth.*
4. *$f_{T,c,\sigma}$ satisfies the $\frac{1}{C_1 T^{1+5c}}$-PL condition, where C_1 is a universal constant.*

Proof. In the proof of Lemma 1, we define

$$\mathbf{b}_{T,c}(\mathbf{x}) = \begin{bmatrix} b_{T,c}(\mathbf{x}_1) \\ \vdots \\ b_{T,c}(\mathbf{x}_T) \end{bmatrix}. \tag{47}$$

1. We have
$$\nabla f_{T,c,\sigma}(\mathbf{x}) = \mathbf{A}\mathbf{x} + \sigma\mathbf{x} - \sigma\mathbf{b}_{T,c}(\mathbf{x}). \tag{48}$$
When $\mathbf{x}_i = \cdots = \mathbf{x}_T = 0$, $b_{T,c}(\mathbf{x}_j) = 1$ for $j \geq i$. When $k \geq i+1$, $(\nabla g_{T,c,\sigma}(\mathbf{x}))_k = (\mathbf{A}(\mathbf{1} - \mathbf{x}))_k = 0$. Therefore, $\operatorname{supp}(\nabla g_{T,c,\sigma}(\mathbf{x})) \in \{1, 2, \cdots, i+1\}$, which implies that $g_{T,c,\sigma}$ is a zero-chain.

2. $q(\mathbf{x})$ attains its minimum at $\mathbf{x} = \mathbf{0}$, and $v_{T,c}(x)$ attains its minimum at $x = 0$. Therefore, $f_{T,c,\sigma}(\mathbf{x}) = q(\mathbf{x}) + \sigma\sum_{t=1}^T v_{T,c}(\mathbf{x}_i)$ attains its minimum at $\mathbf{x} = \mathbf{0}$, and $f^*_{T,c,\sigma} = 0$. From the definition of $v_{T,c}$ in (14), we have $v_{T,c}(x) \leq \frac{1}{2}x^2$, which implies $f_{T,c,\sigma}(\mathbf{x}) \leq \frac{1}{2}\mathbf{x}^T(\mathbf{A} + \sigma\mathbf{I})\mathbf{x}$.

3. Let
$$v(\mathbf{x}) = \sum_{i=1}^T v_{T,c}(\mathbf{x}_i). \tag{49}$$
For $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^T$,
$$\begin{aligned} \|\nabla v(\mathbf{x}_1) - \nabla v(\mathbf{x}_2)\| &= \|\mathbf{x}_1 - \mathbf{b}_{T,c}(\mathbf{x}_1) - \mathbf{x}_2 + \mathbf{b}_{T,c}(\mathbf{x}_2)\| \\ &\leq \|\mathbf{x}_1 - \mathbf{x}_2\| + \|\mathbf{b}_{T,c}(\mathbf{x}_1) - \mathbf{b}_{T,c}(\mathbf{x}_2)\| \\ &\leq (1 + 32T^c)\|\mathbf{x}_1 - \mathbf{x}_2\|. \end{aligned} \tag{50}$$
The last inequality of (50) is due to the definition of $b_{T,c}$, which implies that $b_{T,c}$ is $32T^c$-Lipschitz. Consequently, $f_{T,c,\sigma}$ is $(4 + \sigma + 32\sigma T^c)$-smooth because $\mathbf{A} \preceq 4\mathbf{I}$.

4. We define $C_1 = 460800$. Let μ be the PL constant of $f_{T,c,\sigma}$.
$$\begin{aligned} \mu &= \inf_{\mathbf{x}\in\mathbb{R}^T} \frac{\|\nabla f_{T,c,\sigma}(\mathbf{x})\|^2}{2(f_{T,c,\sigma}(x) - f^*_{T,c,\sigma})} \\ &= \inf_{\mathbf{x}\in\mathbb{R}^T} \frac{\|\nabla f_{T,c,\sigma}(\mathbf{x})\|^2}{2f_{T,c,\sigma}(\mathbf{x})} \\ &\geq \inf_{\mathbf{x}\in\mathbb{R}^T} \frac{\|(\mathbf{A} + \sigma\mathbf{I})\mathbf{x} - \sigma\mathbf{b}_{T,c}(\mathbf{x})\|^2}{\mathbf{x}^T(\mathbf{A} + \sigma\mathbf{I})\mathbf{x}}, \end{aligned} \tag{51}$$

where the last inequality follows from property 2. We estimate μ by dividing $\mathbb{R}^T$, the domain of $f_{T,c,\sigma}$, into three parts. Define

$$\tilde{\mu}(\mathbf{x}) = \frac{\|(\mathbf{A}+\sigma\mathbf{I})\mathbf{x} - \mathbf{b}_{T,c}(\mathbf{x})\|^2}{\mathbf{x}^T(\mathbf{A}+\sigma\mathbf{I})\mathbf{x}}. \tag{52}$$

(a) If $\max_{i=0,\cdots,T+1} |\tilde{\mathbf{x}}_i| < 1 - \frac{1}{32}T^{-c}$, then $\mathbf{b}_{T,c}(\mathbf{x}) = \mathbf{0}$. Therefore,

$$\tilde{\mu}(\mathbf{x}) = \frac{\|(\mathbf{A}+\sigma\mathbf{I})\mathbf{x}\|^2}{\mathbf{x}^T(\mathbf{A}+\sigma\mathbf{I})\mathbf{x}} \geq \sigma, \tag{53}$$

which follows by $(\mathbf{A}+\mathbf{I})^2 = \mathbf{A}^2 + 2\sigma\mathbf{A} + \sigma^2\mathbf{I} \succeq \sigma(\mathbf{A}+\mathbf{I})$.

(b) If $\max_{i=0,\cdots,T+1} |\tilde{\mathbf{x}}_i| > 1 + \frac{1}{32}T^{-c}$, according to Lemma 6, we have

$$\begin{aligned}\tilde{\mu}(\mathbf{x}) &= \frac{\|\nabla f_{T,c,\sigma}(\mathbf{x})\|^2}{\mathbf{x}^T(\mathbf{A}+\sigma\mathbf{I})\mathbf{x}} \\ &\geq \frac{T^{-(1+2c)}\sigma^2\|(\mathbf{A}+\sigma\mathbf{I})\mathbf{x}\|^2}{57600\mathbf{x}^T(\mathbf{A}+\sigma\mathbf{I})\mathbf{x}} \\ &\geq \frac{\sigma^3}{57600T^{1+2c}},\end{aligned} \tag{54}$$

where the last inequality is due to $(\mathbf{A}+\mathbf{I})^2 = \mathbf{A}^2 + 2\sigma\mathbf{A} + \sigma^2\mathbf{I} \succeq \sigma(\mathbf{A}+\mathbf{I})$.

(c) If $1 - \frac{1}{32}T^{-c} \leq \max_{i=0,\cdots,T+1} |\tilde{\mathbf{x}}_i| \leq 1 + \frac{1}{32}T^{-c}$, and $\|\nabla f_{T,c,\sigma}(\mathbf{x})\| \geq \frac{1}{240}T^{-(1+2c)/2}\sigma\|(\mathbf{A}+\sigma\mathbf{I})\mathbf{x}\|$, we get the same lower bound of $\tilde{\mu}$ as in the case (b). If there does not exist i such that $\tilde{\mathbf{x}}_i \geq 1 - \frac{1}{32}T^{-c}$, then $\mathbf{b}_{T,c}(\mathbf{x}) = \mathbf{0}$, and we obtain the same lower bound in case (a). Otherwise, the the assumptions of Lemma 7 holds. The assumptions together with the result of Lemma 7 is exactly the assumptions of Lemma 8, so according to Lemma 8, $\|\nabla f_{T,c,\sigma}(\mathbf{x})\| \geq \frac{\sigma}{16}$. Therefore,

$$\begin{aligned}\tilde{\mu}(\mathbf{x}) &= \frac{\|\nabla f_{T,c,\sigma}(\mathbf{x})\|^2}{\mathbf{x}^T(\mathbf{A}+\sigma\mathbf{I})\mathbf{x}} \\ &\overset{(a)}{\geq} \frac{\sigma^2}{256} \cdot \frac{1}{5T\left(1+\frac{1}{32}T^{-c}\right)^2} \\ &\overset{(b)}{\geq} \frac{\sigma^2}{256} \cdot \frac{4}{15T} \\ &= \frac{\sigma^2}{960T},\end{aligned} \tag{55}$$

where (a) uses $(\mathbf{A}+\sigma\mathbf{I}) \preceq 5\mathbf{I}$, and (b) uses $T^{-c} < 1$. Now we establish a lower bound of $\tilde{\mu}(\mathbf{x})$ with (53), (54) and (55). We have

$$\tilde{\mu}(\mathbf{x}) \geq \min\left\{\sigma, \frac{\sigma^3}{57600T^{1+2c}}, \frac{\sigma^2}{960T}\right\}. \tag{56}$$

Because $T^c \geq \frac{1}{2}\sigma^{-1}$, $\frac{1}{460800T^{1+5c}}$ is a uniform lower bound of $\tilde{\mu}(\mathbf{x})$, which serves as the PL-constant of $f_{T,c,\sigma}$.

□

B.4 Proof of Lemma 6, 7 and 8

In this section, we first introduce some notations to simplify the proofs. Let

$$\tilde{\mathbf{x}} = \begin{bmatrix} 0 \\ \mathbf{x} \\ \mathbf{x}_T \end{bmatrix} \in \mathbb{R}^{T+2}, \tag{57}$$

$$\tilde{\mathbf{A}} = \begin{bmatrix} -\mathbf{e}^{(1)} & \mathbf{A}+\mathbf{E}_{T,T} & -\mathbf{e}^{(T)} \end{bmatrix} \in \mathbb{R}^{T\times(T+2)}, \tag{58}$$

and

$$\tilde{\mathbf{I}} = [\mathbf{0} \quad \mathbf{I}_T \quad \mathbf{0}] \in \mathbb{R}^{T\times(T+2)}, \tag{59}$$

where $\mathbf{e}^{(i)}$ is the ith column of $\mathbf{I}$, and $\mathbf{E}_{T,T}$ is an $T \times T$ whose (T, T) entry is 1 and other entries are 0. We have $\mathbf{A}\mathbf{x} = \tilde{\mathbf{A}}\tilde{\mathbf{x}}$, and $\mathbf{x} = \tilde{\mathbf{I}}\tilde{\mathbf{x}}$. We use $\tilde{\mathbf{x}}_0, \cdots, \tilde{\mathbf{x}}_{T+1}$ to denote the coordinates of $\tilde{\mathbf{x}}$, i.e. $\tilde{\mathbf{x}}_0 = 0$, $\tilde{\mathbf{x}}_i = \mathbf{x}_i$ for $i = 1, \cdots, T$, $\tilde{\mathbf{x}}_{T+1} = \mathbf{x}_T$. With these newly defined notations, we can check that

$$(\nabla f_{T,c,\sigma}(\mathbf{x}))_k = -\tilde{\mathbf{x}}_{k-1} + 2\tilde{\mathbf{x}}_k - \tilde{\mathbf{x}}_{k+1} + \sigma\tilde{\mathbf{x}}_k - \sigma b_{T,c}(\tilde{\mathbf{x}}_k) \tag{60}$$

for $k = 1, 2, \cdots, T$.

Before we present formal proofs below, we point out that $\mathbf{0} \preceq \mathbf{A} \preceq 4\mathbf{I}$. This fact will be used frequently to estimate norms and Lipschitz constants.

Lemma 6. *If $\mathbf{x}_k = \max_i \mathbf{x}_i := m$ and $m \geq 1 + \frac{1}{32}T^{-c}$, then*

$$\|\nabla f_{T,c,\sigma}(\mathbf{x})\| \geq \left|(\nabla f_{T,c,\sigma}(\mathbf{x}))_k\right| \geq \frac{1}{240}T^{-(1+2c)/2}\sigma\|(\mathbf{A} + \sigma\mathbf{I})\mathbf{x}\|. \tag{61}$$

Proof. We have

$$(\nabla f_{T,c,\sigma}(\mathbf{x}))_k = -\tilde{\mathbf{x}}_{k-1} + 2\tilde{\mathbf{x}}_k - \tilde{\mathbf{x}}_{k+1} + \sigma\tilde{\mathbf{x}}_k - \sigma b_{T,c}(\tilde{\mathbf{x}}_k). \tag{62}$$

For k such that $|\mathbf{x}_k| = \max|\mathbf{x}_i|$, we have $|\tilde{\mathbf{x}}_k| \geq |\tilde{\mathbf{x}}_{k-1}|$ and $|\tilde{\mathbf{x}}_k| \geq |\tilde{\mathbf{x}}_{k+1}|$, so $-\tilde{\mathbf{x}}_{k-1}+2\tilde{\mathbf{x}}_k-\tilde{\mathbf{x}}_{k+1}$ has the same sign with $\tilde{\mathbf{x}}_k$. From the definition of $b_{T,c}$, we have $|b_{T,c}(x)| \leq |x|$ for all $x \in \mathbb{R}$, so $\sigma\tilde{\mathbf{x}}_k - \sigma b_{T,c}(\tilde{\mathbf{x}}_k)$ also has the same sign with $\tilde{\mathbf{x}}_k$. Therefore, $-\tilde{\mathbf{x}}_{k-1} + 2\tilde{\mathbf{x}}_k - \tilde{\mathbf{x}}_{k+1}$ and $\sigma\tilde{\mathbf{x}}_k - \sigma b_{T,c}(\tilde{\mathbf{x}}_k)$ have the same sign. Therefore,

$$\begin{aligned}
\|\nabla f_{T,c,\sigma}(\mathbf{x})\| \geq \left|(\nabla f_{T,c,\sigma}(\mathbf{x}))_k\right| &\geq \sigma|\tilde{\mathbf{x}}_k - b_{T,c}(\tilde{\mathbf{x}}_k)| \\
&\geq \sigma\left(|\tilde{\mathbf{x}}_k| - |b_{T,c}(\tilde{\mathbf{x}}_k)|\right) \\
&\overset{(a)}{\geq} \sigma(m-1) \\
&\overset{(b)}{\geq} \frac{1}{48}T^{-c}\sigma m \\
&\geq \frac{1}{48}T^{-(1+2c)/2}\sigma\sqrt{Tm^2} \\
&\geq \frac{1}{48}T^{-(1+2c)/2}\sigma\|\mathbf{x}\| \\
&\overset{(c)}{\geq} \frac{1}{240}T^{-(1+2c)/2}\sigma\|(\mathbf{A} + \sigma\mathbf{I})\mathbf{x}\|,
\end{aligned} \tag{63}$$

□

where (a) is due to $b_{T,c}(x) \leq 1$, (b) is due to $\frac{m-1}{m} \geq \frac{1/32\cdot T^{-c}}{1+1/32\cdot T^{-c}} \geq \frac{1}{48}T^{-c}$, and (c) is due to $(\mathbf{A} + \sigma\mathbf{I}) \preceq 5\mathbf{I}$.

Lemma 7. *If $\sigma < 1$, $T^c \geq \frac{1}{2}\sigma^{-1}$, $1 - \frac{1}{32}T^{-c} \leq \max|\mathbf{x}_i| \leq 1 + \frac{1}{32}T^{-c}$, and*

$$\|\nabla f_{T,c,\sigma}(\mathbf{x})\| < \frac{1}{240}T^{-(1+2c)/2}\sigma\|(\mathbf{A} + \sigma\mathbf{I})\mathbf{x}\|, \tag{64}$$

then there exists a consecutive sequence of integers $k, k+1, \cdots, l$ ($k = l$ is allowed), such that

$$\mathbf{x}_{l+1} \geq 1 - \frac{1}{32}T^{-c}, \tag{65}$$

and

$$\mathbf{x}_{i+1} \geq \left(1 + \frac{\sigma}{2}\right)\mathbf{x}_i \tag{66}$$

for $i = k, \cdots, l$.

Proof. $\{\tilde{\mathbf{x}}_i\}_{i=0}^{T+1}$ is a sequence starting from 0, and has a maximum value between $1-\frac{1}{32}T^{-c}$ and $1+\frac{1}{32}T^{-c}$. Let l be the first index such that $\tilde{\mathbf{x}}_{l+1} \geq 1-\frac{1}{32}T^{-c}$. If $\tilde{\mathbf{x}}_l < 0.5$, letting $k=l$, (65) and (66) holds. Otherwise, there exists a consecutive sequence of integers $k, k+1, \cdots, l$, such that $\tilde{\mathbf{x}}_{k-1} \leq 0.5$, $0.5 < \tilde{\mathbf{x}}_i < 1-\frac{1}{32}T^{-c}$ for $i=k,\cdots,l$, and $\tilde{\mathbf{x}}_{l+1} \geq 1-\frac{1}{32}T^{-c}$. For $i=k-1,k,\cdots,l$, we have $b_{T,c}(\tilde{\mathbf{x}}_i)=0$, so

$$\begin{aligned}
|-\tilde{\mathbf{x}}_{i-1}+(2+\sigma)\tilde{\mathbf{x}}_i-\tilde{\mathbf{x}}_{i+1}| &= |-\tilde{\mathbf{x}}_{i-1}+(2+\sigma)\tilde{\mathbf{x}}_i-\tilde{\mathbf{x}}_{i+1}-\sigma b_{T,c}(\mathbf{x}_i)| \\
&= \left|(\nabla f_{T,c,\sigma}(\mathbf{x}))_i\right| \\
&\leq \|\nabla f_{T,c,\sigma}(\mathbf{x})\| \\
&\overset{(a)}{\leq} \frac{1}{240}T^{-(1+2c)/2}\sigma\|(\mathbf{A}+\sigma\mathbf{I})\mathbf{x}\| \\
&\overset{(b)}{\leq} \frac{\sigma^2}{4} \\
&\overset{(c)}{<} \frac{\sigma}{4},
\end{aligned} \tag{67}$$

where (a) is due to (64), (b) is due to $\|(\mathbf{A}+\sigma\mathbf{I})\mathbf{x}\| \preceq 5\mathbf{I}$ and $\|\mathbf{x}\| < \frac{6}{5}\sqrt{T}$, and (c) is due to $\sigma < 1$. We use induction on i.

1. When $i=k$,

$$\begin{aligned}
\tilde{\mathbf{x}}_{k+1} &\overset{(a)}{\geq} (2+\sigma)\tilde{\mathbf{x}}_k-\tilde{\mathbf{x}}_{k-1}-\frac{1}{4}\sigma \\
&\overset{(b)}{\geq} (1+\sigma)\tilde{\mathbf{x}}_k-\frac{1}{2}\sigma\tilde{\mathbf{x}}_k \\
&= \left(1+\frac{\sigma}{2}\right)\tilde{\mathbf{x}}_k,
\end{aligned} \tag{68}$$

   where (a) is due to (67), and (b) is due to $\tilde{\mathbf{x}}_{k-1} < \frac{1}{2} < \tilde{\mathbf{x}}_k$.

2. If (66) holds for $i=n$, then for $i=n+1$,

$$\begin{aligned}
\tilde{\mathbf{x}}_{n+1} &\overset{(a)}{\geq} (2+\sigma)\tilde{\mathbf{x}}_n-\tilde{\mathbf{x}}_{n-1}-\frac{1}{4}\sigma \\
&\overset{(b)}{\geq} (2+\sigma)\tilde{\mathbf{x}}_n-\frac{1}{1+\frac{1}{2}\sigma}\tilde{\mathbf{x}}_n-\frac{1}{2}\sigma\tilde{\mathbf{x}}_n \\
&\geq \left(1+\frac{1}{2}\sigma\right)\tilde{\mathbf{x}}_n,
\end{aligned} \tag{69}$$

   where (a) is due to (67), and (b) is due to $\tilde{\mathbf{x}}_n \geq (1+\frac{\sigma}{2})\tilde{\mathbf{x}}_{n-1}$ and $\tilde{\mathbf{x}}_n \geq \frac{1}{2}$.

□

Lemma 8. *If $\sigma < 1$, $T^c \geq \frac{1}{2}\sigma^{-1}$, $\mathbf{x}_{l+1} \geq 1-\frac{1}{32}T^{-c}$, $\max_i|\tilde{\mathbf{x}}_i| \leq 1+\frac{1}{32}T^{-c}$ and $\mathbf{x}_{l+1} \geq \left(1+\frac{\sigma}{2}\right)\mathbf{x}_l$, then*

$$\|\nabla f_{T,c,\sigma}(\mathbf{x})\| \geq (\nabla f_{T,c,\sigma}(\mathbf{x}))_{l+1} \geq \frac{11\sigma}{128}. \tag{70}$$

Proof.

$$
\begin{aligned}
\|\nabla f_{T,c,\sigma}(\mathbf{x})\| &\geq (\nabla f_{T,c,\sigma}(\mathbf{x}))_{l+1} \\
&\overset{(a)}{=} -\tilde{\mathbf{x}}_l + (2+\sigma)\tilde{\mathbf{x}}_{l+1} - \tilde{\mathbf{x}}_{l+2} - \sigma b_{T,c}(\tilde{\mathbf{x}}_{l+1}) \\
&\overset{(b)}{\geq} -\tilde{\mathbf{x}}_l + 2\tilde{\mathbf{x}}_{l+1} - \tilde{\mathbf{x}}_{l+2} + \sigma(1 - \frac{1}{32}T^{-c}) - \sigma \\
&\overset{(c)}{\geq} \left(2 - \frac{1}{1+\sigma/2}\right)\tilde{\mathbf{x}}_{l+1} - \left(1 + \frac{1}{32}T^{-c}\right) - \frac{1}{32}\sigma T^{-c} \\
&\overset{(d)}{\geq} \left(1 + \frac{\sigma}{4}\right)\tilde{\mathbf{x}}_{l+1} - 1 - \frac{1}{32}T^{-c} - \frac{1}{32}\sigma T^{-c} \\
&\overset{(e)}{\geq} \left(1 + \frac{\sigma}{4}\right)\left(1 - \frac{1}{32}T^{-c}\right) - 1 - \frac{1}{32}T^{-c} - \frac{1}{32}\sigma T^{-c} \\
&= \frac{1}{4}\sigma - \frac{1}{16}T^{-c} - \frac{5}{128}T^{-c}\sigma \\
&\overset{(f)}{\geq} \frac{11}{128}\sigma,
\end{aligned}
\tag{71}
$$

where (a) is due to the definition of $f_{T,c,\sigma}$, (b) is due to $\tilde{\mathbf{x}}_{l+1} \geq 1 - \frac{1}{32}T^{-c}$ and $b_{T,c}(x) \leq 1$, (c) is due to $\tilde{\mathbf{x}}_{l+1} \geq (1+\frac{\sigma}{2})\tilde{\mathbf{x}}_l$ and $\tilde{\mathbf{x}}_{l+2} \leq 1 + \frac{1}{32}T^{-c}$, (d) is due to $(1+\frac{\sigma}{2})(1-\frac{\sigma}{4}) \geq 1$ for $0 \leq \sigma \leq 1$, (e) is due to $\tilde{\mathbf{x}}_{l+1} \geq 1 - \frac{1}{32}T^{-c}$, and (f) is due to $T^{-c} \leq 2\sigma$ and $T^{-c} \leq 1$. □

B.5 Proof of Theorem 2

Theorem 2. *Given $L \geq \mu > 0$ and $c < 0.01$. When $\kappa = \frac{L}{\mu} > C_2$ where C_2 is a universal constant, we let*

$$
T \in \left(\left(\frac{100}{C_2}\kappa\right)^{1/(1+6c)}, \left(\frac{200}{C_2}\kappa\right)^{1/(1+5c)}\right) \cap \mathbb{Z}, \tag{72}
$$

and

$$
\sigma = \frac{100}{C_2}\kappa T^{-(1+6c)}. \tag{73}
$$

Then $\tilde{f}$ is L-smooth and μ-PL. Moreover, any first-order zero-respecting algorithm with $\mathbf{x}^{(0)} = \mathbf{0}$ needs at least $\Omega\left(\kappa^{1/(1+6c)}\right)$ gradient computations to find a point $\mathbf{x}$ satisfying $\tilde{f}(\mathbf{x}) - \tilde{f}^* \leq \frac{1}{16}(\tilde{f}(\mathbf{x}^{(0)}) - \tilde{f}^*)$.

Proof. Define $C_2 = 4200C_1$, where C_1 is the constant defined in Lemma 1. Because $\kappa > C_2$ and $c < 0.01$, we have

$$
\begin{aligned}
&\left(\frac{200}{C_2}\kappa\right)^{1/(1+5c)} - \left(\frac{100}{C_2}\kappa\right)^{1/(1+6c)} \\
&\geq 200^{1/1.05} - 100 \\
&> 50,
\end{aligned}
\tag{74}
$$

so there exists T satisfying (20).

It can be verified that T and σ in (20) and (21) satisfy the assumptions in Lemma 1, which is $\sigma < 1$ and $T^c \geq \frac{1}{2}\sigma^{-1}$. Therefore, we can use Lemma 1 to calculate the smoothness constant and PL constant of $\tilde{f}$.

According to Lemma 1, $f_{T,c,\sigma}$ is $(4 + \sigma + 32\sigma T^c)$-smooth. As $\sigma < 1$ and $\sigma T^c \geq \frac{1}{2}$, $f_{T,c,\sigma}$ is $42\sigma T^c$-smooth. The smoothness constant of $\tilde{f}$ is:

$$
\frac{LT^{-1}D^2}{42\sigma T^c} \cdot TD^{-2} \cdot 42\sigma T^c = L, \tag{75}
$$

and the PL constant of $\tilde{f}$ is:

$$\begin{aligned}\frac{LT^{-1}D^2}{42\sigma T^c}\cdot TD^{-2}\cdot\frac{1}{C_1T^{1+5c}} &= \frac{100L}{C_2T^{1+6c}\sigma}\\ &\overset{a}{=} \frac{100L}{C_2T^{1+6c}}\cdot\frac{C_2T^{1+6c}}{100\kappa}\\ &= \frac{L}{\kappa}\\ &= \mu,\end{aligned} \tag{76}$$

where $\overset{a}{=}$ uses (21).

Finally, by Lemma 2, any first-order zero-respecting algorithm needs at least $\frac{T}{2}$ access to gradient to find $\mathbf{x}$ such that $\tilde{f}(x)-\tilde{f}^*\le\frac{1}{16}\left(\tilde{f}(\mathbf{x}^{(0)})-\tilde{f}^*\right)$. By (20),

$$\frac{T}{2}\ge\frac{1}{2}\left(\frac{100\kappa}{C_2}\right)^{1/(1+6c)}=\Omega\left(\kappa^{1/(1+6c)}\right), \tag{77}$$

which completes the proof. □

C Omitted proof in Section 4.2

C.1 Proof of Lemma 5

Lemma 5. *Assume that $\varepsilon<0.01$ and let $t=2\left\lfloor\log_{\frac{8}{7}}\frac{3}{2\varepsilon}\right\rfloor$. A first-order zero-respecting algorithm with $\mathbf{x}^{(0)}=\mathbf{0}$ needs at least $\frac{1}{2}Tt$ gradient computations to find a point $\mathbf{x}$ satisfying $\tilde{g}(\mathbf{x})-\tilde{g}^*\le\varepsilon(\tilde{g}(\mathbf{x}^{(0)})-\tilde{g}^*)$.*

Proof. By induction, we have $\operatorname{supp}(\mathbf{x}^{(k)})\in\{1,\cdots,k\}$. By the definition of $\tilde{g}$ and v_y, we have

$$\begin{aligned}\tilde{g}(\mathbf{0}) &= \frac{1}{2}+\sum_{i=0}^{t-1}Tv_{(\frac{7}{8})^i}\left(\left(\frac{7}{8}\right)^i\right)\\ &= \frac{1}{2}+T\sum_{i=0}^{t-1}\frac{31}{64}\left(\frac{7}{8}\right)^{2i}\\ &= \frac{1}{2}+T\cdot\frac{31}{15}\cdot\left(1-\left(\frac{7}{8}\right)^{2t}\right)\\ &\le 3T.\end{aligned} \tag{78}$$

For $k\le\frac{1}{2}Tt$,

$$\begin{aligned}\tilde{g}(\mathbf{x}^{(k)}) &\ge \sum_{i=\frac{t}{2}}^{t}Tv_{(\frac{7}{8})^i}\left(\left(\frac{7}{8}\right)^i\right)\\ &= T\sum_{i=\frac{t}{2}}^{t}\frac{31}{64}\left(\frac{7}{8}\right)^{2i}\\ &= T\cdot\left(\frac{7}{8}\right)^{\frac{t}{2}}\cdot\frac{31}{15}\left(1-\left(\frac{7}{8}\right)^{t}\right)\\ &\ge 2T\cdot\frac{3\varepsilon}{2}\\ &= 3T\varepsilon.\end{aligned} \tag{79}$$

Therefore, for $k\le\frac{1}{2}Tt$, $\tilde{g}(\mathbf{x})-\tilde{g}^*\le\varepsilon(\tilde{g}(\mathbf{x}^{(0)})-\tilde{g}^*)$. □

C.2 PROOF OF LEMMA 4

Lemma 4. *$f_{T,c,\sigma}$ satisfies the following.*

1. *$g_{T,t}(\mathbf{y}-\mathbf{x})$ is a zero-chain.*
2. *$\mathbf{x}^* = \mathbf{0}$, $g^*_{T,t} = 0$, $g_{T,t}(\mathbf{x}) \le \frac{1}{2}\mathbf{x}^T(\mathbf{B}+\mathbf{I})\mathbf{x}$.*
3. *$g_{T,t}$ is 37-smooth.*
4. *$g_{T,t}$ satisfies the $\frac{1}{C_3 T}$-PL condition, where C_3 is a universal constant.*

Proof. In the proof of Lemma 4, we define

$$\mathbf{b}(\mathbf{x}) = \begin{bmatrix} b_{\mathbf{y}_1}(\mathbf{x}_1) \\ \vdots \\ b_{\mathbf{y}_{Tt}}(\mathbf{x}_{Tt}) \end{bmatrix}. \tag{80}$$

1. We have
$$\nabla g_{T,t}(\mathbf{y}-\mathbf{x}) = -(\mathbf{B}(\mathbf{y}-\mathbf{x}) + (\mathbf{y}-\mathbf{x}) - \mathbf{b}(\mathbf{y}-\mathbf{x})). \tag{81}$$
When $\mathbf{x}_i = \cdots = \mathbf{x}_{Tt} = 0$, $b_{\mathbf{y}_j}(\mathbf{y}_j - \mathbf{x}_j) = \mathbf{y}_j$ for $j \ge i$. When $k \ge i+1$, $(\nabla g_{T,t}(\mathbf{x}))_k = (\mathbf{B}(\mathbf{y}-\mathbf{x}))_k = 0$. Therefore, $\operatorname{supp}(\nabla g_{T,t}(\mathbf{y}-\mathbf{x})) \in \{1, 2, \cdots, i+1\}$, which implies that $g_{T,t}(\mathbf{y}-\mathbf{x})$ is a zero-chain with respect to $\mathbf{x}$.

2. $q_{T,t}(\mathbf{x})$ attains its minimum at $\mathbf{x} = \mathbf{0}$, and $v_y(x)$ attains its minimum at $x = 0$. Therefore, $g_{T,t}(\mathbf{x}) = q_{T,t}(\mathbf{x}) + \sum_{t=1}^{Tt} v_{\mathbf{y}_i}(\mathbf{x}_i)$ attains its minimum at $\mathbf{x} = \mathbf{0}$, and $g^*_{T,t} = 0$. From the definition of $v_{T,c}$ in (24), we have $v_y(x) \le \frac{1}{2}x^2$, which implies $g_{T,t}(\mathbf{x}) \le \frac{1}{2}\mathbf{x}^T(\mathbf{B}+\mathbf{I})\mathbf{x}$.

3. Let
$$v(\mathbf{x}) = \sum_{i=1}^{Tt} v_{\mathbf{y}_i}(\mathbf{x}_i). \tag{82}$$
For $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^{Tt}$,
$$\begin{aligned} \|\nabla v(\mathbf{x}_1) - \nabla v(\mathbf{x}_2)\| &= \|\mathbf{x}_1 - \mathbf{b}_{T,c}(\mathbf{x}_1) - \mathbf{x}_2 + \mathbf{b}_{T,c}(\mathbf{x}_2)\| \\ &\le \|\mathbf{x}_1 - \mathbf{x}_2\| + \|\mathbf{b}_{T,c}(\mathbf{x}_1) - \mathbf{b}_{T,c}(\mathbf{x}_2)\| \\ &\le 33\|\mathbf{x}_1 - \mathbf{x}_2\|. \end{aligned} \tag{83}$$
The last inequality of (50) is due to the definition of $b_{T,c}$, which implies that $b_{T,c}$ is 32-Lipschitz. Consequently, $f_{T,c,\sigma}$ is 37-smooth because $\mathbf{B} \preceq 4\mathbf{I}$.

4. Similar to the proof in B.4, we first introduce some notations to simplify the proofs. For $\mathbf{x} \in \mathbb{R}^{Tt}$, let
$$\tilde{\mathbf{x}} = \begin{bmatrix} 0 \\ \mathbf{x} \\ \mathbf{x}_{Tt} \end{bmatrix} \in \mathbb{R}^{Tt+2}, \tag{84}$$
$$\tilde{\mathbf{B}} = \begin{bmatrix} -\mathbf{e}^{(1)} & \mathbf{B} + \mathbf{E}_{Tt,Tt} & -\mathbf{e}^{(Tt)} \end{bmatrix} \in \mathbb{R}^{Tt \times (Tt+2)}, \tag{85}$$
and
$$\tilde{\mathbf{I}} = [\mathbf{0} \quad \mathbf{I}_{Tt} \quad \mathbf{0}] \in \mathbb{R}^{Tt \times (Tt+2)}, \tag{86}$$
where $\mathbf{e}^{(i)}$ is the ith column of $\mathbf{I}$, and $\mathbf{E}_{Tt,Tt}$ is an $Tt \times Tt$ whose (Tt, Tt) entry is 1 and other entries are 0. We have $\mathbf{B}\mathbf{x} = \tilde{\mathbf{B}}\tilde{\mathbf{x}}$, and $\mathbf{x} = \tilde{\mathbf{I}}\tilde{\mathbf{x}}$. We use $\tilde{\mathbf{x}}_0, \cdots, \tilde{\mathbf{x}}_{Tt+1}$ to denote the coordinates of $\tilde{\mathbf{x}}$, i.e. $\tilde{\mathbf{x}}_0 = 0$, $\tilde{\mathbf{x}}_i = \mathbf{x}_i$ for $i = 1, \cdots, Tt$, and $\tilde{\mathbf{x}}_{Tt+1} = \mathbf{x}_{Tt}$. Similarly, we define $\tilde{\mathbf{y}}$ by $\tilde{\mathbf{y}}_0 = 1$, $\tilde{\mathbf{y}}_i = \mathbf{y}_i$ for $i = 1, \cdots, Tt$, and $\tilde{\mathbf{y}}_{Tt+1} = \mathbf{y}_{Tt}$. With these newly defined notations, we can check that
$$(\nabla g_{T,t}(\mathbf{x}))_k = (\tilde{\mathbf{B}}\tilde{\mathbf{x}})_k + \tilde{\mathbf{x}}_k - b_{\tilde{\mathbf{y}}_k}(\tilde{\mathbf{x}}_k) \tag{87}$$
for $k = 1, 2, \cdots, Tt$.

We define $C_3 = \frac{1067220}{1083}$. Let μ be the PL constant of $g_{T,t}$.

$$\begin{aligned}
\mu &= \inf_{\mathbf{x}\in\mathbb{R}^{Tt}} \frac{\|\nabla g_{T,t}(\mathbf{x})\|^2}{2(g_{T,t}(x) - g^*_{T,t})} \\
&= \inf_{\mathbf{x}\in\mathbb{R}^{Tt}} \frac{\|\nabla g_{T,t}(\mathbf{x})\|^2}{2g_{T,t}(\mathbf{x})} \\
&\overset{(a)}{\geq} \inf_{\mathbf{x}\in\mathbb{R}^{Tt}} \frac{\|(\mathbf{B}+\mathbf{I})\mathbf{x} - \mathbf{b}_{T,c}(\mathbf{x})\|^2}{\mathbf{x}^T(\mathbf{B}+\mathbf{I})\mathbf{x}} \\
&\overset{(b)}{\geq} \inf_{\mathbf{x}\in\mathbb{R}^{Tt}} \frac{\|(\mathbf{B}+\mathbf{I})\mathbf{x} - \mathbf{b}_{T,c}(\mathbf{x})\|^2}{5\|\mathbf{x}\|^2}
\end{aligned} \tag{88}$$

where (a) follows from property 2 of Lemma 4, and (b) follows from $\mathbf{B}+\mathbf{I} \preceq 5\mathbf{I}$. Define

$$\tilde{\mu}(\mathbf{x}) = \frac{\|\nabla g_{T,t}(\mathbf{x})\|^2}{5\|\mathbf{x}\|^2} = \frac{\|(\mathbf{B}+\mathbf{I})\mathbf{x} - \mathbf{b}_{T,c}(\mathbf{x})\|^2}{5\|\mathbf{x}\|^2}. \tag{89}$$

We only need to give a lower bound on $\tilde{\mu}$. We estimate the proportion $\frac{\|(\mathbf{B}+\mathbf{I})\mathbf{x}-\mathbf{b}_{T,c}(\mathbf{x})\|}{\|\mathbf{x}\|}$ by computing each $\left|\frac{((\mathbf{B}+\mathbf{I})\mathbf{x})_k - b_{\tilde{\mathbf{y}}_k}(\tilde{\mathbf{x}}_k)}{\tilde{\mathbf{x}}_k}\right|$. If $\left|\frac{((\mathbf{B}+\mathbf{I})\mathbf{x})_k - b_{\tilde{\mathbf{y}}_k}(\tilde{\mathbf{x}}_k)}{\tilde{\mathbf{x}}_k}\right|$ is small, we will upper bound $\tilde{\mathbf{x}}_k^2$ by one of the nearby terms.

We define a function $\mathsf{Dominate} : \{1, \cdots, Tt\} \to \{\text{True}, \text{False}\}$ as follows:

$$\mathsf{Dominate}(k) = \begin{cases} \text{True}, & \text{If } \left|\frac{((\mathbf{B}+\mathbf{I})\mathbf{x})_k - b_{\tilde{\mathbf{y}}_k}(\tilde{\mathbf{x}}_k)}{\tilde{\mathbf{x}}_k}\right| > 0.19 \text{ or } \tilde{\mathbf{x}}_k = 0, \\ \text{False}, & \text{otherwise.} \end{cases} \tag{90}$$

Define $\tilde{\mathbf{z}}_k = \frac{\tilde{\mathbf{x}}_k}{\tilde{\mathbf{y}}_k}$, we present four auxiliary lemmas below. We prove them in Section C.3.

Lemma 9. *For $k = 1. \cdots .Tt$, if $\tilde{\mathbf{z}}_k \notin [\frac{31}{32}, \frac{33}{32}]$ and $\neg\mathsf{Dominate}(k)$, then $\frac{\tilde{\mathbf{z}}_{k-1}}{\tilde{\mathbf{z}}_k} \geq \frac{4}{3}$ or $\frac{\tilde{\mathbf{z}}_{k+1}}{\tilde{\mathbf{z}}_k} \geq \frac{4}{3}$.*

Lemma 10. *If $\tilde{\mathbf{z}}_{k-1} < \frac{5}{7}$, $\tilde{\mathbf{z}}_k \in [\frac{31}{32}, \frac{33}{32}]$, $\tilde{\mathbf{z}}_{k+1} \leq \frac{33}{32}$, then $\mathsf{Dominate}(k) = \text{True}$.*

Lemma 11. *If $\tilde{\mathbf{z}}_{k-2} < \frac{31}{32}$, $\frac{5}{7} \leq \tilde{\mathbf{z}}_{k-1} < \frac{31}{32}$, $\tilde{\mathbf{z}}_k \in [\frac{31}{32}, \frac{33}{32}]$, then $\mathsf{Dominate}(k-1) = \text{True}$.*

Lemma 12. *If $\tilde{\mathbf{z}}_n \in [\frac{31}{32}, \frac{33}{32}]$, $\neg\mathsf{Dominate}(n)$ and $\tilde{\mathbf{z}}_{n-1}, \tilde{\mathbf{z}}_{n+1} \leq \frac{33}{32}$, there exist k such that $k < n$, $\mathsf{Dominate}(k)$ and $\tilde{\mathbf{z}}_k > \frac{5}{7}$.*

Note that if we alter the "'$+$" and "$-$" sign in Lemma 10 or 11, the conclusion still holds.

Now we define an operator Next on indices on which $\mathsf{Dominate}$ operator is false. Intuitively, Next finds the direction in which $\tilde{\mathbf{z}}$ grows exponentially.

By Lemma 9, if $\tilde{\mathbf{z}}_k \notin [\frac{31}{32}, \frac{33}{32}]$ and $\neg\mathsf{Dominate}(k)$, define $\mathsf{Next}(k) \in \{k-1, k+1\}$ to be one of the coordinate satisfying $\frac{\tilde{\mathbf{z}}_{\mathsf{Next}(k)}}{\tilde{\mathbf{z}}_k} \geq \frac{4}{3}$.

Next, we define how the Next operator acts on the index n satisfying the condition below, which can be intuitively thought as an index in a chain of exponential growth. Without the loss of generality (if we alter the "$+$" and "$-$", the following conclusion still holds), if $\tilde{\mathbf{z}}_{n+1} > \frac{33}{32}$, we define $\mathsf{Next}(n) = n+1$. Therefore, if $\mathsf{Next}(n-1) = n$, $\tilde{\mathbf{z}}_n \in [\frac{31}{32}, \frac{33}{32}]$ and $\neg\mathsf{Dominate}(n)$, then by Lemma 11, we have $\tilde{\mathbf{z}}_{n+1} \geq \frac{33}{32}$. Therefore, if $n = \mathsf{Next}(m)$, then $\mathsf{Next}(n)$ is defined. We can apply the Next operator recursively, and will finally reach a index n such that $\mathsf{Dominate}(n) = \text{True}$. This process will terminate because $\tilde{\mathbf{z}}_0 = 0$ and $\tilde{\mathbf{z}}_{Tt+1} = \tilde{\mathbf{z}}_{Tt}$, ensuring that if the recursive Next operation reaches 0 or Tt, it terminates.

For other n such that $\neg\mathsf{Dominate}(n)$ and $\tilde{\mathbf{z}}_n \in [\frac{31}{32}, \frac{33}{32}]$ ($\tilde{\mathbf{z}}_{n-1}, \tilde{\mathbf{z}}_{n+1} \leq \frac{33}{32}$), the operator Next is undefined, and we will use Lemma 12 to tackle it

Algorithm 1 Algorithm to find $\mathsf{UBIndex}(n)$

Require: $\neg\mathsf{Dominate}(n)$
if $\mathsf{Next}(n)$ is defined **then**
  $m \leftarrow n$
  while $\neg\mathsf{Dominate}(m)$ **do**
    $m \leftarrow \mathsf{Next}(m)$ ▷ This process is well-defined and will terminate.
  end while
  $\mathsf{UBIndex}(n) \leftarrow m$
else if $\mathsf{Next}(n)$ is not defined **then**
  $\mathsf{UBIndex}(n) \leftarrow \operatorname{argmax}_{k<n} \{\tilde{\mathbf{z}}_k : \mathsf{Dominate}(k)\}$. ▷ $\mathsf{UBIndex}(n)$ exists and $\tilde{\mathbf{z}}_{\mathsf{UBIndex}(n)} > \frac{5}{7}$ (Lemma 12).
end if

For n such that $\neg\mathsf{Dominate}(n)$, we define an operator $\mathsf{UBIndex}$, and use a proportion of $\|g_{T,t}(\mathbf{x})\|_{\mathsf{UBIndex}(n)}$ to upper bound $\mathbf{x}_n$. The process of finding $\mathsf{UBIndex}$ is provided in Algorithm 1.

Define $\mathsf{Dist}(n) = |n - \mathsf{UBIndex}(n)|$, and define $\mathrm{UB}(k)$ to be a proportion of $\|g_{T,t}(\mathbf{x})\|_{\mathsf{UBIndex}(k)^2}$ as follows:

$$\mathrm{UB}(n) = \begin{cases} \frac{1}{4}|(\nabla g_{T,t}(\mathbf{x}))_n|^2, & \\ & \mathsf{Dominate}(n), \\ \frac{1}{4}\cdot\frac{13}{49}\cdot\left(\frac{36}{49}\right)^{\mathsf{Dist}(n)-1}|(\nabla g_{T,t}(\mathbf{x}))_{\mathsf{UBIndex}(n)}|^2, & \\ & \neg\mathsf{Dominate}(n), \mathsf{Next}(n) \text{ exists}, \\ \frac{1}{4T}\cdot\frac{15}{64}\cdot\frac{\tilde{\mathbf{y}}_n^2}{\tilde{\mathbf{y}}_{\mathsf{UBIndex}(n)}^2}|(\nabla g_{T,t}(\mathbf{x}))_{\mathsf{UBIndex}(n)}|^2, & \\ & \neg\mathsf{Dominate}(n), \mathsf{Next}(n) \text{ does not exist}. \end{cases} \tag{91}$$

Let $A = \{n : \mathsf{Dominate}(n)\}$, $B = \{n : \neg\mathsf{Dominate}(n), \mathsf{Next}(n) \text{ exists}\}$, and $C = \{n : \neg\mathsf{Dominate}(n), \mathsf{Next}(n) \text{ does not exist}\}$. Now we calculate $\sum_{i=1}^{Tt}\mathrm{UB}(n)$, and show that it is smaller than $\|g_{T,t}(\mathbf{x})\|^2$.

$$\begin{aligned} \sum_{i=1}^{Tt}\mathrm{UB}(n) &= \sum_{n\in A}\mathrm{UB}(n) + \sum_{n\in B}\mathrm{UB}(n) + \sum_{n\in C}\mathrm{UB}(n) \\ &= \frac{1}{4}\sum_{n\in A}|(\nabla g_{T,t}(\mathbf{x}))_n|^2 \\ &\quad + \sum_{n\in B}\frac{13}{196}\left(\frac{36}{49}\right)^{\mathsf{Dist}(n)-1}|(\nabla g_{T,t}(\mathbf{x}))_{\mathsf{UBIndex}}|^2 \\ &\quad + \sum_{n\in C}\frac{15}{256T}\left(\frac{\tilde{\mathbf{y}}_n}{\tilde{\mathbf{y}}_{\mathsf{UBIndex}(n)}}\right)^2|(\nabla g_{T,t}(\mathbf{x}))_{\mathsf{UBIndex}}|^2. \end{aligned} \tag{92}$$

By alternating the order of summation, we have

$$\begin{aligned} &\sum_{n\in B}\frac{13}{196}\left(\frac{36}{49}\right)^{\mathsf{Dist}(n)-1}|(\nabla g_{T,t}(\mathbf{x}))_{\mathsf{UBIndex}}|^2 \\ \le\ &\frac{13}{196}\sum_{n\in A}2\cdot\sum_{i=1}^{\infty}\left(\frac{36}{49}\right)^{i-1}|(\nabla g_{T,t}(\mathbf{x}))_n|^2 \\ =\ &\frac{1}{2}\sum_{n\in A}|(\nabla g_{T,t}(\mathbf{x}))_n|^2, \end{aligned} \tag{93}$$

and

$$
\begin{aligned}
&\sum_{n\in C}\frac{15}{256T}\left(\frac{\tilde{\mathbf{y}}_n}{\tilde{\mathbf{y}}_{\mathsf{UBIndex}(n)}}\right)^2|(\nabla g_{T,t}(\mathbf{x}))_{\mathsf{UBIndex}}|^2\\
\le&\frac{15}{256T}\sum_{n\in A}\sum_{m\ge n}\left(\frac{\tilde{\mathbf{y}}_m}{\tilde{\mathbf{y}}_n}\right)^2|(\nabla g_{T,t}(\mathbf{x}))_n|^2\\
=&\frac{15}{256T}\sum_{n\in A}T\sum_{i=0}^{\infty}\left(\frac{7}{8}\right)^i|(\nabla g_{T,t}(\mathbf{x}))_n|^2\\
=&\frac{1}{4}\sum_{n\in A}|(\nabla g(\mathbf{x}))_n|^2.
\end{aligned}
\tag{94}
$$

Summing up (92), (93) and (94), we have

$$
\begin{aligned}
\sum_{i=1}^{Tt}\mathrm{UB}(n)&=\frac{1}{4}\sum_{n\in A}|(\nabla g_{T,t}(\mathbf{x}))_n|^2+\frac{1}{2}\sum_{n\in A}|(\nabla g_{T,t}(\mathbf{x}))_n|^2+\frac{1}{4}\sum_{n\in A}|(\nabla g_{T,t}(\mathbf{x}))_n|^2\\
&=\sum_{n\in A}|(\nabla g_{T,t}(\mathbf{x}))_n|^2\\
&\le\|\nabla g_{T,t}(\mathbf{x})\|^2.
\end{aligned}
\tag{95}
$$

Finally, we calculate $\frac{\mathrm{UB}(n)}{\mathbf{x}_n^2}$ to give an universal lower bound of the PL constant. For $n\in A$, we have

$$
\begin{aligned}
\frac{\mathrm{UB}(n)}{\mathbf{x}_n^2}&=\frac{1}{4}\cdot\frac{|(\nabla g_{T,t}(\mathbf{x}))_n|^2}{\mathbf{x}_n^2}\\
&\overset{(a)}{\ge}\frac{361}{40000},
\end{aligned}
\tag{96}
$$

where (a) is due to $\mathsf{Dominate}(n)=\mathrm{True}$.

For $n\in B$, by Lemma 9 we have $\left|\frac{\tilde{\mathbf{z}}_{\mathsf{Next}(m)}}{\tilde{\mathbf{z}}_m}\right|\ge\frac{4}{3}$ for $m=n,\mathsf{Next}(n),\mathsf{Next}(\mathsf{Next}(n))\cdots$, with one possible exception when $\tilde{\mathbf{z}}_m\in[\frac{31}{32},\frac{33}{32}]$, in which case there is $|\tilde{\mathbf{z}}_{\mathsf{Next}(m)}|\ge|\tilde{\mathbf{z}}_m|$. Therefore, $|\tilde{\mathbf{z}}_{\mathsf{UBIndex}(n)}|\ge\left(\frac{4}{3}\right)^{\mathsf{Dist}(n)-1}|\tilde{\mathbf{z}}_n|$, so we have

$$
\begin{aligned}
\frac{\mathrm{UB}(n)}{\mathbf{x}_n^2}&=\frac{1}{4}\cdot\frac{13}{49}\cdot\left(\frac{36}{49}\right)^{\mathsf{Dist}(n)-1}|(\nabla g_{T,t}(\mathbf{x}))_{\mathsf{UBIndex}(n)}|^2\\
&\overset{(a)}{\ge}\frac{13}{196}\cdot\left(\frac{36}{49}\right)^{\mathsf{Dist}(n)-1}\frac{361}{10000}\frac{\mathbf{x}_{\mathsf{UBIndex}(n)}^2}{\mathbf{x}_n^2}\\
&\overset{(b)}{=}\frac{4693}{1960000}\left(\frac{36}{49}\right)^{\mathsf{Dist}(n)-1}\cdot\left(\frac{\tilde{\mathbf{z}}_{\mathsf{UBIndex}(n)}}{\tilde{\mathbf{z}}_n}\right)^2\cdot\left(\frac{\tilde{\mathbf{y}}_{\mathsf{UBIndex}(n)}}{\tilde{\mathbf{y}}_n}\right)^2\\
&\overset{(c)}{\ge}\frac{4693}{1960000}\left(\frac{36}{49}\right)^{\mathsf{Dist}(n)-1}\cdot\left(\frac{16}{9}\right)^{\mathsf{Dist}(n)-1}\cdot\left(\frac{49}{64}\right)^{\mathsf{Dist}(n)}\\
&=\frac{4693}{2560000},
\end{aligned}
\tag{97}
$$

where (a) is due to $\mathsf{Dominate}(\mathsf{UBIndex}(n))=\mathrm{True}$, (b) is due to $\tilde{\mathbf{z}}_m=\frac{\tilde{\mathbf{x}}_m}{\tilde{\mathbf{y}}_m}$, (c) is due to $|\tilde{\mathbf{z}}_{\mathsf{UBIndex}(n)}|\ge\left(\frac{4}{3}\right)^{\mathsf{Dist}(n)-1}|\tilde{\mathbf{z}}_n|$ and $\frac{|\tilde{\mathbf{y}}_{m+1}|}{|\tilde{\mathbf{y}}_m|}\ge\frac{7}{8}$.

For $n \in C$, we have

$$\begin{aligned}
\frac{\mathrm{UB}(n)}{\mathbf{x}_n^2} &= \frac{1}{4T} \cdot \frac{15}{64} \cdot \frac{\tilde{\mathbf{y}}_n^2}{\tilde{\mathbf{y}}_{\mathsf{UBIndex}(n)}^2} \cdot \frac{|(\nabla g_{T,t}(\mathbf{x}))_{\mathsf{UBIndex}(n)}|^2}{\mathbf{x}_n^2} \\
&\overset{(a)}{\geq} \frac{15}{256T} \cdot \frac{\tilde{\mathbf{y}}_n^2}{\tilde{\mathbf{y}}_{\mathsf{UBIndex}(n)}^2} \cdot \frac{361\tilde{\mathbf{x}}_{\mathsf{UBIndex}(n)}^2}{10000\tilde{\mathbf{x}}_n^2} \\
&\overset{(b)}{\geq} \frac{1083}{512000T} \frac{\tilde{\mathbf{z}}_{\mathsf{UBIndex}(n)}^2}{\tilde{\mathbf{z}}_n^2} \\
&\overset{(c)}{\geq} \frac{1083}{512000T} \left(\frac{5/7}{33/32}\right)^2 \\
&= \frac{1083}{1067220T},
\end{aligned} \tag{98}$$

where (a) is due to $\mathsf{Dominate}(\mathsf{UBIndex}(n)) = \text{True}$, (b) is due to $\tilde{\mathbf{z}}_m = \frac{\tilde{\mathbf{x}}_m}{\tilde{\mathbf{y}}_m}$, and (c) is due to Lemma 12 and $\tilde{\mathbf{z}}_n \in [\frac{31}{32}, \frac{33}{32}]$. Therefore,

$$\begin{aligned}
\mu &\geq \inf_{\mathbf{x}} \frac{\|\nabla g_{T,t}(\mathbf{x})\|^2}{5\mathbf{x}^2} \\
&\geq \inf_{\mathbf{x}} \frac{\sum_{n=1}^{Tt} \mathrm{UB}(n)}{5\sum_{n=1}^{Tt} \mathbf{x}_n^2} \\
&\geq \inf_{\mathbf{x}} \min_n \frac{\mathrm{UB}(n)}{5\mathbf{x}_n^2} \\
&\geq \inf_{\mathbf{x}} \frac{1}{5} \cdot \min\left\{\frac{361}{40000}, \frac{4693}{2560000}, \frac{1083}{1067220T}\right\} \\
&= \frac{1083}{1067220T}.
\end{aligned} \tag{99}$$

□

C.3 Proof of Lemma 9, 10, 11 and 12

Lemma 9. *For $k = 1. \cdots .Tt$, if $\tilde{\mathbf{z}}_k \notin [\frac{31}{32}, \frac{33}{32}]$ and $\neg\mathsf{Dominate}(k)$, then $\frac{\tilde{\mathbf{z}}_{k-1}}{\tilde{\mathbf{z}}_k} \geq \frac{4}{3}$ or $\frac{\tilde{\mathbf{z}}_{k+1}}{\tilde{\mathbf{z}}_k} \geq \frac{4}{3}$.*

Proof. If $\tilde{\mathbf{z}}_k \notin [\frac{31}{32}, \frac{33}{32}]$, then $b_{\tilde{\mathbf{y}}_k}(\tilde{\mathbf{x}}_k) = 0$, and $(\nabla g_{T,t}(\mathbf{x}))_k = (\tilde{\mathbf{B}}\tilde{\mathbf{x}})_k + \tilde{\mathbf{x}}_k$. By $\neg\mathsf{Dominate}(k)$, we have $\left|(\tilde{\mathbf{B}}\tilde{\mathbf{x}})_k + \tilde{\mathbf{x}}_k\right| < 0.19\,|\tilde{\mathbf{x}}_k|$. We consider three cases:

1. If $k = T+1, 2T+1, \cdots, T(t-1)+1$, $(\tilde{\mathbf{B}}\tilde{\mathbf{x}})_k = -\frac{7}{8}\tilde{\mathbf{x}}_{k-1} + 2\tilde{\mathbf{x}}_k - \tilde{\mathbf{x}}_{k+1}$, $(\tilde{\mathbf{B}}\tilde{\mathbf{x}})_k + \tilde{\mathbf{x}}_k = \frac{7}{8}\tilde{\mathbf{x}}_{k-1} + 3\tilde{\mathbf{x}}_k - \tilde{\mathbf{x}}_{k+1}$, and $\frac{7}{8}\tilde{\mathbf{y}}_{k-1} = \tilde{\mathbf{y}}_k = \tilde{\mathbf{y}}_{k+1}$. We have

$$\begin{aligned}
3|\tilde{\mathbf{x}}_k| - \frac{7}{8}|\tilde{\mathbf{x}}_{k-1}| - |\tilde{\mathbf{x}}_{k+1}| &\leq |(\tilde{\mathbf{B}}\tilde{\mathbf{x}})_k + \tilde{\mathbf{x}}_k| \\
&< 0.19|\tilde{\mathbf{x}}_k|.
\end{aligned} \tag{100}$$

   Dividing both sides by $\tilde{\mathbf{y}}_k$, we have

$$3|\tilde{\mathbf{z}}_k| - |\tilde{\mathbf{z}}_{k-1}| - |\tilde{\mathbf{z}}_{k+1}| < 0.19|\tilde{\mathbf{z}}_k|. \tag{101}$$

   Thus, we have $|\tilde{\mathbf{z}}_{k-1}| + \tilde{\mathbf{z}}_{k+1} > \frac{8}{3}|\tilde{\mathbf{z}}_k|$, which indicates that $\frac{\tilde{\mathbf{z}}_{k-1}}{\tilde{\mathbf{z}}_k} \geq \frac{4}{3}$ or $\frac{\tilde{\mathbf{z}}_{k+1}}{\tilde{\mathbf{z}}_k} \geq \frac{4}{3}$.

2. If $k = T, 2T, \cdots, T(t-1)$, $(\tilde{\mathbf{B}}\tilde{\mathbf{x}})_k + \tilde{\mathbf{x}}_k = -\tilde{\mathbf{x}}_{k-1} + \frac{177}{64}\tilde{\mathbf{x}}_k - \frac{7}{8}\tilde{\mathbf{x}}_{k+1}$, and $\tilde{\mathbf{y}}_{k-1} = \tilde{\mathbf{y}}_k = \frac{8}{7}\tilde{\mathbf{y}}_{k+1}$. Thus, we have

$$\frac{177}{64}|\tilde{\mathbf{x}}_k| - |\tilde{\mathbf{x}}_{k-1}| - \frac{7}{8}|\tilde{\mathbf{x}}_{k+1}| < 0.19|\tilde{\mathbf{x}}_k|. \tag{102}$$

Dividing both sides by $\tilde{\mathbf{y}}_k$, we have

$$\frac{177}{64}|\tilde{\mathbf{z}}_k| - |\tilde{\mathbf{z}}_{k-1}| - \frac{49}{64}|\tilde{\mathbf{z}}_{k+1}| < 0.19|\tilde{\mathbf{z}}_k|. \tag{103}$$

Thus, we have

$$|\tilde{\mathbf{z}}_{k-1}| + \frac{49}{64}|\tilde{\mathbf{z}}_{k+1}| > \left(\frac{177}{64} - 0.19\right)|\tilde{\mathbf{z}}_k| > \frac{4}{3}\cdot\left(1 + \frac{49}{64}\right)|\tilde{\mathbf{z}}_k|. \tag{104}$$

3. For other k, $(\tilde{\mathbf{B}}\tilde{\mathbf{x}})_k + \tilde{\mathbf{x}}_k = -\tilde{\mathbf{x}}_{k-1} + 3\tilde{\mathbf{x}}_k - \tilde{\mathbf{x}}_{k+1}$, and $\tilde{\mathbf{y}}_{k-1} = \tilde{\mathbf{y}}_k = \tilde{\mathbf{y}}_{k+1}$. Therefore, we have (by dividing $\tilde{\mathbf{y}}_k$ to both sides of (101))

$$|\tilde{\mathbf{z}}_{k-1}| + |\tilde{\mathbf{z}}_{k+1}| > 2.81|\tilde{\mathbf{z}}_k| > \frac{8}{3}|\tilde{\mathbf{z}}_k|. \tag{105}$$

□

Lemma 10. *If* $\tilde{\mathbf{z}}_{k-1} < \frac{5}{7}$, $\tilde{\mathbf{z}}_k \in [\frac{31}{32}, \frac{33}{32}]$, $\tilde{\mathbf{z}}_{k+1} \leq \frac{33}{32}$, *then* $\mathsf{Dominate}(k) = \text{True}$.

Proof. For any $y > 0$, we have $0 \leq b_y(x) \leq y$. Therefore, $(\nabla g_{T,t}(\mathbf{x}))_k \geq (\tilde{\mathbf{B}}\tilde{\mathbf{x}})_k$.

1. If $k = T+1, 2T+1, \cdots, T(t-1)+1$, $(\tilde{\mathbf{B}}\tilde{\mathbf{x}})_k = -\frac{7}{8}\tilde{\mathbf{x}}_{k-1} + 2\tilde{\mathbf{x}}_k - \tilde{\mathbf{x}}_{k+1}$ and $\frac{7}{8}\tilde{\mathbf{y}}_{k-1} = \tilde{\mathbf{y}}_k = \tilde{\mathbf{y}}_{k+1}$.

$$\begin{aligned}
\frac{(\nabla g_{T,t}(\mathbf{x}))_k}{\tilde{\mathbf{x}}_k} &\geq \frac{(\tilde{\mathbf{B}}\tilde{\mathbf{x}})_k}{\tilde{\mathbf{x}}_k} \\
&= \frac{-\frac{7}{8}\tilde{\mathbf{x}}_{k-1} + 2\tilde{\mathbf{x}}_k = \tilde{\mathbf{x}}_{k+1}}{\tilde{\mathbf{x}}_k} \\
&\overset{(a)}{=} \frac{-\tilde{\mathbf{z}}_{k-1} + \tilde{\mathbf{z}}_k - \tilde{\mathbf{z}}_{k+1}}{\tilde{\mathbf{z}}_k} \\
&\overset{(b)}{>} -\frac{5/7}{31/32} + 2 - \frac{33/32}{31/32} \\
&> 0.19,
\end{aligned} \tag{106}$$

where (a) holds by dividing $\tilde{\mathbf{y}}_k$ to the numerator and denominator, (b) holds by the assumptions on $\tilde{\mathbf{z}}_{k-1}, \tilde{\mathbf{z}}_k, \tilde{\mathbf{z}}_{k+1}$.

2. If $k = T, 2T, \cdots, T(t-1)$, $(\tilde{\mathbf{B}}\tilde{\mathbf{x}})_k = -\tilde{\mathbf{x}}_{k-1} + \frac{113}{64}\tilde{\mathbf{x}}_k - \frac{7}{8}\tilde{\mathbf{x}}_{k+1}$, and $\tilde{\mathbf{y}}_{k-1} = \tilde{\mathbf{y}}_k = \frac{8}{7}\tilde{\mathbf{y}}_{k+1}$.

$$\begin{aligned}
\frac{(\nabla g_{T,t}(\mathbf{x}))_k}{\tilde{\mathbf{x}}_k} &\geq \frac{(\tilde{\mathbf{B}}\tilde{\mathbf{x}})_k}{\tilde{\mathbf{x}}_k} \\
&= \frac{-\tilde{\mathbf{x}}_{k-1} + \frac{113}{64}\tilde{\mathbf{x}}_k - \frac{7}{8}\tilde{\mathbf{x}}_{k+1}}{\tilde{\mathbf{x}}_k} \\
&\overset{(a)}{=} \frac{-\tilde{\mathbf{z}}_{k-1} + \frac{113}{64}\tilde{\mathbf{z}}_k - \frac{49}{64}\tilde{\mathbf{z}}_{k+1}}{\tilde{\mathbf{z}}_k} \\
&\overset{(b)}{>} -\frac{5/7}{31/32} + \frac{113}{64} - \frac{49}{64}\cdot\frac{33/32}{31/32} \\
&> 0.19,
\end{aligned} \tag{107}$$

where (a) holds by dividing $\tilde{\mathbf{y}}_k$ to the numerator and denominator, and (b) holds by the assumptions on $\tilde{\mathbf{z}}_{k-1}, \tilde{\mathbf{z}}_k$, and $\tilde{\mathbf{z}}_{k+1}$.

3. For other k, $(\tilde{\mathbf{B}}\tilde{\mathbf{x}})_k = -\tilde{\mathbf{x}}_{k-1} + 2\tilde{\mathbf{x}}_k - \tilde{\mathbf{x}}_{k+1}$, and $\tilde{\mathbf{y}}_{k-1} = \tilde{\mathbf{y}}_k = \tilde{\mathbf{y}}_{k+1}$. Therefore,

$$\begin{aligned}\frac{(\nabla g_{T,t}(\mathbf{x}))_k}{\tilde{\mathbf{x}}_k} &\geq \frac{(\tilde{\mathbf{B}}\tilde{\mathbf{x}})_k}{\tilde{\mathbf{x}}_k} \\ &= \frac{-\tilde{\mathbf{x}}_{k-1} + 2\tilde{\mathbf{x}}_k = \tilde{\mathbf{x}}_{k+1}}{\tilde{\mathbf{x}}_k} \\ &\overset{(a)}{=} \frac{-\tilde{\mathbf{z}}_{k-1} + \tilde{\mathbf{z}}_k - \tilde{\mathbf{z}}_{k+1}}{\tilde{\mathbf{z}}_k} \\ &\overset{(b)}{>} -\frac{5/7}{31/32} + 2 - \frac{33/32}{31/32} \\ &> 0.19,\end{aligned} \tag{108}$$

where (a) holds by dividing $\tilde{\mathbf{y}}_k$ to the numerator and denominator, and (b) holds by the assumptions on $\tilde{\mathbf{z}}_{k-1}, \tilde{\mathbf{z}}_k$, and $\tilde{\mathbf{z}}_{k+1}$.

□

Lemma 11. *If* $\tilde{\mathbf{z}}_{k-2} < \frac{31}{32}$, $\frac{5}{7} \leq \tilde{\mathbf{z}}_{k-1} < \frac{31}{32}$, $\tilde{\mathbf{z}}_k \in [\frac{31}{32}, \frac{33}{32}]$, *then* $\mathsf{Dominate}(k-1) = \text{True}$.

Proof. Like the proof of Lemma 10, we directly compute $\frac{(\nabla g_{T,t}(\mathbf{x}))_{k-1}}{\tilde{\mathbf{x}}_{k-1}}$. Because $\tilde{\mathbf{z}}_{k-1} < \frac{31}{32}$, $b_{\tilde{\mathbf{y}}_{k-1}}(\tilde{\mathbf{x}}_{k-1}) = 0$.

$$\begin{aligned}\frac{(\nabla g_{T,t}(\tilde{\mathbf{x}}))_{k-1}}{\tilde{\mathbf{x}}_{k-1}} &= \frac{(\tilde{\mathbf{B}}\tilde{\mathbf{x}})_{k-1} + \tilde{\mathbf{x}}_{k-1}}{\tilde{\mathbf{x}}_{k-1}} \\ &\overset{(a)}{=} -a\frac{\tilde{\mathbf{z}}_{k-2}}{\tilde{\mathbf{z}}_{k-1}} + b - c\frac{\tilde{\mathbf{z}}_k}{\tilde{\mathbf{z}}_{k-1}},\end{aligned} \tag{109}$$

where (a) holds by dividing $\tilde{\mathbf{y}}_k$ to the numerator and denominator. For $k-1 = T, 2T, \cdots, T(t-1)$, $a = -1$, $b - \frac{177}{64}$ and $c = \frac{49}{64}$. For other k, $a = c = 1$ and $b = 3$. By the assumptions on $\tilde{\mathbf{z}}_{k-2}, \tilde{\mathbf{z}}_{k-1}$ and $\tilde{\mathbf{z}}_k$, we have $\frac{\tilde{\mathbf{z}}_{k-2}}{\tilde{\mathbf{z}}_{k-1}} < \frac{31/32}{5/7}$, and $\frac{\tilde{\mathbf{z}}_k}{\tilde{\mathbf{z}}_{k-1}} < \frac{33/32}{5/7}$. Plugging everything into (109), we have $\frac{(\nabla g_{T,t}(\tilde{\mathbf{x}}))_{k-1}}{\tilde{\mathbf{x}}_{k-1}} > 0.2$. □

Lemma 12. *If* $\tilde{\mathbf{z}}_n \in [\frac{31}{32}, \frac{33}{32}]$, $\neg\mathsf{Dominate}(n)$ *and* $\tilde{\mathbf{z}}_{n-1}, \tilde{\mathbf{z}}_{n+1} \leq \frac{33}{32}$, *there exist* k *such that* $k < n$, $\mathsf{Dominate}(k)$ *and* $\tilde{\mathbf{z}}_k > \frac{5}{7}$.

Proof. Define $m = \operatorname{argmax}_{m'} \left\{\tilde{\mathbf{z}}_{m'} \notin [\frac{31}{32}, \frac{33}{32}] : m' < n\right\}$. If $\tilde{\mathbf{z}}_m > \frac{33}{32}$, let k be m if $\mathsf{Dominate}(m)$ and $\mathsf{UBIndex}(m)$ if $\neg\mathsf{Dominate}(m)$. If $\tilde{\mathbf{z}}_m \in [\frac{5}{7}, \frac{33}{32}]$, let k be m if $\mathsf{Dominate}(m)$ and $\mathsf{UBIndex}(m)$ if $\neg\mathsf{Dominate}(m)$. Finally, if $\tilde{\mathbf{z}}_m < \frac{5}{7}$, let $k = m + 1$. □

C.4 Proof of Theorem 4

Theorem 4. *Given* $L \geq \mu > 0$. *When* $\kappa = \frac{L}{\mu} > C_4$ *where* C_4 *is a universal constant, there exists* T *and* t *such that* $\tilde{g}$ *is* L*-smooth and* μ*-PL. Moreover, any first-order zero-respecting algorithm with* $\mathbf{x}^{(0)} = \mathbf{0}$ *needs at least* $\Omega\left(\kappa \log \frac{1}{\varepsilon}\right)$ *gradient computations to find a point* $\mathbf{x}$ *satisfying* $\tilde{g}(\mathbf{x}) - \tilde{g}^* \leq \varepsilon(\tilde{g}(\mathbf{x}^{(0)}) - \tilde{g}^*)$.

Proof. We let $C_4 = 370C_3$. Given $\kappa > C_4$, we let

$$T = \left\lfloor \frac{\kappa}{37C_3} \right\rfloor, \tag{110}$$

and

$$t = 2\left\lfloor \log_{\frac{8}{7}} \frac{3}{2\varepsilon} \right\rfloor. \tag{111}$$

We use Lemma 4 to calculate the smoothness constant and PL constant of $\tilde{f}$. The smoothness constant of $\tilde{f}$ is:

$$\frac{LT^{-1}D^2}{37} \cdot TD^{-2} \cdot 37 = L, \tag{112}$$

and the PL constant of $\tilde{f}$ is:

$$\begin{aligned} \frac{LT^{-1}D^2}{37} \cdot TD^{-2} \cdot \frac{1}{C_3 T} &= \frac{L}{37C_3 T} \\ &\overset{a}{=} \frac{L}{37C_3} \cdot \frac{1}{\left\lfloor \frac{\kappa}{37C_3} \right\rfloor} \\ &\geq \frac{L}{\kappa} \\ &= \mu, \end{aligned} \tag{113}$$

where $\overset{a}{=}$ uses (110).

Finally, by Lemma 2, any first-order zero-respecting algorithm needs at least $\frac{Tt}{2}$ accesses to gradient to find $\mathbf{x}$ such that $\tilde{g}(x) - \tilde{g}^* \leq \frac{1}{16}\left(\tilde{g}(\mathbf{x}^{(0)}) - \tilde{g}^*\right)$. By (110) and (111),

$$\frac{Tt}{2} \geq \left(\frac{\kappa}{37C_1} - 1\right) \cdot \left(\log_{\frac{8}{7}} \frac{3}{2\varepsilon}\right) = \Omega\left(\kappa \log \frac{1}{\varepsilon}\right), \tag{114}$$

which completes the proof. □