

# On Finding Small Hyper-Gradients in Bilevel Optimization: Hardness Results and Improved Analysis

**Lesi Chen\***

*IIIS, Tsinghua University  
Shanghai Qizhi Institute*

CHENLC23@MAILS.TSINGHUA.EDU.CN

**Jing Xu\***

*IIIS, Tsinghua University*

XUJING21@MAILS.TSINGHUA.EDU.CN

**Jingzhao Zhang<sup>†</sup>**

*IIIS, Tsinghua University  
Shanghai AI Lab  
Shanghai Qizhi Institute*

JINGZHAOZ@MAIL.TSINGHUA.EDU.CN

**Editors:** Shipra Agrawal and Aaron Roth

## Abstract

Bilevel optimization reveals the inner structure of otherwise oblique optimization problems, such as hyperparameter tuning, neural architecture search, and meta-learning. A common goal in bilevel optimization is to minimize a hyper-objective that implicitly depends on the solution set of the lower-level function. Although this hyper-objective approach is widely used, its theoretical properties have not been thoroughly investigated in cases where *the lower-level functions lack strong convexity*. In this work, we first provide hardness results to show that the goal of finding stationary points of the hyper-objective for nonconvex-convex bilevel optimization can be intractable for zero-respecting algorithms. Then we study a class of tractable nonconvex-nonconvex bilevel problems when the lower-level function satisfies the Polyak-Łojasiewicz (PL) condition. We show a simple first-order algorithm can achieve complexity bounds of  $\tilde{O}(\epsilon^{-2})$ ,  $\tilde{O}(\epsilon^{-4})$  and  $\tilde{O}(\epsilon^{-6})$  in the deterministic, partially stochastic, and fully stochastic setting respectively.

**Keywords:** bilevel optimization, optimization theory, oracle complexity

## 1. Introduction

The goal of bilevel optimization is to minimize the upper-level function  $f(x, y)$  under the constraint that  $y$  is minimized with respect to the lower-level function  $g(x, y)$ . Formally, it is defined as,

$$\min_{x \in \mathbb{R}^{d_x}, y \in Y^*(x)} f(x, y), \quad Y^*(x) = \arg \min_{y \in \mathbb{R}^{d_y}} g(x, y). \quad (1)$$

Bilevel optimization in this form has received increasing attention due to its wide applications in many machine learning problems, including hyperparameter tuning (Franceschi et al., 2018; Pedregosa, 2016), neural architecture search (Liu et al., 2019; Wang et al., 2022; Zoph and Le, 2016; Zhang et al., 2021), meta-learning (Franceschi et al., 2018; Hospedales et al., 2021; Ravi and Larochelle, 2017; Pham et al., 2021; Rajeswaran et al., 2019), out-of-distribution learning (Zhou et al., 2022), adversarial training (Goodfellow et al., 2020; Sinha et al., 2018; Lin et al., 2020a,b),

---

\*. Equal contributions.

<sup>†</sup>. The corresponding author.

reinforcement learning (Konda and Tsitsiklis, 1999; Hong et al., 2023), causal learning (Jiang and Veitch, 2022; Arjovsky et al., 2019).

The hyper-objective approach (Dempe, 2002) reformulates Problem (1) by a minimization problem defined below,

$$\min_{x \in \mathbb{R}^d} \varphi(x), \text{ where } \varphi(x) = \min_{y \in Y^*(x)} f(x, y) \quad (2)$$

is called the hyper-objective. When  $\varphi(x)$  has Lipschitz continuous gradients, a common is to find almost stationary points of  $\varphi(x)$ .

Finding stationary points can be especially easy when the lower-level function is strongly convex, because Equation (2) can be simplified to the composite optimization problem below. Specifically, since  $Y^*(x)$  has only one element when the lower-level function is strongly convex, we have  $Y^*(x) = \{y^*(x)\}$  and

$$\min_{x \in \mathbb{R}^{d_x}} \varphi(x) := f(x, y^*(x)), \text{ where } y^*(x) = \arg \min_{y \in \mathbb{R}^{d_y}} g(x, y). \quad (3)$$

Further, the implicit function theorem (Dontchev et al., 2009) implies

$$\nabla \varphi(x) = \underbrace{\nabla_x f(x, y^*(x))}_{\text{Explicit Gradient}} + \underbrace{(\nabla y^*(x))^\top \nabla_y f(x, y^*(x))}_{\text{Implicit Gradient}}. \quad (4)$$

This equation enables one to estimate the hyper-gradient  $\nabla \varphi(x)$  and perform gradient descent on  $\varphi(x)$ . AID (Ghadimi and Wang, 2018) and ITD (Ji et al., 2021) estimate  $\nabla \varphi(x)$  with Hessian-vector-product oracles. The more recently proposed F<sup>2</sup>BA (Chen et al., 2023) estimate  $\nabla \varphi(x)$  with gradient oracles. All these methods can find a stationary point of  $\varphi(x)$ .

However, due to the prevalence of nonconvex functions in real-world scenarios, the strong convexity assumption may limit the applicability of algorithms. Therefore, our work first aims to whether convexity, as a relaxation of strong convexity, suffices for finding small hyper-gradients efficiently. Specifically, we hope to answer the question below.

*Can we find stationary points of  $\varphi(x)$  when the lower level function  $g(x, y)$  is (strictly) convex but not strongly convex in  $y$ ?*

We provide a negative answer to the above question. We first prove in Example 3.1 that in the case where  $g(x, \cdot)$  is merely convex,  $\varphi(x)$  may not have stationary points since  $\varphi(x)$  may be discontinuous. Furthermore, we show that the continuity of  $\varphi(x)$  can be fully characterized by the Pompeiu–Hausdorff continuity of  $Y^*(x)$  in Theorem 3.1.

We then study the cases when  $\nabla \varphi(x)$  exists, *e.g.* when the lower-level function is strictly convex. We demonstrate that the stationary points of  $\varphi(x)$  may still be computationally hard for any zero-respecting algorithms. This algorithm class contains a broad range of existing algorithms, including all the algorithms mentioned in this paper.

**Theorem 1.1 (Informal version of Theorem 3.2)** *There exists a bilevel problem instance whose lower-level function is strictly convex and both  $f, g$  satisfy regular smoothness conditions, such that any zero-respecting algorithm gets stuck at the initialization  $x_0$ .*

Table 1: We present the complexities of different methods for nonconvex-PL bilevel problems.

Oracle	Method	Deterministic	Partially Stochastic	Fully Stochastic	Reference
2nd	GALET <sup>(a)</sup>	$\tilde{\mathcal{O}}(\kappa^5 \epsilon^{-2})$	-	-	Xiao et al. (2023)
1st	Prox-F <sup>2</sup> BA <sup>(b)</sup>	$\tilde{\mathcal{O}}(\kappa^{p_1} \epsilon^{-3})$	$\tilde{\mathcal{O}}(\kappa^{p_2} \epsilon^{-5})$	$\tilde{\mathcal{O}}(\kappa^{p_3} \epsilon^{-7})$	Kwon et al. (2024b)
1st	F <sup>2</sup> BA	$\tilde{\mathcal{O}}(\kappa^4 \epsilon^{-2})$	$\tilde{\mathcal{O}}(\kappa^5 \epsilon^{-4})$	$\tilde{\mathcal{O}}(\kappa^{11} \epsilon^{-6})$	This Paper

(a) Although Xiao et al. (2023) did not provide the dependency on  $\kappa$  in the complexity, we can calculate by the way in our Remark C.1. Their analysis additionally requires the smallest singular value of  $\nabla_{yy}^2 g(x, y)$  has a constant gap between zero, but such assumption seems to be redundant by our Remark C.2.

(b) We use  $p_1, p_2, p_3$  to denote the polynomial dependency in  $\kappa$  since they are not provided by Kwon et al. (2024b).

Given the negative result, we aim to study intermediate function classes that are not strongly convex but computationally tractable. In particular, we study the cases when the lower-level function satisfies the Polyak-Łojasiewicz (PL) conditions. This condition allows global nonconvexity but ensures local strong convexity uniformly in a subspace.

The PL condition can pose nontrivial challenges since we have neither Equation (3) nor (4) in this case. Researchers have provided novel analyses in this case. Xiao et al. (2023) proposed a Hessian-vector-product-based method GALET with non-asymptotic convergence to KKT points of the gradient-based reformulated problem (Equation (11)) when  $g$  is PL in  $y$ . Kwon et al. (2024b) proved the differentiability of  $\varphi(x)$  when the penalty function  $\sigma f + g$  is uniformly PL in  $y$  for all  $\sigma$  in the neighborhood of zero, and showed a proximal variant of F<sup>2</sup>BA, which we call Prox-F<sup>2</sup>BA, can find a stationary point of  $\varphi(x)$ . Based on their differentiability result, we show that GALET also converges to a stationary point of  $\varphi(x)$  in Appendix C.

Although Prox-F<sup>2</sup>BA (Kwon et al., 2024b) has been shown to converge to an  $\epsilon$ -stationary point of  $\varphi(x)$  with  $\tilde{\mathcal{O}}(\epsilon^{-3})$  first-order oracle calls, the rate is worse than the  $\mathcal{O}(\epsilon^{-2})$  rate of gradient descent on nonconvex single-level optimization. Therefore, listed as an important future direction, Kwon et al. (2024b) asked the question below.

*Can one achieve the  $\mathcal{O}(\epsilon^{-2})$  rate for nonconvex-PL bilevel problems with gradient oracles?*

We give a positive answer to this question. Interestingly, we prove that the original F<sup>2</sup>BA (Chen et al., 2023) can already achieve this fast convergence rate. Our improvement over Kwon et al. (2024b) comes from establishing a tighter bound on the smoothness constant of  $\varphi(x)$ .

**Theorem 1.2 (Informal version of Theorem 4.1)** *Under regular conditions as in Kwon et al. (2024b), F<sup>2</sup>BA can provably find an  $\epsilon$ -stationary point of  $\varphi(x)$  with  $\tilde{\mathcal{O}}(\epsilon^{-2})$  first-order oracle calls for nonconvex-PL bilevel problems.*

We also extend our analysis to stochastic bilevel problems when an algorithm only has access to a noisy estimator of  $\nabla f$  and  $\nabla g$ . Under the partially stochastic setting when the noise is only in  $\nabla g$ , we prove the stochastic F<sup>2</sup>BA has the  $\tilde{\mathcal{O}}(\epsilon^{-4})$  first-order oracle complexity, matching SGD for single-level nonconvex minimization (Arjevani et al., 2023). Under the more general fully stochastic setting when noise appears both in  $\nabla f$  and  $\nabla g$ , we prove the method has an  $\tilde{\mathcal{O}}(\epsilon^{-6})$  complexity.

Compared with Xiao et al. (2023), our deterministic F<sup>2</sup>BA achieves the near-optimal rate on  $\epsilon$  without the assistance of Hessian-vector-product oracles. Our method also has a better dependency on  $\kappa$  since we do not use a squared trick on the Hessian of  $g$ . And we additionally study

stochastic problems that have not been studied by [Xiao et al. \(2023\)](#). Compared with [Kwon et al. \(2024b\)](#), we strictly improve the complexities for both deterministic and stochastic cases under the same assumptions. We compare our results with prior works in [Table 1](#), and leave a more detailed introduction of related works in [Appendix A](#).

**Notations.** Throughout this paper, we use  $\|\cdot\|$  to denote the  $\ell_2$ -norm of a vector or the operator norm of a matrix. We use  $\mathbb{B}_\delta(z) = \{z' : \|z' - z\| \leq \delta\}$  to denote the  $\ell_2$ -ball centered at  $z$  with radius  $\delta$ . We use notation  $\tilde{\mathcal{O}}(\cdot)$  to hide logarithmic factors in notation  $\mathcal{O}(\cdot)$ . For a matrix  $A$ , we use  $A^\dagger$  to denote the Moore–Penrose inverse. Notations  $\text{Ker}(A) = \{x : Ax = 0\}$  and  $\text{Range}(A) = \{Ax\}$  denote the kernel space and range space of  $A$ , respectively. For a vector  $v$ , we use the subscript  $v_{[j]}$  to denote its  $j$ -th coordinate.

## 2. Preliminaries

This section presents some basic definitions that are commonly used in optimization ([Nesterov, 2018](#)). To start with, the following definitions describe different orders of smoothness and levels of convexity for a function.

**Definition 2.1** We say an operator  $\mathcal{T}(x) : \mathbb{R}^{d_1} \rightarrow \mathbb{R}^{d_2 \times d_3}$  is  $C$ -Lipschitz for some  $C > 0$  if

$$\|\mathcal{T}(x) - \mathcal{T}(x')\| \leq C\|x - x'\|, \quad \forall x, x' \in \mathbb{R}^{d_1}.$$

For a function  $h(x) : \mathbb{R}^d \rightarrow \mathbb{R}$ , we say it has  $L$ -Lipschitz gradients (or it is  $L$ -smooth) if it is differentiable and  $\nabla h(x)$  is  $L$ -Lipschitz; we say it has  $\rho$ -Lipschitz Hessians if it is twice differentiable and  $\nabla^2 h(x)$  is  $\rho$ -Lipschitz.

**Definition 2.2** We say a function  $h(x) : \mathbb{R}^d \rightarrow \mathbb{R}$  is  $\mu$ -strongly convex for some  $\mu > 0$  if for any  $x, x' \in \mathbb{R}^d$  and  $t \in (0, 1)$ , we have that

$$h(tx + (1-t)x') \leq th(x) + (1-t)h(x') - \frac{1}{2}\mu t(1-t)\|x - x'\|^2.$$

We say  $h(x)$  is convex if  $\mu = 0$ .

As one relaxation of the above strong convexity condition, the Polyak–Łojasiewicz (PL) condition, independently introduced by [Polyak \(1967\)](#) and [Łojasiewicz \(1963\)](#), is formally defined as follows.

**Definition 2.3** We say a function  $h(x) : \mathbb{R}^d \rightarrow \mathbb{R}$  is  $\mu$ -PL for some  $\mu > 0$  if it has a non-empty solution set and for any  $x \in \mathbb{R}^d$  it holds that  $\|\nabla h(x)\|^2 \geq 2\mu(h(x) - \min_{x \in \mathbb{R}^d} h(x))$ .

Compared to the strong convexity, the PL condition allows nonconvexity and multiple minima. The PL condition has wide applications in establishing global convergence of many nonconvex learning problems, including neural network training ([Charles and Papailiopoulos, 2018](#); [Liu et al., 2022b](#); [Hardt and Ma, 2016](#); [Li et al., 2018](#)) and optimal control ([Fazel et al., 2018](#)).

Recall our goal is to minimize the hyper-objective  $\varphi(x)$ . Since  $\varphi(x)$  is typically nonconvex in bilevel optimization, the common goal is to find an  $\epsilon$ -stationary point, defined as follows.

**Definition 2.4** We say  $x$  is an  $\epsilon$ -stationary point of a differentiable function  $\varphi(x)$  if  $\|\nabla \varphi(x)\| \leq \epsilon$ .

### 3. Negative Results for General Convex Lower-Level Functions

This section formally describes the challenges for bilevel optimization without lower-level strong convexity assumption. In Section 3.1, we show that  $\varphi(x)$  may not have stationary points and analyze the underlying reasons behind it. In Section 3.2, we demonstrate that even if a stationary point of  $\varphi(x)$  exists, a zero-respecting algorithm may not be able to find it within a finite time.

#### 3.1. Stationary Points May not Exist

The following example shows that when the lower-level function only has convexity,  $\varphi(x)$  (Equation (2)) can be discontinuous and has no stationary points.

**Example 3.1 (Lucchetti et al. (1987))** Consider a bilevel problem as Problem (1) with  $d_x = 1$ ,  $d_y = 1$ . Let  $f(x, y) = x^2 + y^2$ ,  $g(x, y) = xy + I_C(y)$ , where  $I_C(\cdot)$  is the indicator function of the set  $\{y : 0 \leq y \leq 1\}$ . In this example  $g(x, y)$  is convex in  $y$ . But the hyper-objective  $\varphi(x)$  is discontinuous at  $x = 0$ , because  $\lim_{x \rightarrow 0^+} \varphi(x) = 0$ ,  $\lim_{x \rightarrow 0^-} \varphi(x) = 1$ .

**Remark 1** In the above example, the lower-level problem is a constrained optimization in  $y$ . We can also give a similar counter-example for unconstrained problems by replacing  $I_C(y)$  with a smoothed surrogate  $h(y) = (y - 1)\mathbb{I}[y \geq 1] - y\mathbb{I}[y \leq 0]$  and then letting  $g(x, y) = xy + h(y)$ .

In this example, the discontinuity of  $\varphi(x)$  comes from the discontinuity of  $Y^*(x)$ . Below, we prove that this statement and its reverse generally holds. As  $Y^*(x)$  is a set-valued mapping, we introduce the Hausdorff distance and use it to define different types of continuity.

**Definition 3.1** The Hausdorff distance between two sets  $S_1, S_2 \subseteq \mathbb{R}^d$  is defined as

$$\text{dist}(S_1, S_2) = \max \left\{ \sup_{x_1 \in S_1} \inf_{x_2 \in S_2} \|x_1 - x_2\|, \sup_{x_2 \in S_2} \inf_{x_1 \in S_1} \|x_1 - x_2\| \right\}.$$

We also denote  $\text{dist}(v, S) = \text{dist}(\{v\}, S)$  for  $v \in \mathbb{R}^d$ ,  $S \subseteq \mathbb{R}^d$ .

**Definition 3.2** We say a set-valued mapping  $S(x) : \mathbb{R}^n \rightrightarrows \mathbb{R}^m$  is (Pompeiu–Hausdorff) continuous if for any  $x \in \mathbb{R}^n$  and any  $\epsilon > 0$ , there exists  $\delta > 0$ , such that for any  $x' \in \mathbb{R}^n$  satisfying  $\|x' - x\| \leq \delta$ , we have  $\text{dist}(S(x), S(x')) \leq \epsilon$ .

**Definition 3.3** We say a set-valued mapping  $S(x) : \mathbb{R}^n \rightrightarrows \mathbb{R}^m$  is (Pompeiu–Hausdorff) locally Lipschitz if for any  $x \in \mathbb{R}^n$ , there exists  $\delta > 0$  and  $L > 0$  such that for any  $x' \in \mathbb{B}_\delta(x)$  we have  $\text{dist}(S(x), S(x')) \leq L\|x - x'\|$ . We say  $S(x)$  is (globally) Lipschitz if it holds for  $\delta \rightarrow \infty$ .

Based on the definitions, we show the following connections between  $\varphi(x)$  and  $Y^*(x)$ .

**Theorem 3.1** Suppose that for any given  $x \in \mathbb{R}^{d_x}$  the set  $Y^*(x)$  is non-empty and compact.

- (a) If  $f(x, y)$  and  $Y^*(x)$  are continuous, then  $\varphi(x)$  is continuous.
- (b) Conversely, if  $\varphi(x)$  is continuous for any continuous  $f(x, y)$ , then  $Y^*(x)$  is continuous.
- (c) If  $f(x, y)$  and  $Y^*(x)$  are locally Lipschitz, then  $\varphi(x)$  is locally Lipschitz.

---

**Algorithm 1** Zero-respecting algorithms for Problem (5)

---

- 1: **inputs:** initialization  $x_0, y_0$ , number of outer loops  $T$ , number of inner loops  $K$
  - 2: **for**  $t = 0, \dots, T - 1$
  - 3:   Generate  $y_t^0$  such that  $\text{supp}(y_t^0) \subseteq \bigcup_{0 \leq s < t, 0 \leq k \leq K} \text{supp}(y_s^k)$ .
  - 4:   **for**  $k = 0, \dots, K - 1$
  - 5:     Generate  $y_t^{k+1}$  such that  $\text{supp}(y_t^{k+1}) \subseteq \bigcup_{0 \leq i \leq k, h \in \{f, g\}} \text{supp}(\nabla_y h(x_t, y_t^i))$ .
  - 6:   **end for**
  - 7:   Generate  $x_{t+1}$  such that  $\text{supp}(x_{t+1}) \subseteq \bigcup_{0 \leq s \leq t} \text{supp}(\nabla_x f(x_s, y_s^K))$ .
  - 8: **end for**
- 

- (d) Conversely, if  $\varphi(x)$  is locally Lipschitz for any locally Lipschitz function  $f(x, y)$ , then  $Y^*(x)$  is locally Lipschitz.
- (e) If  $f(x, y)$  is  $C_f$ -Lipschitz and  $Y^*(x)$  is  $\kappa$ -Lipschitz, then we have that  $\varphi(x)$  is  $C_\varphi$ -Lipschitz with  $C_\varphi = (\kappa + 1)C_f$ .
- (f) Conversely, if  $\varphi(x)$  is  $C_\varphi$ -Lipschitz for any  $C_f$ -Lipschitz  $f(x, y)$ , then we know that  $Y^*(x)$  is  $\kappa$ -Lipschitz with  $\kappa = C_\varphi/C_f$ .

The theorem implies that continuity of the hyper-objective  $\varphi(x)$  requires a strong assumption on the set of minima  $Y^*(x)$  for  $g$  and suggests that one would need local strong convexity of  $g$  for the hyper-gradients to exist. However, we will see in the next subsection that even in such cases, finding a small hyper-gradient can be hard.

### 3.2. Stationary Points May be Intractable to Find

This subsection shows that even for nonconvex-strictly-convex bilevel problems where  $\nabla\varphi(x)$  is guaranteed to exist, finding a point with a small hyper-gradient can still be intractable. We prove the negative result on the following simplified case of Problem (1) when the lower-level function does not depend on  $x$ :

$$\min_{x \in \mathbb{R}^{d_x}, y \in Y^*} f(x, y), \quad \text{s.t.} \quad Y^* = \arg \min_{y \in \mathbb{R}^{d_y}} g(y). \quad (5)$$

This problem is much simpler because the implicit gradient in Equation (4) disappears as  $Y^*$  is a fixed set. However, we can show that this problem is hard enough for the following algorithm class.

**Zero-respecting algorithms for bilevel problems.** We define an algorithm class that covers a wide range of existing algorithms designed for bilevel optimization. This definition is inspired by the classical definitions in [Nesterov \(2018\)](#) but has an additional structure for bilevel problems (5). We first recall the definition of “zero-respecting” algorithms.

**Definition 3.4 (Definition 1 [Carmon et al. \(2021\)](#))** For a vector  $v \in \mathbb{R}^d$ , we use  $\text{supp}(v) = \{j \in [d] : v_{[j]} \neq 0\}$  to denote its support. We say an algorithm  $\mathcal{A}$  is zero-respecting to oracle  $\mathbb{O} : \mathbb{R}^d \rightarrow \mathbb{R}^d$  if the sequence  $\{z_t\}$  generated by algorithm  $\mathcal{A}$  only explore coordinates in the support of the previous oracles, i.e. the sequence satisfies  $\text{supp}(z_{t+1}) \subseteq \bigcup_{0 \leq s \leq t} \text{supp}(\mathbb{O}(z_s))$ .

We then define the zero-respecting algorithms for Problem (5).

**Definition 3.5** We say a (first-order) algorithm for Problem (5) is zero-respecting if it has the procedure as Algorithm 1. Such an algorithm consists of two loops: In the inner loop, it generates a sequence  $\{y_t^k\}_{k=1}^K$  that is zero-respecting to  $\nabla_y f(x_t, y)$  and  $\nabla g(y)$  for a fixed  $x_t$ ; In the outer loop, it generates a sequence  $\{x_t\}_{t=1}^T$  that is zero-respecting to  $\nabla_x f(x, y^K)$ .

The above zero-respecting algorithm class for bilevel problems subsumes many known algorithms when applied to Problem (5), including: AID (Ghadimi and Wang, 2018), ITD (Ji et al., 2021), GALET (Xiao et al., 2023), (Prox)-F<sup>2</sup>BA (Kwon et al., 2024b, 2023; Chen et al., 2023), Fde-HBO (Yang et al., 2024), BGS-Opt (Arbel and Mairal, 2022), BDA (Liu et al., 2020), BVFIM (Liu et al., 2021a), PDBO (Sow et al., 2022), SLM (Lu, 2023), LV-HBA (Yao et al., 2024).

Below, we give a hard instance such that all these algorithms cannot find small hypergradients as they get stuck at the initialization  $x_0$ .

**Theorem 3.2** Without loss of generality, suppose  $x_0 = y_0 = 0$  (otherwise we can translate the functions and the result still holds). Fix  $T$  and  $K$ . Let  $d_x = 1$ ,  $d_y = q = 2TK$ , and

$$f(x, y) = 2(x + 1)^2 \sum_{j=q/2}^q \psi(y_{[j]}), \quad g(y) = \frac{1}{8}(y_{[1]} - 1/\sqrt{q})^2 + \frac{1}{8} \sum_{j=1}^{q-1} (y_{[j+1]} - y_{[j]})^2,$$

where  $\psi(\cdot) : \mathbb{R} \rightarrow \mathbb{R}$  (defined in Equation (14)) is a function with  $\psi'(0) = \psi(0) = 0$ . Define the sublevel set  $\mathcal{X} := \{x : \varphi(x) \leq \varphi(0)\}$ . There exists numerical constants  $c_1, c_2 > 0$  such that

1.  $f(x, y)$  is  $c_1$ -Lipschitz in  $y$  on  $\mathcal{X} \times \mathbb{R}^{d_y}$ ;
2.  $f(x, y)$  has  $c_2$ -Lipschitz gradients on  $\mathcal{X} \times \mathbb{R}^{d_y}$ ;
3.  $g(y)$  is a strictly convex quadratic and has 1-Lipschitz gradients;
4. The resulting hyper-objective is  $\varphi(x) = (x + 1)^2/2$ .

For this problem, any algorithm with a procedure as Algorithm 1 stays at  $x_t = 0$  for any iteration number  $t \leq T$ .

In our construction, we let  $\psi(t)$  be a function which is  $\frac{1}{2}t^2$  near zero, while remaining bounded when  $|t|$  is large. We assign different dimensions in  $y$  to functions  $\psi$  and  $g$  separately. Note that  $g$  is designed such that any zero-respecting algorithm can only make progress at most one dimension per oracle calls. As the progress in  $g$  is slow, it will not affect the dimensions that  $\psi$  depends on.

Below we discuss the insights brought by the hard instances that we have constructed.

**Remark 3.1** Our results in these two subsections, from two complementary perspectives, motivate us to focus on more well-behaved lower-level functions that: (1) can confer the continuity of  $Y^*(x)$ , to avoid the case as Example 3.1. (2) an algorithm can converge rapidly to a neighborhood of  $Y^*(x)$ , to avoid the case as Theorem 3.2. In the next section, we will show that the PL condition simultaneously satisfies both of these requirements.

## 4. Positive Results for Lower-Level Functions Satisfying PL Conditions

As we have shown finding small hyper-gradients of nonconvex-convex bilevel problems is intractable for nonconvex-convex bilevel problems, we turn our attention to the tractable cases. Motivated by recent works (Shen and Chen, 2023; Kwon et al., 2024b; Arbel and Mairal, 2022), we study the case when the lower-level problem satisfies the PL condition.

### 4.1. The Assumptions for Nonconvex-PL Bilevel Problems

Without the typical lower-level strong convexity assumption, it is difficult to directly analyze the implicit gradient as Equation (4) since we can not directly use the implicit function theorem. Recently, Kwon et al. (2024b) proposed a novel way to study the differentiability of  $\varphi(x)$  for nonconvex-PL bilevel problems. Instead of directly studying the original hyper-objective  $\varphi(x)$ , they studied the following value-function penalized hyper-objective as a bridge:

$$\varphi_\sigma(x) := \min_{y \in \mathcal{Y}} \left\{ f(x, y) + \frac{g(x, y) - g^*(x)}{\sigma} \right\}, \quad (6)$$

where  $g^*(x) = \min_{y \in \mathcal{Y}} g(x, y)$  is the lower-level value-function and  $\mathcal{Y} \subseteq \mathbb{R}^{d_y}$ . They first studied the differentiability of  $\varphi_\sigma(x)$  and then showed the limit of  $\sigma \rightarrow 0^+$  exists.

However, the PL condition on  $g$  is not sufficient to guarantee the differentiability of  $\varphi_\sigma(x)$ . Below, we give a concrete example to illustrate this.

**Example 4.1** Consider Problem (1) with  $d_x = 1$ ,  $d_y = 2$ . Let the upper-level function  $f(x, y) = xy_{[1]}$ , and the lower-level function  $g(x, y) = \frac{1}{2}y_{[2]}^2$ . The lower-level function is 1-PL in  $y$ . Consider the penalized hyper-objective  $\varphi_\sigma(x)$  in Equation (6). For any  $\sigma \geq 0$ , if taking minimum in  $y$  over all the domain  $y \in \mathbb{R}^2$ ,  $\varphi_\sigma(x)$  is not well-defined since  $\varphi_\sigma(x) = -\infty$  for any  $x \neq 0$ ; if taking minimum in  $y$  over a compact domain such as  $\mathcal{Y} = [0, 1] \times [0, 1]$ ,  $\varphi_\sigma(x) = \min\{x, 0\}$  is not differentiable at the point  $x = 0$ .

To guarantee the differentiability of  $\varphi_\sigma(x)$ , Kwon et al. (2024b) assumed the PL condition not only holds for  $g$ , but also holds for the penalty function  $h_\sigma = \sigma f + g$  uniformly for all  $\sigma$  near zero. In addition to this assumption, the authors also imposed other standard smoothness assumptions which are typically required in previous works. These assumptions, as stated in Kwon et al. (2024b), are formally presented below.

**Assumption 4.1** Recall the bilevel problem defined in Equation (1), where  $f$  is the upper-level problem,  $g$  is the lower-level problem. Let  $h_\sigma = \sigma f + g$  be the penalty function. Suppose that

- (a) The penalty function  $h_\sigma(x, y)$  is  $\mu$ -PL in  $y$  for any  $0 \leq \sigma \leq \bar{\sigma}$ ;
- (b) The upper-level function  $f(x, y)$  is  $C_f$ -Lipschitz in  $y$  and has  $L_f$ -Lipschitz gradients;
- (c) The lower-level function  $g(x, y)$  has  $L_g$ -Lipschitz gradients;
- (d) The upper-level function  $f(x, y)$  has  $\rho_f$ -Lipschitz Hessians;
- (e) The lower-level function  $g(x, y)$  has  $\rho_g$ -Lipschitz Hessians.



Under this assumption, we define the largest smoothness constant  $\ell = \max\{C_f, L_f, L_g, \rho_g\}$  and the condition number  $\kappa := \ell/\mu$ .

**Remark 4.1** Note that all the assumptions are the same as [Kwon et al. \(2024b\)](#), except (a) may seem different from the Prox-EB assumption (Assumption 1 in [Kwon et al. \(2024b\)](#)). However, in [Appendix D](#) we show these two assumptions can imply each other. Our narrative uses the PL condition and is therefore more convenient. As a by-product, we show the proximal operator in [Kwon et al. \(2024b\)](#) is unnecessary, and the original F<sup>2</sup>BA ([Chen et al., 2023](#)) can also converge under PL conditions.

Because we focus on the behavior when  $\sigma$  is close to zero, whenever we mention  $\sigma$  in the context, we always assume that  $\sigma \in [0, \bar{\sigma}]$  even when this condition is not explicitly stated.

Assumption 4.1 ensures that the solution set  $Y^*$  is stable under perturbations of  $\sigma$  and  $x$ .

**Lemma 4.1** Let  $Y_\sigma^* := \arg \min_{y \in \mathbb{R}^{d_y}} h_\sigma(x, y)$  denote the set of minima for the penalty function  $h_\sigma(x, y) = \sigma f(x, y) + g(x, y)$ . Under Assumption 4.1, we have that

$$\text{dist}(Y_{\sigma_1}^*(x_1), Y_{\sigma_2}^*(x_2)) \leq \frac{C_f}{\mu} |\sigma_1 - \sigma_2| + \frac{\sigma L_f + L_g}{\mu} \|x_1 - x_2\|.$$

Then we can use [Theorem 3.1](#) to get the continuity of hyper-objective  $\varphi(x)$ .

## 4.2. F<sup>2</sup>BA Can also be Applied to Nonconvex-PL Bilevel Problems

Although F<sup>2</sup>BA is originally proposed for nonconvex-strongly-convex bilevel problems, we show that F<sup>2</sup>BA can also be applied to nonconvex-PL bilevel problems in this section.

Our starting points are the following lemmas that hold once we have Lipschitz continuity of solution set from [Lemma 4.1](#), which unnecessarily requires strong convexity. Firstly, we can obtain the following result by using the generalized Danskin's theorem ([Shen and Chen, 2023](#)) twice, specifically, in both  $Y^*(x)$  and  $Y_\sigma^*(x)$ .

**Lemma 4.2** ([Shen and Chen \(2023\)](#)) Recall that  $\varphi_\sigma(x)$  is the penalized hyper-objective defined in [Equation \(6\)](#). Under Assumption 4.1,  $\nabla \varphi_\sigma(x)$  exists and takes the form of

$$\nabla \varphi_\sigma(x) = \nabla_x f(x, y) + \frac{\nabla_x g(x, y_\sigma^*(x)) - \nabla_x g(x, y^*(x))}{\sigma}, \quad (7)$$

where  $y^*(x)$ ,  $y_\sigma^*(x)$  can be arbitrary elements in  $Y^*(x)$  and  $Y_\sigma^*(x)$ , respectively.

Secondly, the stability of  $Y_\sigma^*$  under perturbations of  $\sigma$  by [Lemma 4.1](#) implies the stability of  $\varphi_\sigma$  and  $\nabla \varphi_\sigma$  by invoking the result by [Kwon et al. \(2024b\)](#).

**Lemma 4.3** ([Kwon et al. \(2024b\)](#)) Recall that  $\varphi(x)$  is the original hyper-objective in [Equation \(2\)](#), while  $\varphi_\sigma(x)$  is the penalized hyper-objective in [Equation \(6\)](#). Under Assumption 4.1,  $\nabla \varphi(x)$  exists and can be defined as the limit  $\lim_{\sigma \rightarrow 0^+} \nabla \varphi_\sigma(x)$ . Moreover,  $\varphi_\sigma(x)$  is close to  $\varphi(x)$ . Formally, for any  $0 \leq \sigma \leq \min\{\rho_g/\rho_f, \bar{\sigma}\}$ , we have that

$$|\varphi_\sigma(x) - \varphi(x)| = \mathcal{O}(\sigma \ell \kappa), \quad \text{and} \quad \|\nabla \varphi_\sigma(x) - \nabla \varphi(x)\| = \mathcal{O}(\sigma \ell \kappa^3).$$

**Algorithm 2** F<sup>2</sup>BA ( $x_0, y_0, \eta, \tau, \sigma, T, K$ )

---

```

1:  $z_0 = y_0$ 
2: for  $t = 0, 1, \dots, T - 1$ 
3:    $y_t^0 = y_t, z_t^0 = z_t$ 
4:   for  $k = 0, 1, \dots, K - 1$ 
5:      $z_t^{k+1} = z_t^k - \tau \nabla_y g(x_t, z_t^k)$ 
6:      $y_t^{k+1} = y_t^k - \tau (\sigma \nabla_y f(x_t, y_t^k) + \nabla_y g(x_t, y_t^k))$ 
7:   end for
8:    $\hat{\nabla} \varphi(x_t) = \nabla_x f(x_t, y_t^K) + (\nabla_x g(x_t, y_t^K) - \nabla_x g(x_t, z_t^K)) / \sigma$ 
9:    $x_{t+1} = x_t - \eta \hat{\nabla} \varphi(x_t)$ 
10: end for

```

---

These two lemmas make it reasonable to apply the F<sup>2</sup>BA (Algorithm 2), which applies gradient descent on  $\varphi_\sigma(x)$  according to Equation (7). Lemma 4.3 shows that to find an  $\mathcal{O}(\epsilon)$ -stationary point of  $\varphi(x)$ , it suffices to find an  $\mathcal{O}(\epsilon)$ -stationary point of  $\varphi_\sigma(x)$  for  $\sigma = \mathcal{O}(\epsilon)$ . Note that gradient descent can find an  $\epsilon$ -stationary point of a nonconvex  $L$ -smooth function within  $\mathcal{O}(L\epsilon^{-2})$  complexity and that  $\varphi_\sigma(x)$  is  $\mathcal{O}(\sigma^{-1})$ -smooth in the worst-case. Such an analysis illustrates an  $\tilde{\mathcal{O}}(\sigma^{-1}\epsilon^{-2}) = \tilde{\mathcal{O}}(\epsilon^{-3})$  complexity as Kwon et al. (2024b).

### 4.3. Near-Optimal Rates for the Deterministic Case

Interestingly, we can show a near-optimal rate of  $\tilde{\mathcal{O}}(\epsilon^{-2})$  of Algorithm 2. Our improvement comes from a similar technique by Chen et al. (2023), which shows that  $\varphi(x)$  is  $\mathcal{O}(1)$ -smooth for nonconvex-strongly-convex problems. The intuition is that we can restrict our analysis to the strongly convex subspace induced by the PL condition, and apply the result of Chen et al. (2023).

**Lemma 4.4** *Under Assumption 4.1,  $\varphi(x)$  has  $\mathcal{O}(\ell\kappa^3)$ -Lipschitz gradients.*

We sketch the proof as follows. First, we give the explicit form of  $\nabla \varphi(x)$  in Lemma G.7. Next, we show that the PL condition ensures a strongly convex subspace near any minimum in Lemma G.6. Finally, We apply Lemma G.3 and G.5 to project the functions onto the strongly convex subspace, where  $\varphi(x)$  has  $\mathcal{O}(1)$ -Lipschitz gradients to complete the proof.

Based on this lemma, we can readily prove our main theorem in this subsection, stated below.

**Theorem 4.1** *Suppose Assumption 4.1 holds. Define  $\Delta := \varphi(x_0) - \inf_{x \in \mathbb{R}^{d_x}} \varphi(x)$ ,  $R := \text{dist}^2(y_0, Y^*(x))$  and supposed both  $\Delta, R$  are bounded. Set the parameters in Algorithm 2 as*

$$\eta \asymp \ell^{-1} \kappa^{-3}, \quad \sigma \asymp \min \left\{ \frac{R}{\kappa}, \frac{\epsilon}{\ell \kappa^3}, \frac{L_g}{L_f}, \frac{\rho_g}{\rho_f}, \bar{\sigma} \right\},$$

$$\tau = \frac{1}{\sigma L_f + L_g}, \quad K \asymp \frac{L_g}{\mu} \log \left( \frac{L_g}{\mu \sigma} \right),$$

*then it can find an  $\epsilon$ -first-order stationary point of  $\varphi(x)$  within  $T = \mathcal{O}(\ell\kappa^3\epsilon^{-2})$  iterations, where  $\ell, \kappa$  are defined in Assumption 4.1. The total number of first-order oracle calls is bounded by  $\mathcal{O}(\ell\kappa^4\epsilon^{-2} \log(\rho_f \ell \kappa / \epsilon))$ .*

---

**Algorithm 3** F<sup>2</sup>BSA  $(x_0, y_0, \delta_0, \eta, \tau, \sigma, T, B)$ 


---

```

1:  $z_0 = y_0$ 
2: for  $t = 0, 1, \dots, T - 1$ 
3:    $y_t^0 = y_t, z_t^0 = z_t$ 
4:   Set  $K_t$  as Equation (8) based on the value of  $\delta_t$ 
5:   for  $k = 0, 1, \dots, K_t - 1$ 
6:      $z_t^{k+1} = z_t^k - \tau \nabla_y g(x_t, z_t^k; B)$ 
7:      $y_t^{k+1} = y_t^k - \tau (\sigma \nabla_y f(x_t, y_t^k; B) + \nabla_y g(x_t, y_t^k; B))$ 
8:   end for
9:    $\hat{\nabla} \varphi(x_t) = \nabla_x f(x_t, y_t^K; B) + (\nabla_x g(x_t, y_t^K; B) - \nabla_x g(x_t, z_t^K; B)) / \sigma$ 
10:   $x_{t+1} = x_t - \eta \hat{\nabla} \varphi(x_t)$ 
11:  Calculate  $\delta_{t+1}$  as Equation (9) based on  $\delta_t, x_{t+1}$  and  $x_t$ 
12: end for
    
```

---

The above complexity matches gradient descent for single-level nonconvex minimization by Carmon et al. (2021), except we require  $f$  to be  $\rho_f$ -Hessian Lipschitz in Assumption 4.1d. However, our dependency on  $\rho_f$  is only logarithmic, rather than polynomial.

**Remark 4.2** Here we note some small differences between the assumptions in Theorem 4.1 and those of the constructed hard instance in Theorem 3.2.

One difference is that the hard instance only satisfies the Lipschitz-type conditions in  $\mathcal{X} \times \mathbb{R}^{d_y}$ , where  $\mathcal{X}$  is the sublevel set of  $x_0$ . However, indeed, the upper bound also only requires the Lipschitz-type conditions in the sublevel set by the descent lemma  $\varphi(x_{t+1}) \leq \varphi(x_t)$ .

Another difference is that the upper bound requires  $f$  to have  $\rho_f$ -Lipschitz Hessians while the Hessian Lipschitz constant in the hard instance may depend on  $d_y$ . However, the upper bound only has a logarithmic dependency on  $\rho_f$ . Such a difference is also mild.

Besides the above two differences, the only remaining difference is that the lower-level function in the hard instance is only convex, while Assumption 4.1a required by the upper bound subsumes local strong convexity. Such a difference leads to a separation of intractability and tractability.

#### 4.4. The Extensions of F<sup>2</sup>BA to the Stochastic Case

Besides the near-optimal rate in the deterministic case, our novel viewpoint can also lead to better bounds on the stochastic case, where  $f(x, y), g(x, y)$  are both the expectation of some stochastic components  $F(x, y; \xi)$  and  $G(x, y; \zeta)$ , indexed by random vectors  $\xi$  and  $\zeta$ :

$$f(x, y) := \mathbb{E}_\xi [F(x, y; \xi)], \quad g(x, y) := \mathbb{E}_\zeta [G(x, y; \zeta)].$$

An algorithm has access to the stochastic gradients, with the following assumptions.

**Assumption 4.2** We study the stochastic problem under the two assumptions below.

(a) Suppose the stochastic gradients are unbiased:

$$\mathbb{E}_\xi [\nabla F(x, y; \xi)] = \nabla f(x, y), \quad \mathbb{E}_\zeta [\nabla G(x, y; \zeta)] = \nabla g(x, y);$$

(b) Suppose the stochastic gradients have bounded variance. In other words, there exist some constants  $M_f, M_g > 0$  such that

$$\mathbb{E}_\xi [\|\nabla F(x, y; \xi) - \nabla f(x, y)\|^2] \leq M_f^2, \quad \mathbb{E}_\zeta [\|\nabla G(x, y; \zeta) - \nabla g(x, y)\|^2] \leq M_g^2.$$

Under these assumptions, the natural extension of Algorithm 2 to the stochastic setting is to replace the full-batch gradient in Algorithm 2 to the mini-batch gradient, defined as follows.

**Definition 4.1** Given mini-batch size  $B$ . We define the gradient estimators for the upper-level and lower-level functions using the notations below,

$$\nabla f(x, y; B) = \frac{1}{B} \sum_{i=1}^B \nabla f(x, y; \xi_i), \quad \nabla g(x, y; B) = \frac{1}{B} \sum_{i=1}^B \nabla g(x, y; \zeta_i),$$

where both  $\xi_i$  and  $\zeta_i$  are sampled i.i.d.

By replacing all the full-batch gradients in deterministic Algorithm 2, we get the stochastic counterpart Algorithm 3, namely F<sup>2</sup>BSA (Fully First-order Bilevel Stochastic Approximation). By additionally taking into account the error from stochastic gradients, we can extend Theorem 4.1 to also tackle the stochastic case, yielding the following theorem.

**Theorem 4.2** Suppose Assumption 4.1 and 4.2 hold. Define  $\Delta := \varphi(x_0) - \inf_{x \in \mathbb{R}^{d_x}} \varphi(x)$ ,  $R := \text{dist}^2(y_0, Y^*(x))$  and supposed both  $\Delta, R$  are bounded. Set the parameters in Algorithm 3 as

$$\begin{aligned} \eta &\asymp \ell^{-1} \kappa^{-3}, \quad \sigma \asymp \min \left\{ \frac{R}{\kappa}, \frac{\epsilon}{\ell \kappa^3}, \frac{L_g}{L_f}, \frac{\rho_g}{\rho_f}, \bar{\sigma} \right\}, \\ \tau &= \frac{1}{\sigma L_f + L_g}, \quad B \asymp \frac{L_g (\sigma^2 M_f^2 + M_g^2)}{\mu \sigma^2 \epsilon^2}, \quad K_t \asymp \frac{L_g}{\mu} \log \left( \frac{L_g^3 \delta_t}{\mu \sigma^2 \epsilon^2} \right), \end{aligned} \quad (8)$$

where  $\delta_t$  is defined via the recursion

$$\delta_{t+1} \leq \frac{1}{2} \delta_t + \frac{8L_g^2}{\mu^2} \|x_{t+1} - x_t\|^2 + \mathcal{O} \left( \frac{\sigma^2 \epsilon^2}{L_g^2} \right), \quad \delta_0 \asymp R. \quad (9)$$

Then Algorithm 3 can find an  $\epsilon$ -first-order stationary point of  $\varphi(x)$  in expectation within  $T = \mathcal{O}(\ell \kappa^3 \epsilon^{-2})$  iterations, where  $\ell, \kappa$  are defined in Definition 4.1. The total number of stochastic first-order oracle calls is bounded by

$$\begin{cases} \mathcal{O}(\ell \kappa^4 \epsilon^{-2} \log(\rho_f \ell \kappa / \epsilon)), & M_f = 0, M_g = 0; \\ \mathcal{O}(\ell \kappa^5 \epsilon^{-4} \log(\rho_f \ell \kappa / \epsilon)), & M_f > 0, M_g = 0; \\ \mathcal{O}(\ell^3 \kappa^{11} \epsilon^{-6} \log(\rho_f \ell \kappa / \epsilon)), & M_f > 0, M_g > 0. \end{cases}$$

In the deterministic case ( $M_f = 0, M_g = 0$ ), this result recovers the  $\tilde{\mathcal{O}}(\epsilon^{-2})$  rate by Algorithm 2. In the partially stochastic case ( $M_f > 0, M_g = 0$ ), the  $\tilde{\mathcal{O}}(\epsilon^{-4})$  complexity also matches SGD (Arjevani et al., 2023). In the fully stochastic case ( $M_f > 0, M_g > 0$ ), our  $\tilde{\mathcal{O}}(\epsilon^{-6})$  upper bound is also better than the  $\tilde{\mathcal{O}}(\epsilon^{-7})$  upper bound by Kwon et al. (2024b).

## 5. Conclusions and Future Works

This paper investigates bilevel optimization without the typical lower-level strong convexity assumption. We have shown that finding a point with small a hyper-gradient is computationally hard for nonconvex-convex bilevel problems, but easy for nonconvex-PL bilevel problems where simple first-order algorithms can achieve fast convergence rates. In Appendix B, we discuss some future directions that have not been addressed by this work.

## References

- Michael Arbel and Julien Mairal. Amortized implicit differentiation for stochastic bilevel optimization. In *ICLR*, 2021.
- Michael Arbel and Julien Mairal. Non-convex bilevel games with critical point selection maps. In *NeurIPS*, 2022.
- Yossi Arjevani, Yair Carmon, John C. Duchi, Dylan J. Foster, Nathan Srebro, and Blake Woodworth. Lower bounds for non-convex stochastic optimization. *Mathematical Programming*, 199(1-2): 165–214, 2023.
- Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- Jérôme Bolte, Tam Le, Edouard Pauwels, and Tony Silveti-Falls. Nonsmooth implicit differentiation for machine-learning and optimization. In *NeurIPS*, 2021.
- 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.
- Zachary Charles and Dimitris Papailiopoulos. Stability and generalization of learning algorithms that converge to global optima. In *ICML*, 2018.
- Lesi Chen, Yaohua Ma, and Jingzhao Zhang. Near-optimal fully first-order algorithms for finding stationary points in bilevel optimization. *arXiv preprint arXiv:2306.14853*, 2023.
- Mathieu Dagréou, Pierre Ablin, Samuel Vaiter, and Thomas Moreau. A framework for bilevel optimization that enables stochastic and global variance reduction algorithms. In *NeurIPS*, 2022.
- Stephan Dempe. *Foundations of bilevel programming*. Springer Science & Business Media, 2002.
- Justin Domke. Generic methods for optimization-based modeling. In *AISTATS*, 2012.
- Asen L Dontchev, R Tyrrell Rockafellar, and R Tyrrell Rockafellar. *Implicit functions and solution mappings: A view from variational analysis*, volume 616. Springer, 2009.
- Ilyas Fatkhullin, Jalal Etesami, Niao He, and Negar Kiyavash. Sharp analysis of stochastic optimization under global kurdyka-lojasiewicz inequality. In *NeurIPS*, 2022.
- Maryam Fazel, Rong Ge, Sham M. Kakade, and Mehran Mesbahi. Global convergence of policy gradient methods for the linear quadratic regulator. In *ICML*, 2018.

- Luca Franceschi, Michele Donini, Paolo Frasconi, and Massimiliano Pontil. Forward and reverse gradient-based hyperparameter optimization. In *ICML*, 2017.
- Luca Franceschi, Paolo Frasconi, Saverio Salzo, Riccardo Grazi, and Massimiliano Pontil. Bilevel programming for hyperparameter optimization and meta-learning. In *ICML*, 2018.
- Saeed Ghadimi and Mengdi Wang. Approximation methods for bilevel programming. *arXiv preprint arXiv:1802.02246*, 2018.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.
- Stephen Gould, Basura Fernando, Anoop Cherian, Peter Anderson, Rodrigo Santa Cruz, and Edison Guo. On differentiating parameterized argmin and argmax problems with application to bi-level optimization. *arXiv preprint arXiv:1607.05447*, 2016.
- Riccardo Grazi, Luca Franceschi, Massimiliano Pontil, and Saverio Salzo. On the iteration complexity of hypergradient computation. In *ICML*, 2020.
- Moritz Hardt and Tengyu Ma. Identity matters in deep learning. *arXiv preprint arXiv:1611.04231*, 2016.
- Mingyi Hong, Hoi-To Wai, Zhaoran Wang, and Zhuoran Yang. A two-timescale stochastic algorithm framework for bilevel optimization: Complexity analysis and application to actor-critic. *SIAM Journal on Optimization*, 33(1):147–180, 2023.
- Timothy Hospedales, Antreas Antoniou, Paul Micaelli, and Amos Storkey. Meta-learning in neural networks: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(9): 5149–5169, 2021.
- Feihu Huang. On momentum-based gradient methods for bilevel optimization with nonconvex lower-level. *arXiv preprint arXiv:2303.03944*, 2023.
- Kaiyi Ji, Junjie Yang, and Yingbin Liang. Bilevel optimization: Convergence analysis and enhanced design. In *ICML*, 2021.
- Yibo Jiang and Victor Veitch. Invariant and transportable representations for anti-causal domain shifts. In *NeurIPS*, 2022.
- 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*, 2016.
- Vijay Konda and John Tsitsiklis. Actor-critic algorithms. In *NeurIPS*, 1999.
- Jeongyeol Kwon, Dohyun Kwon, Stephen Wright, and Robert Nowak. A fully first-order method for stochastic bilevel optimization. In *ICML*, 2023.
- Jeongyeol Kwon, Dohyun Kwon, and Hanbaek Lyu. On the complexity of first-order methods in stochastic bilevel optimization. In *ICML*, 2024a.

- Jeongyeol Kwon, Dohyun Kwon, Steve Wright, and Robert Nowak. On penalty methods for non-convex bilevel optimization and first-order stochastic approximation. In *ICLR*, 2024b.
- Jiajin Li, Linglingzhi Zhu, and Anthony Man-Cho So. Nonsmooth composite nonconvex-concave minimax optimization. *arXiv preprint arXiv:2209.10825*, 2022.
- Yuanzhi Li, Tengyu Ma, and Hongyang Zhang. Algorithmic regularization in over-parameterized matrix sensing and neural networks with quadratic activations. In *COLT*, 2018.
- Tianyi Lin, Chi Jin, and Michael I. Jordan. On gradient descent ascent for nonconvex-concave minimax problems. In *ICML*, 2020a.
- Tianyi Lin, Chi Jin, and Michael I. Jordan. Near-optimal algorithms for minimax optimization. In *COLT*, 2020b.
- Bo Liu, Mao Ye, Stephen Wright, Peter Stone, and Qiang Liu. BOME! bilevel optimization made easy: A simple first-order approach. In *NeurIPS*, 2022a.
- 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*, 59:85–116, 2022b.
- Chengchang Liu and Luo Luo. Quasi-newton methods for saddle point problems. In *NeurIPS*, 2022.
- Hanxiao Liu, Karen Simonyan, and Yiming Yang. DARTS: Differentiable architecture search. In *ICLR*, 2019.
- Risheng Liu, Pan Mu, Xiaoming Yuan, Shangzhi Zeng, and Jin Zhang. A generic first-order algorithmic framework for bi-level programming beyond lower-level singleton. In *ICML*, 2020.
- Risheng Liu, Xuan Liu, Xiaoming Yuan, Shangzhi Zeng, and Jin Zhang. A value-function-based interior-point method for non-convex bi-level optimization. In *ICML*, 2021a.
- Risheng Liu, Yaohua Liu, Shangzhi Zeng, and Jin Zhang. Towards gradient-based bilevel optimization with non-convex followers and beyond. *NeurIPS*, 2021b.
- Stanislaw Lojasiewicz. A topological property of real analytic subsets. *Coll. du CNRS, Les équations aux dérivées partielles*, 117(87-89):2, 1963.
- Songtao Lu. SLM: A smoothed first-order lagrangian method for structured constrained nonconvex optimization. In *NeurIPS*, 2023.
- Roberto Lucchetti, F Mignanego, and G Pieri. Existence theorems of equilibrium points in stackelberg. *Optimization*, 18(6):857–866, 1987.
- Yurii Nesterov. *Lectures on convex optimization*, volume 137. Springer, 2018.
- Fabian Pedregosa. Hyperparameter optimization with approximate gradient. In *ICML*, 2016.
- Quang Pham, Chenghao Liu, Doyen Sahoo, and HOI Steven. Contextual transformation networks for online continual learning. In *ICLR*, 2021.

- Boris Teodorovich Polyak. A general method for solving extremal problems. In *Doklady Akademii Nauk*, volume 174, pages 33–36. Russian Academy of Sciences, 1967.
- Aravind Rajeswaran, Chelsea Finn, Sham M. Kakade, and Sergey Levine. Meta-learning with implicit gradients. In *NeurIPS*, 2019.
- Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *ICLR*, 2017.
- Amirreza Shaban, Ching-An Cheng, Nathan Hatch, and Byron Boots. Truncated back-propagation for bilevel optimization. In *AISTATS*, 2019.
- Han Shen and Tianyi Chen. On penalty-based bilevel gradient descent method. In *ICML*, 2023.
- Aman Sinha, Hongseok Namkoong, and John C. Duchi. Certifying some distributional robustness with principled adversarial training. In *ICLR*, 2018.
- Daouda Sow, Kaiyi Ji, Ziwei Guan, and Yingbin Liang. A primal-dual approach to bilevel optimization with multiple inner minima. *arXiv preprint arXiv:2203.01123*, 2022.
- Xiaoxing Wang, Wenxuan Guo, Jianlin Su, Xiaokang Yang, and Junchi Yan. ZARTS: On zero-order optimization for neural architecture search. In *NeurIPS*, 2022.
- Quan Xiao, Songtao Lu, and Tianyi Chen. An alternating method for bilevel optimization under the polyak-łojasiewicz condition. In *NeurIPS*, 2023.
- Yifan Yang, Peiyao Xiao, and Kaiyi Ji. Achieving  $\mathcal{O}(\epsilon^{-1.5})$  complexity in hessian/jacobian-free stochastic bilevel optimization. In *NeurIPS*, 2024.
- Wei Yao, Chengming Yu, Shangzhi Zeng, and Jin Zhang. Constrained bi-level optimization: Proximal lagrangian value function approach and hessian-free algorithm. In *ICLR*, 2024.
- Jane J. Ye, DL Zhu, and Qiji Jim Zhu. Exact penalization and necessary optimality conditions for generalized bilevel programming problems. *SIAM Journal on optimization*, 7(2):481–507, 1997.
- Miao Zhang, Steven W. Su, Shirui Pan, Xiaojun Chang, Ehsan M Abbasnejad, and Reza Haffari. iDARTS: Differentiable architecture search with stochastic implicit gradients. In *ICML*, 2021.
- Taoli Zheng, Linglingzhi Zhu, Anthony Man-Cho So, Jose Blanchet, and Jiajin Li. Doubly smoothed gda: Global convergent algorithm for constrained nonconvex-nonconcave minimax optimization. In *NeurIPS*, 2023.
- Xiao Zhou, Yong Lin, Renjie Pi, Weizhong Zhang, Renzhe Xu, Peng Cui, and Tong Zhang. Model agnostic sample reweighting for out-of-distribution learning. In *ICML*, 2022.
- Barret Zoph and Quoc Le. Neural architecture search with reinforcement learning. In *ICLR*, 2016.



## Appendix A. Related Works

There have been many works studying bilevel optimization. Here we only review related works focusing on non-asymptotic analysis.

Many existing works for non-asymptotic analysis for bilevel optimization assume the lower-level function is strongly convex. A natural method is to estimate  $\nabla y^*(x)$ , and plug into Equation (4) to get an estimator of  $\nabla \varphi(x)$  to apply gradient descent. ITD (ITerative Differentiation) (Gould et al., 2016; Franceschi et al., 2017; Shaban et al., 2019; Bolte et al., 2021) approximates  $\nabla y^*(x)$  by  $\partial y^K(x)/\partial x$ , where  $y^K(x)$  is  $K$ -steps of gradient descent. AID (Approximate Implicit Differentiation) (Domke, 2012; Ghadimi and Wang, 2018; Pedregosa, 2016; Franceschi et al., 2018; Grazi et al., 2020; Ji et al., 2021; Dagr eou et al., 2022; Arbel and Mairal, 2021) explicitly solve  $\nabla y^*(x) = [(\nabla_{yy}^2 g)^{-1} \nabla_{yx}^2 g](x, y^*(x))$  as a linear equation. However, both AID and ITD require Hessian-vector product oracles, and their convergence analysis is typically restricted to the case when the lower-level function is strongly convex. Arbel and Mairal (2022) extended implicit differentiation to the parametric Morse-Bott function, a class of nonconvex functions with local PL properties. They also proposed an algorithm that combines AID and ITD and showed its limit points must be an equilibrium of the concept of BGS (Bilevel Game with Selection) that they introduced. Yang et al. (2024) proposed a first-order algorithm FdeHBO by estimating Jacobian/Hessian vector-product in AID with gradient differences.

Recently, F<sup>2</sup>BA (Fully First-order Bilevel Approximation) has gained widespread attention. This method uses the value-function-based reformulation (Ye et al., 1997) of Problem (1):

$$\min_{x \in \mathbb{R}^{d_x}, y \in \mathbb{R}^{d_y}} f(x, y), \quad \text{s.t.} \quad g(x, y) = g^*(x) \quad (10)$$

and applies gradient descent on the penalty function. Several studies have demonstrated two notable advantages of this method over AID and ITD: Firstly, it only requires gradient oracles which are much cheaper than Hessian-vector product oracles (Liu et al., 2022a). Secondly, it performs well empirically even in cases where the lower-level function is nonconvex (Liu et al., 2021a; Shen and Chen, 2023). By establishing the equivalence between the KKT (Karush–Kuhn–Tucker) point of Problem (10) and the stationary point of Problem (2), Kwon et al. (2023) showed that F<sup>2</sup>BA can find an  $\epsilon$ -stationary point of  $\varphi(x)$  with  $\tilde{O}(\epsilon^{-3})$  first-order oracle calls when the lower-level function is strongly convex. Chen et al. (2023) further improved the rate to  $\tilde{O}(\epsilon^{-2})$  by a more careful landscape analysis of the penalty function. Kwon et al. (2024b); Yao et al. (2024) proposed proximal variants of F<sup>2</sup>BA by replacing the lower-level value function with its Moreau envelope to tackle lower-level constraints. Kwon et al. (2024b) showed a  $\tilde{O}(\epsilon^{-3})$  complexity when the penalty function satisfies the Prox-EB condition (which is equivalent to the PL condition in the unconstrained case as Proposition D.1).

GALET (Generalized ALternating mEhod for bilevel opTimization) by Xiao et al. (2023) can also tackle nonconvex-PL bilevel optimization problems. It uses Hessian-vector-product oracles to solve the following gradient-based reformulation of Problem (1):

$$\min_{x \in \mathbb{R}^{d_x}, y \in \mathbb{R}^{d_y}} f(x, y), \quad \text{s.t.} \quad \nabla_y g(x, y) = 0. \quad (11)$$

Xiao et al. (2023) showed GALET can converge to a KKT point of Problem (11) with a rate of  $\tilde{O}(\epsilon^{-2})$ . In Appendix C, we show that this also implies the same convergence rate to the stationary point of  $\varphi(x)$  under the assumptions of Kwon et al. (2024b).

Huang (2023) also considers nonconvex-PL bilevel problems. Our work differs from (Huang, 2023) in the following aspects. First, Assumption 2 in (Huang, 2023) supposes that  $\nabla^2 g(x, y)$  is non-singular at the minimizer of  $g$ , while we do not require such an assumption. Secondly, our algorithm only requires first-order oracles, while (Huang, 2023) requires second-order oracles. Thirdly, our algorithm has only  $\mathcal{O}(d)$  complexity at each step, where  $d = \max\{d_x, d_y\}$  is the dimension of the problem. In contrast, each iteration of the algorithm in (Huang, 2023) requires a  $\mathcal{O}(d^3)$  complexity for computing the SVD decomposition of the Hessian matrix.

Some works use a sequential approximation minimization strategy (Liu et al., 2021b,a, 2020) to tackle the discontinuous challenge for bilevel problems without lower-level strong convexity as illustrated by Example 3.1. These works generate a series of continuous functions  $\{\varphi_K\}_{K=0}^\infty$  which converge to  $\varphi$  when  $K \rightarrow \infty$ . Liu et al. (2020) proposed BDA by defining  $\varphi_K(x) = f(x, \text{AggrGD}_{f,g}^K(x))$  for nonconvex-convex problems, where  $\text{AggrGD}_{f,g}^K(x)$  denotes  $K$ -steps of gradient descent on the aggregated function  $\alpha f(x, \cdot) + (1-\alpha)g(x, \cdot)$  for some  $\alpha \in (0, 1)$ . Liu et al. (2021b) defined  $\varphi_K(x) = \min_z \max_{k \leq K} \{f(x, \text{GD}^k(x; z))\}$  for nonconvex-nonconvex problems, where  $\text{GD}^k(x; z)$  denotes  $k$ -steps of gradient descent on  $g(x, \cdot)$  initialized with  $z$ . However, all these works require solving a series of sub-problems and often lack a non-asymptotic analysis.

## Appendix B. Future Works

Below, we discuss possible future directions.

**Problem Complexity.** Assumption 4.1d suppose  $f$  has  $\rho_f$ -Lipschitz Hessians, which comes from Lemma 4.3. Such an assumption is not required by nonconvex-strongly-convex problems (Chen et al., 2023). We wonder if it is possible to remove this assumption in our setting. Also, it is also interesting to study lower bounds for this problem (Kwon et al., 2024a).

**More Lower-Level Problems** It is also interesting to study bilevel problems with lower-level functions beyond the PL condition. One possible direction is to consider the Kurdyka-Łojasiewicz (KL) condition (Fatkhullin et al., 2022), which is more general than the PL condition and it holds for any semialgebraic functions. Recent works have shown non-asymptotic convergence for nonconvex-KL minimax optimization problems (Li et al., 2022; Zheng et al., 2023), but the nonconvex-KL bilevel optimization problem remains challenging. Besides these conditions, we hope future works to exploit the problem structure that arises from the applications of bilevel optimization and identify other tractable conditions that can break the limit of our hardness result.

## Appendix C. Discussion on Xiao et al. (2023)

The recently proposed GALET (Xiao et al., 2023) has been shown to converge with the rate  $\tilde{\mathcal{O}}(\epsilon^{-2})$  to the following definition of stationary points of Problem (11).

**Definition C.1** We say  $x^*$  is an  $\epsilon$ -stationary point of Problem (11) if  $\exists y^*, w^*$  such that

$$\|\mathcal{R}_x(x^*, y^*, w^*)\| \leq \epsilon, \quad \|\mathcal{R}_w(x^*, y^*, w^*)\| \leq \epsilon, \quad \mathcal{R}_y(x^*, y^*) \leq \epsilon^2,$$

where we define

$$\mathcal{R}_x(x, y, w) = \nabla_x f(x, y) + \nabla_{xy}^2 g(x, y)w;$$

$$\begin{aligned}\mathcal{R}_w(x, y, w) &= \nabla_{yy}^2 g(x, y) (\nabla_y f(x, y) + \nabla_{yy}^2 g(x, y) w); \\ \mathcal{R}_y(x, y) &= g(x, y) - \min_{z \in \mathbb{R}^{d_y}} g(x, z).\end{aligned}$$

We can show the  $\epsilon$ -stationary point of  $\varphi(x)$  implies the above stationarity definition.

**Proposition C.1** *Under Assumption 4.1, if  $x^*$  is an  $\epsilon$ -stationary point of  $\varphi(x)$  in Equation (2), then it is also an  $\mathcal{O}(\epsilon)$ -stationary point as Definition C.1.*

**Proof** Given  $x^*$ , take any  $y^* \in Y^*(x)$ , and let  $w^* = -(\nabla_{yy}^2 g(x^*, y^*))^\dagger \nabla_y f(x^*, y^*)$ . Then we have both  $\mathcal{R}_y(x^*, y^*) = 0$  and  $\mathcal{R}_w(x^*, y^*, w^*) = 0$ . By Lemma G.7, we know that  $\mathcal{R}_x(x^*, y^*, w^*) = \nabla \varphi(x^*)$ , therefore we also have  $\|\mathcal{R}_x(x^*, y^*, w^*)\| \leq \epsilon$ . ■

Furthermore, the converse relationship also holds.

**Proposition C.2** *Under Assumption 4.1, if  $\hat{x}$  is an  $\epsilon$ -stationary point as Definition C.1, then it is also an  $\mathcal{O}(\epsilon)$ -stationary point of  $\varphi(x)$ .*

**Proof** Let  $y^*, w^*$  be the corresponding point of  $x^*$  in Definition C.1. By Lemma G.2,

$$\text{dist}(y^*, Y^*(x)) \leq \sqrt{\frac{\mu}{2} \mathcal{R}_y(x^*, y^*)} \leq \sqrt{\frac{\mu}{2}} \cdot \epsilon.$$

Take  $\hat{y} = \arg \min_{y \in Y^*(x)} \|y^* - y\|$  and define the following auxiliary function:

$$\mathcal{L}(w) := \frac{1}{2} \|\nabla_{yy}^2 g(x^*, \hat{y}) w + \nabla_y f(x^*, \hat{y})\|^2. \quad (12)$$

Note that

$$\nabla \mathcal{L}(w) = \nabla_{yy}^2 g(x^*, \hat{y}) (\nabla_y f(x^*, \hat{y}) + \nabla_{yy}^2 g(x^*, \hat{y}) w).$$

When  $\|w^*\| = \mathcal{O}(1)$ , we have  $\|\nabla \mathcal{L}(w^*) - \mathcal{R}_w(x^*, y^*, w^*)\| = \mathcal{O}(\epsilon)$ .

$\mathcal{L}(w)$  has the form of "strongly convex composed with linear". It is  $s^2$ -PL by Karimi et al. (2016), Appendix B, where  $s$  is the smaller singular value of  $\nabla^2 g(x^*, \hat{y})$ . By Lemma G.6,  $s \geq \mu$ .

Let  $W^* = \arg \min_{w \in \mathbb{R}^{d_y}} \mathcal{L}(w)$ . Then by Lemma G.1,

$$\text{dist}(w^*, W^*) \leq \frac{1}{s^2} \|\nabla \mathcal{L}(w^*)\| = \mathcal{O}(\epsilon).$$

Note that  $W^*$  has the explicit form of

$$W^* = -(\nabla_{yy}^2 g(x^*, \hat{y}))^\dagger \nabla_y f(x^*, \hat{y}) + \text{Ker}(\nabla_{yy}^2 g(x^*, \hat{y})).$$

Take  $\hat{w} = \arg \min_{w \in W^*} \|w - w^*\|$  and by Lemma G.3,

$$\nabla \varphi(x^*) = \mathcal{R}_x(x^*, \hat{y}, \hat{w}) = \nabla_x f(x^*, \hat{y}) + \nabla_{xy}^2 g(x^*, \hat{y}) \hat{w}.$$

Since  $w^*$  is bounded by Lemma G.6, we know that

$$\|\nabla \varphi(x^*)\| \leq \|\mathcal{R}_x(x^*, y^*, w^*)\| + \|\mathcal{R}_x(x^*, \hat{y}, \hat{w}) - \mathcal{R}_x(x^*, y^*, w^*)\| = \mathcal{O}(\epsilon). \quad \blacksquare$$

**Remark C.1** In above analysis, the optimal  $w^*$  of GALET (Xiao et al., 2023) is the solution to the linear equation  $\nabla_{yy}^2 g(x, y^*) w = \nabla_y g(x, y^*)$ . However, the Hessian matrix  $\nabla_{yy}^2 g(x, y)$  may be indefinite when the lower-level function is nonconvex. To overcome this issue, Xiao et al. (2023) uses the square trick like Liu and Luo (2022), which solves Equation (12) instead. One drawback of the square trick is the conditional number becomes  $\mathcal{O}(\kappa^2)$  and makes the inner loop slower. Since  $\varphi(x)$  has  $\mathcal{O}(\ell\kappa^3)$ -Lipschitz gradients by Lemma 4.4. According to our above analysis, the reasonable complexity of GALET should be  $\tilde{\mathcal{O}}(\ell\kappa^5\epsilon^{-2})$  though the dependency on  $\kappa$  is not explicitly given by Xiao et al. (2023).

**Remark C.2** Xiao et al. (2023) additionally assumes the smallest singular value of  $\nabla_{yy}^2 g(x, y)$  has a constant gap between zero when solving Equation (12). However, the complexity to solve Equation (12) only depends on the singular value of  $\nabla_{yy}^2 g(x, y^*)$ , which can be bound below by our Lemma G.6. Since  $g(x, y)$  is PL in  $y$ , by Lemma G.8 gradient descent on  $g$  can make  $\text{dist}(y^k, Y^*(x))$  converges linearly. So one can let  $y^K$  sufficiently close to some point in  $Y^*(x)$  with logarithmic large  $K$ . Then by Weyl's inequality the singular value of  $\nabla_{yy}^2 g(x, y^K)$  is also bounded below. Therefore, such an additional assumption in Xiao et al. (2023) seems to be redundant.

## Appendix D. Discussion on Kwon et al. (2024b)

Different from our Assumption 4.1a, Kwon et al. (2024b) uses the following proximal error bound (Prox-EB) assumption.

**Assumption D.1** Let  $h_\sigma(x, y) := \sigma f(x, y) + g(x, y)$  and  $Y_\sigma^*(x) := \arg \min_{y \in \mathbb{R}^{d_y}} h_\sigma(x, y)$ . Suppose that for all  $0 \leq \sigma \leq \bar{\sigma}$ , there exists some  $\mu' > 0$  such that for all  $y \in \mathbb{R}^{d_y}$  we have

$$\rho^{-1} \|y - y_{\sigma, \rho}^+(x)\| \geq \mu' \text{dist}(y, Y_\sigma^*(x)),$$

where  $y_{\sigma, \rho}^+(x)$  is defined via the proximal operator with parameter  $\rho$  for function  $h_\sigma(x, \cdot)$  as

$$y_{\sigma, \rho}^+(x) := \arg \min_{z \in \mathbb{R}^{d_y}} \left\{ h_\sigma(x, z) + \frac{1}{2\rho} \|y - z\|^2 \right\}.$$

Below, we show that this assumption is equivalent to our Assumption 4.1a.

**Proposition D.1** Let  $0 \leq \sigma \leq \min\{L_g/L_f, \bar{\sigma}\}$ . Suppose both  $c$  and  $d$  in Assumption 4.1 hold. Let  $\rho < 1/(2L_g)$ . If Assumption 4.1a holds with constant  $\mu$ , then Assumption D.1 holds with  $\mu' = \mu/(1 + 2L_g\rho)$ . Conversely, if Assumption D.1 holds with constant  $\mu'$ , then Assumption 4.1a holds with  $\mu = (\mu'(1 - 2L_g\rho))^2 / (2L_g)$ .

**Proof** Let  $\sigma \leq L_g/L_f$  then  $h_\sigma(x, y)$  has  $(2L_g)$ -Lipschitz gradients. Define the Moreau envelope for  $h_\sigma(x, y)$  with respect to  $y$  as  $h_{\sigma, \rho}(x, y) := h_\sigma(x, y_{\sigma, \rho}^+(x))$ . Let  $\rho < 1/(2L_g)$  then  $y_{\sigma, \rho}^+(x)$  is uniquely defined. Danskin's theorem implies that  $\nabla_y h_{\sigma, \rho}(x, y) = \rho^{-1}(y - y_{\sigma, \rho}^+(x))$ . Therefore the Prox-EB assumption is equivalent to  $\|\nabla_y h_{\sigma, \rho}(x, y)\| \geq \mu' \text{dist}(y, Y_\sigma^*(x))$ . Note that

$$\begin{aligned} & \|\nabla_y h_{\sigma, \rho}(x, y) - \nabla_y h_\sigma(x, y)\| \\ &= \|\nabla_y h_\sigma(x, y_{\sigma, \rho}^+(x)) - \nabla_y h_\sigma(x, y)\| \\ &\leq 2L_g \|y - y_{\sigma, \rho}^+(x)\| \end{aligned}$$

$$= 2L_g\rho\|\nabla_y h_{\sigma,\rho}(x, y)\|.$$

When  $\rho < 1/(2L_g)$  the triangle inequality implies

$$(1 - 2L_g\rho)\|\nabla_y h_{\sigma,\rho}(x, y)\| \leq \|\nabla_y h_\sigma(x, y)\| \leq (1 + 2L_g\rho)\|\nabla_y h_{\sigma,\rho}(x, y)\|.$$

Therefore, if  $\mu$ -PL condition holds, by Lemma G.1, we have

$$\|\nabla_y h_{\sigma,\rho}(x, y)\| \geq \frac{1}{1 + 2L_g\rho}\|\nabla_y h_\sigma(x, y)\| \geq \frac{\mu}{1 + 2L_g\rho}\text{dist}(y, Y_\sigma^*(x)).$$

Conversely, if  $\mu'$ -Prox-EB condition holds, then we have

$$\|\nabla_y h_\sigma(x, y)\| \geq (1 - 2L_g\rho)\|\nabla_y h_{\sigma,\rho}(x, y)\| \geq \mu'(1 - 2L_g\rho)\text{dist}(y, Y_\sigma^*(x)).$$

Since  $h_\sigma(x, y)$  has  $(2L_g)$ -Lipschitz gradients, then

$$h_\sigma(x, y) - h_\sigma^*(x) \leq L_g\text{dist}^2(y, Y_\sigma^*(x)) \leq \frac{L_g}{(\mu')^2(1 - 2L_g\rho)^2}\|\nabla_y h_\sigma(x, y)\|^2,$$

which is exactly the PL inequality. ■

## Appendix E. Proof of Theorem 3.1

**Theorem 3.1** *Suppose that for any given  $x \in \mathbb{R}^{d_x}$  the set  $Y^*(x)$  is non-empty and compact.*

- (a) *If  $f(x, y)$  and  $Y^*(x)$  are continuous, then  $\varphi(x)$  is continuous.*
- (b) *Conversely, if  $\varphi(x)$  is continuous for any continuous  $f(x, y)$ , then  $Y^*(x)$  is continuous.*
- (c) *If  $f(x, y)$  and  $Y^*(x)$  are locally Lipschitz, then  $\varphi(x)$  is locally Lipschitz.*
- (d) *Conversely, if  $\varphi(x)$  is locally Lipschitz for any locally Lipschitz function  $f(x, y)$ , then  $Y^*(x)$  is locally Lipschitz.*
- (e) *If  $f(x, y)$  is  $C_f$ -Lipschitz and  $Y^*(x)$  is  $\kappa$ -Lipschitz, then we have that  $\varphi(x)$  is  $C_\varphi$ -Lipschitz with  $C_\varphi = (\kappa + 1)C_f$ .*
- (f) *Conversely, if  $\varphi(x)$  is  $C_\varphi$ -Lipschitz for any  $C_f$ -Lipschitz  $f(x, y)$ , then we know that  $Y^*(x)$  is  $\kappa$ -Lipschitz with  $\kappa = C_\varphi/C_f$ .*

**Proof** For given  $x_1, x_2$ , define

$$d_1 := \max_{y_2 \in Y^*(x_2)} \text{dist}(Y^*(x_1), y_2), \quad d_2 := \max_{y_1 \in Y^*(x_1)} \text{dist}(y_1, Y^*(x_2)).$$

Note that we can replace sup with max in Definition 3.1 due to the compactness of  $Y^*(x)$ . Therefore,  $\text{dist}(Y^*(x_1), Y^*(x_2)) = \max\{d_1, d_2\}$ . Below we prove each part of the theorem.

- (a). See Theorem 3B.5 (Dontchev et al., 2009).

(b). It suffices to show for any given  $x_1 \in \mathbb{R}^{d_x}$  and any  $\epsilon > 0$ , there exists  $\delta > 0$  such that for any  $x_2$  satisfying  $\|x_1 - x_2\| \leq \delta$  both  $d_1$  and  $d_2$  are no larger than  $\epsilon$ . We prove this by assigning different  $f(x, y)$  and then applying the continuity of  $\varphi(x) := \min_{y \in Y^*(x)} f(x, y)$ .

Firstly, take  $f(x, y) = -\text{dist}(y, Y^*(x_1))$ . Simple calculus shows  $\varphi(x_1) = 0$  and  $\varphi(x_2) = -d_1$ . By the continuity of  $\varphi(x)$  at  $x_1$ , we know that for given  $\epsilon > 0$ , there exists  $\delta_1 > 0$ , such that for any  $x_2$  satisfying  $\|x_1 - x_2\| \leq \delta_1$ , we have  $d_1 = \varphi(x_1) - \varphi(x_2) \leq \epsilon$ .

Secondly, we want to prove that for any  $\epsilon > 0$  there exists  $\delta_2 > 0$  such that for any  $x_2$  satisfying  $\|x_1 - x_2\| \leq \delta_2$  we have  $d_2 \leq \epsilon$ . We prove this by contradiction. Suppose not, then we can find a sequence  $\{x_n\}$  such that  $x_n \rightarrow x_1$ , but  $\max_{y_1 \in Y^*(x_1)} \text{dist}(y_1, Y^*(x_n)) \geq \epsilon$  for some  $\epsilon > 0$ . We take the corresponding  $y_n = \arg \max_{y_1 \in Y^*(x_1)} \text{dist}(y_1, Y^*(x_n))$ . Since  $\{y_n\}$  is a bounded sequence, there exists a convergent subsequence  $\{y_{n_k}\}$  with some limit point  $y'_1 \in Y^*(x_1)$ . Take  $n^*$  sufficiently large such that for any  $n \geq n^*$  we have  $\|y_{n_k} - y'_1\| \leq \epsilon/2$ . Then by triangle inequality, for any  $n \geq n^*$  we have  $\text{dist}(y'_1, Y^*(x_{n_k})) \geq \epsilon/2$ . Now, take  $f(x, y) = \|y - y'_1\|$ . Simple calculus shows that  $\varphi(x_1) = 0$  and  $\varphi(x_{n_k}) = \text{dist}(y'_1, Y^*(x_{n_k}))$ . However,  $|\varphi(x_{n_k}) - \varphi(x_1)| \geq \epsilon/2$  for the sequence  $\{x_{n_k} : n_k \geq n^*\}$  satisfying  $x_{n_k} \rightarrow x_1$ . This contradicts the continuity of  $\varphi(x)$  at  $x_1$ .

Finally, we take  $\delta = \min\{\delta_1, \delta_2\}$  and conclude that once  $\|x_1 - x_2\| \leq \delta$  we have both  $d_1$  and  $d_2$  are smaller than  $\epsilon$ , implying the continuity of  $Y^*(x)$ .

(c). It suffices to show for any given  $x \in \mathbb{R}^{d_x}$ , there exists  $\delta > 0$  and  $L > 0$  such that  $\varphi(x)$  is  $L$ -Lipschitz on  $\mathbb{B}_\delta(x_1)$ . Firstly, the local Lipschitz continuity of  $Y^*(\cdot)$  implies the existence of  $\delta > 0$  and  $L_1 > 0$  such that  $Y^*(\cdot)$  is  $L_1$ -Lipschitz on  $\mathbb{B}_\delta(x_1)$ . Next, for any  $x_2 \in \mathbb{B}_\delta(x_1)$ , we pick

$$y_1 \in \arg \min_{y \in Y^*(x_1)} f(x_1, y), \quad y_2 \in \arg \min_{y \in Y^*(x_2)} f(x_2, y). \quad (13)$$

There exist  $y'_1 \in Y^*(x_1)$  and  $y'_2 \in Y^*(x_2)$  such that

$$\|y'_1 - y_2\| \leq L_1 \|x_1 - x_2\|, \quad \|y_1 - y'_2\| \leq L_1 \|x_1 - x_2\|.$$

Therefore, both  $y_2, y'_2$  lie in the compact set  $\mathcal{N}_y(x_1) = \{y : \text{dist}(y, Y^*(x_1)) \leq L_1 \delta\}$ .

The local Lipschitz property of  $f(x, y)$  implies that there exists  $L_2 > 0$  such that  $f(x, y)$  is  $L_2$ -Lipschitz on the set  $\mathbb{B}_\delta(x_1) \times \mathcal{N}_y(x_1)$ . Then

$$\begin{aligned} \varphi(x_1) - \varphi(x_2) &\leq f(x_1, y'_1) - f(x_2, y_2) \leq L_2 (\|x_1 - x_2\| + \|y_2 - y'_1\|) \leq (L_1 + 1)L_2 \|x_1 - x_2\|. \\ \varphi(x_2) - \varphi(x_1) &\leq f(x_2, y'_2) - f(x_1, y_1) \leq L_2 (\|x_1 - x_2\| + \|y_1 - y'_2\|) \leq (L_1 + 1)L_2 \|x_1 - x_2\|. \end{aligned}$$

This implies that  $\varphi(x)$  is Lipschitz on  $\mathbb{B}_\delta(x_1)$ .

(d). For any compact set  $K \subseteq \mathbb{R}^{d_x}$ , there exists  $L > 0$  such that  $\varphi(x)$  is  $L$ -Lipschitz on  $K$ . Then for any  $x_1, x_2 \in K$ , taking  $f(x, y) = -\text{dist}(y, Y^*(x_1))$  yields

$$d_1 = \varphi(x_1) - \varphi(x_2) \leq L \|x_1 - x_2\|,$$

By symmetric, we can also show that  $d_2 \leq L \|x_1 - x_2\|$ . Combining them, we show that  $Y^*(x)$  is Lipschitz on any compact set  $K$ . This finishes the proof.

(e). Pick  $y_1, y_2$  as Equation (13). Similarly, there exist  $y'_1 \in Y^*(x_1)$  and  $y'_2 \in Y^*(x_2)$  such that

$$\begin{aligned} \varphi(x_1) - \varphi(x_2) &\leq f(x_1, y'_1) - f(x_2, y_2) \leq C_f (\|x_1 - x_2\| + \|y_2 - y'_1\|) \leq (\kappa + 1)C_f \|x_1 - x_2\|, \\ \varphi(x_2) - \varphi(x_1) &\leq f(x_2, y'_2) - f(x_1, y_1) \leq C_f (\|x_1 - x_2\| + \|y_1 - y'_2\|) \leq (\kappa + 1)C_f \|x_1 - x_2\|, \end{aligned}$$

This establishes the Lipschitz continuity of  $\varphi(x)$ .

(f). Without loss of generality, we assume  $C_f = 1$ , otherwise we can scale  $f(x, y)$  by  $C_f$  to prove the result. Because  $f(x, y)$  is globally Lipschitz, we can take  $K = \mathbb{R}^{d_x}$  in the proof of **d**. Then by the same arguments, we can show that  $Y^*(x)$  is  $C_\varphi$ -Lipschitz. ■

## Appendix F. Proof of Theorem 3.2

Our hard instance is based on the following convex zero-chain. The following function is very similar to the worst-case function given by (Nesterov, 2018), Section 2.1.2. The only difference is that the function by Nesterov (2018) has an additional term  $z_{[q]}^2/8$ .

**Definition F.1 (Worse-Case Zero-Chain)** Consider the family of functions:

$$h_q(z) = \frac{1}{8}(z_{[1]} - 1)^2 + \frac{1}{8} \sum_{j=1}^{q-1} (z_{[j+1]} - z_{[j]})^2.$$

The following properties hold for any  $h_q(z)$  with  $q \in \mathbb{N}^+$ :

- a. It has a unique minimizer  $z^* = \mathbf{1}$ .
- b. It is convex.
- c. It has 1-Lipschitz gradients.
- d. It is a first-order zero-chain, i.e. for any  $z \in \mathbb{R}^q$ ,

$$\text{supp}\{z\} \in \{1, 2, \dots, j\} \Rightarrow \text{supp}\{\nabla h(z)\} \in \{1, 2, \dots, j+1\}.$$

**Proof** We prove each property one by one. It can easily be seen that  $h_q(z) \geq 0$  for all  $z \in \mathbb{R}^q$  and the equality holds if and only if  $z = \mathbf{1}$ . This proves property **a**. Further, note that  $h_q(z)$  is quadratic with Hessian given by

$$A = \frac{1}{4} \begin{bmatrix} 2 & -1 & & & \\ -1 & 2 & -1 & & \\ & -1 & 2 & -1 & \\ & & \ddots & \ddots & \ddots \\ & & & -1 & 1 \end{bmatrix}.$$

As  $A$  is diagonally dominant, we know that  $A \succeq O$ . This proves property **b**. For any  $v \in \mathbb{R}^q$ ,

$$\begin{aligned} v^\top A v &= \frac{1}{4} \left[ v_{[1]}^2 + \sum_{j=1}^{q-1} (v_{[j]} - v_{[j+1]})^2 \right] \\ &\leq \frac{1}{4} \left[ v_{[1]}^2 + \sum_{j=1}^{q-1} (v_{[j]} - v_{[j+1]})^2 + v_{[q]}^2 \right] \end{aligned}$$

$$\begin{aligned} &\leq \frac{1}{4} \left[ v_{[1]}^2 + \sum_{j=1}^{q-1} 2(v_{[j]}^2 + v_{[j+1]}^2) + v_{[q]}^2 \right] \\ &\leq \sum_{j=1}^q v_{[j]}^2 = \|v\|^2. \end{aligned}$$

This proves property **c**. Finally property **d** holds since  $A$  is tridiagonal. ■

In bilevel problems, it is crucial to find a point  $y$  that is close to  $Y^*(x)$ , instead of just achieving a small optimality gap  $g(x, y) - g^*(x)$ . However, it is difficult for any first-order algorithms to “locate” the minimizers of the function class in Definition F.1. Below, we formalize this observation into a rigorous statement.

**Theorem 3.2** *Without loss of generality, suppose  $x_0 = y_0 = 0$  (otherwise we can translate the functions and the result still holds). Fix  $T$  and  $K$ . Let  $d_x = 1$ ,  $d_y = q = 2TK$ , and*

$$f(x, y) = 2(x + 1)^2 \sum_{j=q/2}^q \psi(y_{[j]}), \quad g(y) = \frac{1}{8}(y_{[1]} - 1/\sqrt{q})^2 + \frac{1}{8} \sum_{j=1}^{q-1} (y_{[j+1]} - y_{[j]})^2,$$

where  $\psi(\cdot) : \mathbb{R} \rightarrow \mathbb{R}$  (defined in Equation (14)) is a function with  $\psi'(0) = \psi(0) = 0$ . Define the sublevel set  $\mathcal{X} := \{x : \varphi(x) \leq \varphi(0)\}$ . There exists numerical constants  $c_1, c_2 > 0$  such that

1.  $f(x, y)$  is  $c_1$ -Lipschitz in  $y$  on  $\mathcal{X} \times \mathbb{R}^{d_y}$ ;
2.  $f(x, y)$  has  $c_2$ -Lipschitz gradients on  $\mathcal{X} \times \mathbb{R}^{d_y}$ ;
3.  $g(y)$  is a strictly convex quadratic and has 1-Lipschitz gradients;
4. The resulting hyper-objective is  $\varphi(x) = (x + 1)^2/2$ .

For this problem, any algorithm with a procedure as Algorithm 1 stays at  $x_t = 0$  for any iteration number  $t \leq T$ .

**Proof** Let  $d_x = 1$ ,  $d_y = q = 2TK$ ,  $\beta = 1/\sqrt{q}$  and

$$f(x, y) = 2(x + 1)^2 r(y), \quad g(y) = \beta^2 h_q(y/\beta).$$

where  $h_q(y)$  follows Definition F.1 and  $r(y) = \sum_{j=q/2+1}^q \psi(y_{[j]})$ , where  $\psi : \mathbb{R} \rightarrow \mathbb{R}$  is

$$\psi(y) = \begin{cases} \beta^2, & y > 2\beta; \\ p(y), & \beta < y \leq 2\beta; \\ y^2/2, & -\beta \leq y \leq \beta; \\ p(-y), & -2\beta \leq y < -\beta; \\ \beta^2, & y < -2\beta; \end{cases} \quad (14)$$



where

$$p(y) = -\frac{y^5}{2\beta^3} + \frac{9y^4}{2\beta^2} - \frac{31y^3}{2\beta} + 25y^2 - 18\beta y + 5\beta^2$$

is the Hermite interpolating polynomial that satisfies

$$p(\beta) = \beta^2/2, \quad p'(\beta) = \beta, \quad p''(\beta) = 1 \quad \text{and} \quad p(2\beta) = \beta^2, \quad p'(2\beta) = 0, \quad p''(2\beta) = 0.$$

There must exist numerical constants  $\gamma_0, \gamma_1, \gamma_2, \gamma_3$  such that

$$0 < \psi(y) \leq \gamma_0\beta^2, \quad |\psi'(y)| \leq \gamma_1\beta, \quad |\psi''(y)| \leq \gamma_2, \quad |\psi'''(y)| \leq \gamma_3/\beta.$$

We can verify that  $r(y)$  is both bounded and Lipschitz because

$$r(y) = \sum_{j=q/2+1}^q \psi(y_{[j]}) \leq \frac{\gamma_0}{2},$$

and

$$\|\nabla r(y)\| = \sqrt{\sum_{j=q/2+1}^q (\psi'(y_{[j]}))^2} \leq \frac{\gamma_1}{\sqrt{2}}.$$

Furthermore, we can also prove that  $r(y)$  has Lipschitz gradients. For any  $y, y' \in \mathbb{R}^{d_y}$ ,

$$\|\nabla r(y) - \nabla r(y')\| = \sqrt{\sum_{j=q/2+1}^q (\psi'(y_{[j]}) - \psi'(y'_{[j]}))^2} \leq \gamma_2 \sqrt{\sum_{j=q/2+1}^q (y_{[j]} - y'_{[j]})^2} \leq \gamma_2 \|y - y'\|.$$

Note that  $\nabla_x f(x, y) = 4(x+1)r(y)$ ,  $\nabla_y f(x, y) = 2(x+1)^2 \nabla r(y)$ . For any  $x \in \mathcal{X} = [-2, 0]$ ,

$$\begin{aligned} & |\nabla_x f(x, y) - \nabla_x f(x', y')| \\ & \leq 4|x - x'| \cdot r(y) + 4|x' + 1| \cdot |r(y) - r(y')| \\ & \leq 2\gamma_0|x - x'| + 2\sqrt{2} \gamma_1 \|y - y'\|. \end{aligned}$$

as well as

$$\begin{aligned} & \|\nabla_y f(x, y) - \nabla_y f(x', y')\| \\ & \leq 2((x+1)^2 - (x'+1)^2) \cdot \|\nabla r(y)\| + 2(x'+1)^2 \cdot \|\nabla r(y) - \nabla r(y')\| \\ & \leq 2\sqrt{2} \gamma_1 |x - x'| + 2\gamma_2 \|y - y'\|. \end{aligned}$$

These two inequalities imply  $f(x, y)$  has Lipschitz gradients on  $\mathcal{X} \times \mathbb{R}^{d_y}$ .

Below, prove by induction that  $x_t = 0$ , and  $y_{t,[j]}^k = 0$  for all  $j > tK + k$ .

Suppose  $x_t = 0$ , then  $\nabla_y f(x_t, y) = 2\nabla r(y)$ . If we have  $y_{t,[j]}^k = 0$  for all  $j > tK + k$ , then we have  $\nabla r(y)_{[j]} = \psi'(y_{[j]}) = 0$  for all these coordinates  $j$ . By the property of zero-chain  $g(y)$ , we also have  $\nabla g(y)_{[j]} = 0$  for all  $j > tK + k + 1$ . This indicates that each inner loop iteration step  $k$  increases  $y_t^k$  by at most one non-zero coordinate. Since there are at most  $q/2$  iterations for  $y$  in

total, the last  $q/2$  coordinates of  $y_t^k$  will always remain zero. Since  $r(y)$  only depends on the last  $q/2$  coordinates, we have  $\nabla_x f(x_t, y_t^K) = 4 \sum_{j=q/2}^q \psi(y_{t,[j]}^K) = 0$ . Then by the update rule on  $x_t$ , we have  $x_{t+1} = x_t$  remains unchanged.

The optimal solution of the lower-level function is unique and given by  $y^* = \beta \mathbf{1}$ . By  $r(y^*) = 1/4$ , we know that  $\varphi(x) = (x + 1)^2/2$ . ■

## Appendix G. Proof of Theorem 4

First of all, we recall some useful lemmas for PL conditions (Karimi et al., 2016).

**Lemma G.1** For  $\mu$ -PL function  $h(x) : \mathbb{R}^d \rightarrow \mathbb{R}$  with Lipschitz gradients, for any  $x \in \mathbb{R}^d$ ,

$$\|\nabla h(x)\| \geq \mu \text{dist}(x, X^*), \quad \text{where } X^* = \arg \min_{x \in \mathbb{R}^d} h(x).$$

**Lemma G.2** For a  $\mu$ -PL function  $h(x) : \mathbb{R}^d \rightarrow \mathbb{R}$  with Lipschitz gradient, for any  $x \in \mathbb{R}^d$ ,

$$h(x) - \min_{x \in \mathbb{R}^d} h(x) \geq \frac{\mu}{2} \text{dist}^2(x, X^*).$$

The above two lemmas appear in Theorem 2 (Karimi et al., 2016). Based on these results, we can prove the following lemma.

**Lemma 4.1** Let  $Y_\sigma^* := \arg \min_{y \in \mathbb{R}^{d_y}} h_\sigma(x, y)$  denote the set of minima for the penalty function  $h_\sigma(x, y) = \sigma f(x, y) + g(x, y)$ . Under Assumption 4.1, we have that

$$\text{dist}(Y_{\sigma_1}^*(x_1), Y_{\sigma_2}^*(x_2)) \leq \frac{C_f}{\mu} |\sigma_1 - \sigma_2| + \frac{\sigma L_f + L_g}{\mu} \|x_1 - x_2\|.$$

**Proof** By Lemma G.1, for any  $y_1 \in Y_{\sigma_1}^*(x_1)$ , there exists some  $y_2 \in Y_{\sigma_2}^*(x_2)$  such that

$$\begin{aligned} & \mu \|y_1 - y_2\| \\ & \leq \|\nabla_y h_{\sigma_2}(x_2, y_1)\| \\ & = \|\nabla_y h_{\sigma_2}(x_2, y_1) - \nabla_y h_{\sigma_1}(x_1, y_1)\| \\ & \leq \|\sigma_2 \nabla_y f(x_2, y_1) - \sigma_1 \nabla_y f(x_1, y_1)\| + \|\nabla_y g(x_2, y_1) - \nabla_y g(x_1, y_1)\| \\ & \leq \|\sigma_2 \nabla_y f(x_2, y_1) - \sigma_1 \nabla_y f(x_2, y_1)\| \\ & \quad + \|\sigma_1 \nabla_y f(x_2, y_1) - \sigma_1 \nabla_y f(x_1, y_1)\| + \|\nabla_y g(x_2, y_1) - \nabla_y g(x_1, y_1)\| \\ & \leq |\sigma_1 - \sigma_2| C_f + (\sigma L_f + L_g) \|x_1 - x_2\|. \end{aligned}$$

By symmetry, for any  $y_2 \in Y_{\sigma_2}^*(x_2)$ , there also exists  $y_1 \in Y_{\sigma_1}^*(x_1)$  such that that the above inequality holds for  $y_2, y_1$ . This proves the Pompeiu–Hausdorff continuity. ■

**Lemma 4.2 (Shen and Chen (2023))** Recall that  $\varphi_\sigma(x)$  is the penalized hyper-objective defined in Equation (6). Under Assumption 4.1,  $\nabla \varphi_\sigma(x)$  exists and takes the form of

$$\nabla \varphi_\sigma(x) = \nabla_x f(x, y) + \frac{\nabla_x g(x, y_\sigma^*(x)) - \nabla_x g(x, y^*(x))}{\sigma}, \quad (7)$$

where  $y^*(x), y_\sigma^*(x)$  can be arbitrary elements in  $Y^*(x)$  and  $Y_\sigma^*(x)$ , respectively.

**Proof** It follows the generalized Danskin's theorem proved in [Shen and Chen \(2023\)](#). See also Lemma A.2 in [Kwon et al. \(2024b\)](#). ■

**Lemma 4.3** ([Kwon et al. \(2024b\)](#)) *Recall that  $\varphi(x)$  is the original hyper-objective in Equation (2), while  $\varphi_\sigma(x)$  is the penalized hyper-objective in Equation (6). Under Assumption 4.1,  $\nabla\varphi(x)$  exists and can be defined as the limit  $\lim_{\sigma \rightarrow 0^+} \nabla\varphi_\sigma(x)$ . Moreover,  $\varphi_\sigma(x)$  is close to  $\varphi(x)$ . Formally, for any  $0 \leq \sigma \leq \min\{\rho_g/\rho_f, \bar{\sigma}\}$ , we have that*

$$|\varphi_\sigma(x) - \varphi(x)| = \mathcal{O}(\sigma\ell\kappa), \quad \text{and} \quad \|\nabla\varphi_\sigma(x) - \nabla\varphi(x)\| = \mathcal{O}(\sigma\ell\kappa^3).$$

**Proof** The proof follows Theorem 3.8 [Kwon et al. \(2024b\)](#). The only difference is that Theorem 3.8 [Kwon et al. \(2024b\)](#) states  $\|\nabla\varphi_\sigma(x) - \varphi(x)\| = \mathcal{O}(\sigma\ell\kappa^5)$ . But the additional  $\kappa^2$  dependency comes from the perturbation in the multiplier  $d\lambda/d\sigma \asymp \kappa^2$ . Since we only consider the unconstrained case, there is no need to consider the effect of  $d\lambda/d\sigma$ , and we can improve the bound to  $\mathcal{O}(\sigma\ell\kappa^3)$ . ■

We also recall some technical lemmas from [Kwon et al. \(2024b\)](#).

**Lemma G.3** *Suppose  $Y_\sigma^*(x)$  is Pompeiu–Hausdorff Lipschitz, then for any  $y_\sigma^*(x) \in Y_\sigma^*(x)$ ,*

$$\begin{aligned} \text{Range}(\nabla_{yx}^2 h_\sigma(x, y_\sigma^*(x))) &\subseteq \text{Range}(\nabla_{yy}^2 h_\sigma(x, y_\sigma^*(x))) \\ \nabla_y f(x, y_\sigma^*(x)) &\in \text{Range}(\nabla_{yy}^2 h_\sigma(x, y_\sigma^*(x))). \end{aligned}$$

**Proof** See Proposition 3.1 ([Kwon et al., 2024b](#)). ■

We remark that Proposition 6 ([Arbel and Mairal, 2022](#)) also presents a similar argument as the above lemma for Morse-Bott functions.

**Lemma G.4** *Under Assumption 4.1, there exists some  $\sigma \in [0, \sigma']$  such that*

$$\nabla\varphi_{\sigma'}(x) = \nabla_x f(x, y_\sigma^*(x)) - \nabla_{xy}^2 h_\sigma(x, y_\sigma^*(x)) (\nabla_{yy}^2 h_\sigma(x, y_\sigma^*(x)))^\dagger \nabla_y f(x, y_\sigma^*(x)).$$

for any  $y_\sigma^*(x) \in Y_\sigma^*(x)$ .

**Proof** We can express  $\varphi_{\sigma'}(x)$  by

$$\varphi_{\sigma'}(x) = \frac{l(x, \sigma') - l(x, 0)}{\sigma'} = \frac{\partial}{\partial\sigma} l(x, \sigma), \quad \exists \sigma \in [0, \sigma'],$$

where  $l(x; \sigma') = \min_y h_{\sigma'}(x, y)$  and we apply the mean-value theorem in the second equality. Taking derivative with respect to  $x$  in the above equation yields

$$\nabla\varphi_{\sigma'}(x) = \frac{\partial^2}{\partial x \partial \sigma} l(x, \sigma).$$

Finally, we plug in the explicit form of  $\frac{\partial^2}{\partial x \partial \sigma} l(x, \sigma)$  by Theorem 3.2 ([Kwon et al., 2024b](#)). ■

The following lemma is a fact from linear algebra.

**Lemma G.5** *If  $\text{Range}(U^\top) \subseteq \text{Range}(A^\top)$  and  $\text{Range}(V) \subseteq \text{Range}(B)$ , then*

$$U(A^\dagger - B^\dagger)V = UA^\dagger(B - A)B^\dagger V.$$

**Proof** If there exists some matrix  $P$  such that  $V = BP$ , then

$$V = BP = BB^\dagger BP = BB^\dagger V.$$

Similarly, if there exists some matrix  $Q$  such that  $U = QA$ , then  $U = UAA^\dagger$ . Combining these two identities completes the proof.  $\blacksquare$

Under the PL condition, the smallest eigenvalue of Hessian at any minimum is bounded below.

**Lemma G.6** *For a  $\mu$ -PL function  $h(x) : \mathbb{R}^d \rightarrow \mathbb{R}$  that is twice differentiable, at any  $x^* \in \arg \min_{x \in \mathbb{R}^d} h(x)$ ,*

$$\lambda_{\min}^+(\nabla^2 h(x^*)) \geq \mu,$$

where  $\lambda_{\min}^+(\cdot)$  denotes the smallest non-zero eigenvalue.

**Proof** Let  $\lambda_1, \lambda_2, \dots, \lambda_d$  be the eigenvalues of  $\nabla^2 h(x^*)$  in descending order, and  $v_1, v_2, \dots, v_d$  be the corresponding unit eigenvectors which are mutually orthogonal. Let  $r$  be the rank of  $\nabla^2 h(x^*)$ . Then  $\text{Span}(v_1, \dots, v_r) = \text{Range}(\nabla^2 h(x^*))$ , and  $\text{Span}(v_{r+1}, \dots, v_d) = \text{Ker}(\nabla^2 h(x^*))$ .

Let  $X^* = \arg \min_{x \in \mathbb{R}^d} h(x)$ ,  $x_t = x^* + tv_r$  and  $\hat{x}_t = \arg \min_{x \in X^*} \|x_t - x\|$ . There exist some coefficients  $\alpha_i$  such that  $\hat{x}_t - x^* = \sum_{i=1}^d \alpha_i v_i$ . By the Taylor's expansion

$$\begin{aligned} 0 &= h(\hat{x}_t) - h(x^*) \\ &= \frac{1}{2}(\hat{x}_t - x^*)^\top \nabla^2 h(x^*)(\hat{x}_t - x^*) + o(\|\hat{x}_t - x^*\|^2) \\ &= \frac{1}{2} \sum_{i=1}^r \lambda_i \alpha_i^2 + o(\|\hat{x}_t - x^*\|^2) \\ &\geq \frac{1}{2} \lambda_r \alpha_r^2 + o(\|\hat{x}_t - x^*\|^2). \end{aligned}$$

By triangle inequality and the definition of  $\hat{x}_t$ , we have

$$\|\hat{x}_t - x^*\| \leq \|x_t - x^*\| + \|x_t - \hat{x}_t\| \leq 2\|x_t - x^*\| = 2t.$$

Therefore,

$$\lambda_r \alpha_r^2 = o(\|\hat{x}_t - x^*\|^2) = o(t^2).$$

On the one hand, Lemma G.2 indicates

$$\begin{aligned} &h(x_t) - h(x^*) \\ &\geq \frac{\mu}{2} \|x_t - \hat{x}_t\|^2 \\ &= \frac{\mu}{2} \|x_t - x^* + x^* - \hat{x}_t\|^2 \end{aligned}$$

$$\begin{aligned}
 &= \frac{\mu}{2} \left( \|tv_r - \alpha_r v_r\|^2 + \left\| \sum_{i \neq r} \alpha_i v_i \right\|^2 \right) \\
 &\geq \frac{\mu}{2} (t - \alpha_r)^2.
 \end{aligned}$$

On the other hand, the Taylor's expansion also indicates

$$h(x_t) - h(x^*) = \frac{1}{2} (x_t - x^*)^\top \nabla^2 h(x^*) (x_t - x^*) + o(\|x_t - x^*\|^2) = \frac{\lambda_r}{2} t^2 + o(t^2).$$

Putting these two hands together

$$\frac{\mu}{2} (t - \alpha_r)^2 \leq \frac{\lambda_r}{2} t^2 + o(t^2).$$

Using  $\alpha_r = o(t)$  and letting  $t \rightarrow 0$ , we conclude that  $\lambda_r \geq \mu$ . ■

Lemma 4.3 only claims the existence of  $\nabla \varphi(x)$ . Below, we give the explicit form of  $\nabla \varphi(x)$ .

**Lemma G.7** *Under Assumption 4.1,*

$$\nabla \varphi(x) = \nabla_x f(x, y^*(x)) - \nabla_{xy}^2 h_\sigma(x, y^*(x)) (\nabla_{yy}^2 h(x, y^*(x)))^\dagger \nabla_y f(x, y^*(x)) \quad (15)$$

for any  $y^*(x) \in Y^*(x)$ .

**Proof** Let  $H(x)$  be the right-hand side of Equation (15). Below, we show that  $\nabla \varphi(x) = H(x)$ .

Recall Lemma G.4 that for any  $\sigma' \geq 0$  there exists some  $\sigma \in [0, \sigma']$  such that

$$\nabla \varphi_{\sigma'}(x) = \nabla_x f(x, y_\sigma^*(x)) - \nabla_{xy}^2 h_\sigma(x, y_\sigma^*(x)) (\nabla_{yy}^2 h_\sigma(x, y_\sigma^*(x)))^\dagger \nabla_y f(x, y_\sigma^*(x)).$$

for any  $y_\sigma^*(x) \in Y_\sigma^*(x)$ . Then for any  $y^*(x) \in Y^*(x)$ , by Lemma 4.1, there exists  $y_\sigma^*(x) \in Y_\sigma^*(x)$  such that  $\|y_\sigma^*(x) - y^*(x)\| \leq C_f \sigma / \mu$ . Then by Lemma G.3 and Lemma G.5, we have

$$\begin{aligned}
 &\|\nabla \varphi_{\sigma'}(x) - H(x)\| \\
 &\leq \|\nabla_x f(x, y_\sigma^*(x)) - \nabla_x f(x, y^*(x))\| \\
 &\quad + \left\| (\nabla_{xy}^2 h_\sigma(x, y_\sigma^*(x)) - \nabla_{xy}^2 g(x, y^*(x))) (\nabla_{yy}^2 h_\sigma(x, y_\sigma^*(x)))^\dagger \nabla_y f(x, y^*(x)) \right\| \\
 &\quad + \left\| \nabla_{xy}^2 g(x, y^*(x)) (\nabla_{yy}^2 g(x, y^*(x)))^\dagger \right. \\
 &\quad \quad \left. (\nabla_{yy}^2 h(x, y^*(x)) - \nabla_{yy}^2 h_\sigma(x, y_\sigma^*(x))) (\nabla_{yy}^2 h_\sigma(x, y_\sigma^*(x)))^\dagger \nabla_y f(x, y_\sigma^*(x)) \right\| \\
 &\quad + \left\| \nabla_{xy}^2 g(x, y^*(x)) (\nabla_{yy}^2 g(x, y^*(x)))^\dagger (\nabla_y f(x, y^*(x)) - \nabla_y f(x, y_\sigma^*(x))) \right\|
 \end{aligned}$$

We then use Lemma G.6 to have a further upper bound as

$$\|\nabla \varphi_{\sigma'}(x) - H(x)\| \leq \frac{\sigma C_f}{\mu} \left( 1 + \frac{L_g}{\mu} \right) \left( 2L_f + \frac{\rho_g C_f}{\mu} \right).$$

Taking  $\sigma' \rightarrow 0^+$ , we conclude

$$\nabla \varphi(x) = \lim_{\sigma' \rightarrow 0^+} \nabla \varphi_{\sigma'}(x) = H(x). \quad \blacksquare$$

**Lemma 4.4** Under Assumption 4.1,  $\varphi(x)$  has  $\mathcal{O}(\ell\kappa^3)$ -Lipschitz gradients.

**Proof** Invoking Lemma 4.1, there exists  $y_1 \in Y^*(x_1)$  and  $y_2 \in Y^*(x_2)$  such that  $\|y_1 - y_2\| \leq L_g/\mu$ . Then by Lemma G.3 and Lemma G.5, we have

$$\begin{aligned} & \|\nabla\varphi(x_1) - \nabla\varphi(x_2)\| \\ & \leq \|\nabla_x f(x_1, y_1) - \nabla_x f(x_2, y_2)\| \\ & \quad + \left\| \left( \nabla_{xy}^2 g(x_1, y_1) - \nabla_{xy}^2 g(x_2, y_2) \right) \left( \nabla_{yy}^2 g(x_1, y_1) \right)^\dagger \nabla_y f(x_1, y_1) \right\| \\ & \quad + \left\| \nabla_{xy}^2 g(x_2, y_2) \left( \nabla_{yy}^2 g(x_2, y_2) \right)^\dagger \right. \\ & \quad \quad \left. \left( \nabla_{yy}^2 g(x_2, y_2) - \nabla_{yy}^2 g(x_1, y_1) \right) \left( \nabla_{yy}^2 g(x_1, y_1) \right)^\dagger \nabla_y f(x_1, y_1) \right\| \\ & \quad + \left\| \nabla_{xy}^2 g(x_2, y_2) \left( \nabla_{yy}^2 g(x_2, y_2) \right)^\dagger \left( \nabla_y f(x_2, y_2) - \nabla_y f(x_1, y_1) \right) \right\|. \end{aligned}$$

Further invoking Lemma G.6,

$$\|\nabla\varphi(x_1) - \nabla\varphi(x_2)\| \leq \left( L_f + \frac{C_f \rho_g}{\mu} \right) \left( 1 + \frac{L_g}{\mu} \right) \left( 1 + \frac{L_g}{\mu} \right) \|x_1 - x_2\|.$$

■

The following lemma shows linear-convergence of gradient descent on PL functions.

**Lemma G.8** Suppose  $h(x) : \mathbb{R}^d \rightarrow \mathbb{R}$  is  $\alpha$ -PL and has  $\beta$ -Lipschitz gradients. Consider the following update of gradient descent:

$$x_{t+1} = x_t - \frac{1}{\beta} \nabla h(x_t).$$

Let  $X^* = \arg \min_{x \in \mathbb{R}^d} h(x)$  and  $h^* = \min_{x \in \mathbb{R}^d} h(x)$ . Then it holds that

$$\text{dist}^2(x_T, X^*) \leq \left( 1 - \frac{\alpha}{\beta} \right)^T \frac{\beta}{\alpha} \text{dist}^2(x_0, X^*).$$

**Proof** We first prove the linear convergence on the sub-optimality gap.

$$\begin{aligned} h(x_{t+1}) - h^* & \leq h(x_t) - h^* + \nabla h(x_t)^\top (x_{t+1} - x_t) + \frac{\beta}{2} \|x_{t+1} - x_t\|^2 \\ & = h(x_t) - h^* - \frac{1}{2\beta} \|\nabla h(x_t)\|^2 \\ & \leq \left( 1 - \frac{\alpha}{\beta} \right) (h(x_t) - h^*). \end{aligned}$$

Telescope over  $t = 0, \dots, T-1$

$$h(x_T) - h^* \leq \left( 1 - \frac{\alpha}{\beta} \right)^T (h(x_0) - h^*).$$

We complete the proof by noting that

$$h(x) - h^* \leq \frac{\beta}{2} \text{dist}^2(x, X^*) \quad \text{and} \quad h(x) - h^* \geq \frac{\alpha}{2} \text{dist}^2(x, X^*). \quad (16)$$

■

Then we can easily show that  $\nabla\varphi_\sigma(x)$  can be efficiently approximated in logarithmic time. Combining both the outer and inner iterations yields the following result.

**Theorem 4.1** *Suppose Assumption 4.1 holds. Define  $\Delta := \varphi(x_0) - \inf_{x \in \mathbb{R}^{d_x}} \varphi(x)$ ,  $R := \text{dist}^2(y_0, Y^*(x))$  and supposed both  $\Delta, R$  are bounded. Set the parameters in Algorithm 2 as*

$$\begin{aligned} \eta &\asymp \ell^{-1} \kappa^{-3}, \quad \sigma \asymp \min \left\{ \frac{R}{\kappa}, \frac{\epsilon}{\ell \kappa^3}, \frac{L_g}{L_f}, \frac{\rho_g}{\rho_f}, \bar{\sigma} \right\}, \\ \tau &= \frac{1}{\sigma L_f + L_g}, \quad K \asymp \frac{L_g}{\mu} \log \left( \frac{L_g}{\mu \sigma} \right), \end{aligned}$$

then it can find an  $\epsilon$ -first-order stationary point of  $\varphi(x)$  within  $T = \mathcal{O}(\ell \kappa^3 \epsilon^{-2})$  iterations, where  $\ell, \kappa$  are defined in Assumption 4.1. The total number of first-order oracle calls is bounded by  $\mathcal{O}(\ell \kappa^4 \epsilon^{-2} \log(\rho_f \ell \kappa / \epsilon))$ .

**Proof** Let  $L$  be the gradient Lipschitz constant of  $\varphi(x)$ . Let  $\eta \leq 1/(2L)$ , then

$$\begin{aligned} \varphi(x_{t+1}) &\leq \varphi(x_t) + \langle \nabla\varphi(x_t), x_{t+1} - x_t \rangle + \frac{L}{2} \|x_{t+1} - x_t\|^2 \\ &= \varphi(x_t) - \frac{\eta}{2} \|\nabla\varphi(x_t)\|^2 - \left( \frac{\eta}{2} - \frac{\eta^2 L}{2} \right) \|\hat{\nabla}\varphi(x_t)\|^2 + \frac{\eta}{2} \|\hat{\nabla}\varphi(x_t) - \nabla\varphi(x_t)\|^2 \\ &\leq \varphi(x_t) - \frac{\eta}{2} \|\nabla\varphi(x_t)\|^2 - \frac{1}{4\eta} \|x_{t+1} - x_t\|^2 + \frac{\eta}{2} \|\hat{\nabla}\varphi(x_t) - \nabla\varphi_\sigma(x_t)\|^2 + \mathcal{O}(\eta\epsilon^2). \end{aligned} \quad (17)$$

Note that

$$\|\hat{\nabla}\varphi(x_t) - \nabla\varphi_\sigma(x_t)\| \leq \frac{2L_g}{\sigma} \text{dist}(y_t^K, Y_\sigma^*(x_t)) + \frac{L_g}{\sigma} \text{dist}(z_t^K, Y^*(x_t)). \quad (18)$$

Then by Lemma G.8, we have

$$\|\hat{\nabla}\varphi(x_t) - \nabla\varphi_\sigma(x_t)\|^2 \leq \frac{8L_g^3}{\mu\sigma^2} \exp\left(-\frac{\mu K}{2L_g}\right) (\text{dist}^2(y_t^K, Y_\sigma^*(x_t)) + \text{dist}^2(z_t^K, Y^*(x_t))). \quad (19)$$

By Young's inequality and Lemma 4.1,

$$\begin{aligned} \text{dist}^2(y_{t+1}^0, Y_\sigma^*(x_{t+1})) &\leq 2\text{dist}^2(y_t^K, Y_\sigma^*(x_t)) + 2\text{dist}^2(Y_\sigma^*(x_{t+1}), Y_\sigma^*(x_t)) \\ &\leq \frac{4L_g}{\mu} \exp\left(-\frac{\mu K}{2L_g}\right) \text{dist}^2(y_t^0, Y_\sigma^*(x_t)) + \frac{8L_g^2}{\mu^2} \|x_{t+1} - x_t\|^2, \end{aligned}$$

Similarly, we can derive the recursion about  $\text{dist}^2(z_t^0, Y^*(x_t))$ .

Put them together and let

$$K \geq \frac{2L_g}{\mu} \log \left( \frac{8L_g}{\mu} \right),$$

we have

$$\delta_{t+1} \leq \frac{1}{2} \delta_t + \frac{16L_g^2}{\mu^2} \|x_{t+1} - x_t\|^2,$$

where we define  $\delta_t := \text{dist}^2(y_t^0, Y_\sigma^*(x_t)) + \text{dist}^2(z_t^0, Y^*(x_t))$ . Telescoping over  $t$  yields

$$\delta_t \leq \underbrace{\left( \frac{1}{2} \right)^t \delta_0 + \frac{16L_g^2}{\mu^2} \sum_{j=0}^{t-1} \left( \frac{1}{2} \right)^{t-1-j} \|x_{j+1} - x_j\|^2}_{:= (*)}.$$

Plug into Equation (19), which, in conjunction with Equation (17), yields that

$$\varphi(x_{t+1}) \leq \varphi(x_t) - \frac{\eta}{2} \|\nabla \varphi(x_t)\|^2 - \frac{1}{4\eta} \|x_{t+1} - x_t\|^2 + 4\eta \times \underbrace{\frac{L_g^3}{\mu\sigma^2} \exp\left(-\frac{\mu K}{2L_g}\right)}_{:=\gamma} \times (*) + \mathcal{O}(\eta\epsilon^2).$$

Telescoping over  $t$  further yields

$$\begin{aligned} \frac{\eta}{2} \sum_{t=0}^{T-1} \|\nabla \varphi(x_t)\|^2 &\leq \varphi(x_0) - \inf_{x \in \mathbb{R}^{d_x}} \varphi(x) + 8\eta\gamma\delta_0 \\ &\quad - \left( \frac{1}{4\eta} - \frac{148\eta\gamma L_g^2}{\mu^2} \right) \sum_{t=0}^{T-1} \|x_{t+1} - x_t\|^2 + \mathcal{O}(\eta\epsilon^2). \end{aligned} \tag{20}$$

Let  $K = \mathcal{O}(\kappa \log(\kappa/\sigma)) = \mathcal{O}(\kappa \log(\ell\kappa/\epsilon))$  such that  $\gamma$  is sufficiently small with

$$\gamma \leq \min \left\{ \frac{\mu^2}{1184\eta^2 L_g^2}, \frac{1}{8\eta} \right\}.$$

Then we have,

$$\frac{1}{T} \sum_{t=0}^{T-1} \|\nabla \varphi(x_t)\|^2 \leq \frac{2}{\eta T} (\Delta + \delta_0) + \mathcal{O}(\epsilon^2).$$

By Lemma 4.1 we know that for any

$$\text{dist}^2(y_0, Y_\sigma^*(x)) \leq 2\text{dist}^2(y_0, Y^*(x)) + 2\text{dist}^2(Y_\sigma^*(x), Y^*(x)) = \mathcal{O}(R).$$

Hence  $\delta_0 = R$  is also bounded. This concludes the proof. ■



## Appendix H. Proof of Theorem 4.2

**Lemma H.1** Suppose  $h(x) : \mathbb{R}^d \rightarrow \mathbb{R}$  is  $\alpha$ -PL and has  $\beta$ -Lipschitz gradients. Consider the following update of stochastic gradient descent:

$$x_{t+1} = x_t - \frac{1}{\beta} \nabla h(x_t; \mathcal{B}_t),$$

where the mini-batch gradient satisfies

$$\mathbb{E}_{\mathcal{B}_t} [\nabla h(x_t; \mathcal{B}_t)] = \nabla h(x_t), \quad \mathbb{E}_{\mathcal{B}_t} \|\nabla h(x_t; \mathcal{B}_t) - \nabla h(x_t)\|^2 \leq \frac{M^2}{B}.$$

Let  $X^* = \arg \min_{x \in \mathbb{R}^d} h(x)$  and  $h^* = \min_{x \in \mathbb{R}^d} h(x)$ . Then it holds that

$$\mathbb{E} [\text{dist}^2(x_T, X^*)] \leq \left(1 - \frac{\alpha}{\beta}\right)^T \frac{\beta}{\alpha} \text{dist}^2(x_0, X^*) + \frac{M^2}{\alpha\beta B}.$$

**Proof** The proof is similar to the deterministic case. Conditional on  $x_t$ , we have

$$\begin{aligned} \mathbb{E} [h(x_{t+1}) - h^*] &\leq \mathbb{E} \left[ h(x_t) - h^* + \nabla h(x_t)^\top (x_{t+1} - x_t) + \frac{\beta}{2} \|x_{t+1} - x_t\|^2 \right] \\ &= \mathbb{E} \left[ h(x_t) - h^* - \frac{1}{2\beta} \|\nabla h(x_t)\|^2 + \frac{1}{2\beta} \|\nabla h(x_t; \mathcal{B}_t) - \nabla h(x_t)\|^2 \right] \\ &\leq \left(1 - \frac{\alpha}{\beta}\right) (h(x_t) - h^*) + \frac{M^2}{2\beta B}. \end{aligned}$$

Telescope,

$$\mathbb{E} [h(x_{t+1}) - h^*] \leq \left(1 - \frac{\alpha}{\beta}\right)^T (h(x_0) - h^*) + \frac{M^2}{2\beta B}.$$

We conclude the proof by using Equation (16). ■

**Theorem 4.2** Suppose Assumption 4.1 and 4.2 hold. Define  $\Delta := \varphi(x_0) - \inf_{x \in \mathbb{R}^{d_x}} \varphi(x)$ ,  $R := \text{dist}^2(y_0, Y^*(x))$  and supposed both  $\Delta, R$  are bounded. Set the parameters in Algorithm 3 as

$$\begin{aligned} \eta &\asymp \ell^{-1} \kappa^{-3}, \quad \sigma \asymp \min \left\{ \frac{R}{\kappa}, \frac{\epsilon}{\ell \kappa^3}, \frac{L_g}{L_f}, \frac{\rho_g}{\rho_f}, \bar{\sigma} \right\}, \\ \tau &= \frac{1}{\sigma L_f + L_g}, \quad B \asymp \frac{L_g (\sigma^2 M_f^2 + M_g^2)}{\mu \sigma^2 \epsilon^2}, \quad K_t \asymp \frac{L_g}{\mu} \log \left( \frac{L_g^3 \delta_t}{\mu \sigma^2 \epsilon^2} \right), \end{aligned} \tag{8}$$

where  $\delta_t$  is defined via the recursion

$$\delta_{t+1} \leq \frac{1}{2} \delta_t + \frac{8L_g^2}{\mu^2} \|x_{t+1} - x_t\|^2 + \mathcal{O} \left( \frac{\sigma^2 \epsilon^2}{L_g^2} \right), \quad \delta_0 \asymp R. \tag{9}$$

Then Algorithm 3 can find an  $\epsilon$ -first-order stationary point of  $\varphi(x)$  in expectation within  $T = \mathcal{O}(\ell\kappa^3\epsilon^{-2})$  iterations, where  $\ell, \kappa$  are defined in Definition 4.1. The total number of stochastic first-order oracle calls is bounded by

$$\begin{cases} \mathcal{O}(\ell\kappa^4\epsilon^{-2}\log(\rho_f\ell\kappa/\epsilon)), & M_f = 0, M_g = 0; \\ \mathcal{O}(\ell\kappa^5\epsilon^{-4}\log(\rho_f\ell\kappa/\epsilon)), & M_f > 0, M_g = 0; \\ \mathcal{O}(\ell^3\kappa^{11}\epsilon^{-6}\log(\rho_f\ell\kappa/\epsilon)), & M_f > 0, M_g > 0. \end{cases}$$

**Proof** Define  $\delta_t := \mathbb{E} [\text{dist}^2(y_t^0, Y_\sigma^*(x_t)) + \text{dist}^2(z_t^0, Y^*(x_t))]$ . By Equation (18), letting

$$\text{dist}^2(y_t^K, Y_\sigma^*(x_t)) + \text{dist}^2(z_t^K, Y^*(x_t)) \leq \mathcal{O}\left(\frac{\sigma^2\epsilon^2}{L_g^2}\right)$$

ensures  $\mathbb{E}\|\hat{\nabla}\varphi(x_t) - \nabla\varphi_\sigma(x_t)\|^2 \leq \mathcal{O}(\epsilon)$ . Then telescoping over Equation (17) shows that one can find an  $\epsilon$ -stationary point of  $\varphi(x)$  with  $T = \mathcal{O}(\epsilon^{-2})$  outer-loop iterations. By Lemma H.1, it suffices to set the parameters in the inner loop as Equation (8). But  $K_t$  requires the knowledge of  $\delta_t$ . Next, we bound  $\delta_t$  via a recursion which allows us to get rid of the prior knowledge of  $\delta_t$ . By Lemma H.1 and Lemma 4.1,

$$\begin{aligned} & \text{dist}^2(y_{t+1}^0, Y_\sigma^*(x_{t+1})) \\ & \leq 2\text{dist}^2(y_t^K, Y_\sigma^*(x_t)) + 2\text{dist}^2(Y_\sigma^*(x_{t+1}), Y_\sigma^*(x_t)) \\ & \leq \frac{4L_g}{\mu} \exp\left(-\frac{\mu K_t}{2L_g}\right) \text{dist}^2(y_t^0, Y_\sigma^*(x_t)) + \frac{8L_g^2}{\mu^2} \|x_{t+1} - x_t\|^2 + \mathcal{O}\left(\frac{\sigma^2\epsilon^2}{L_g^2}\right). \end{aligned}$$

Similarly, we can derive the recursion about  $\text{dist}^2(z_t^0, Y^*(x_t))$ . Put them together and let

$$K_t \geq \frac{2L_g}{\mu} \log\left(\frac{8L_g}{\mu}\right),$$

we can get Equation (9). Telescoping over  $t$ ,

$$\sum_{t=0}^{T-1} \delta_t \leq 2\delta_0 + \frac{32L_g^2}{\mu^2} \sum_{t=0}^{T-1} \|x_{t+1} - x_t\|^2 + \frac{\sigma^2\epsilon^2}{2L_g^2},$$

where on the right-hand side we know  $\sum_{t=0}^{T-1} \|x_{t+1} - x_t\|$  must also be bounded by Equation (17).

With this recursion, the total iterations can be bounded above by

$$\sum_{t=0}^{T-1} K_t \leq \sum_{t=0}^{T-1} \frac{2L_g}{\mu} \log\left(\frac{32L_g^3\delta_t}{\mu\sigma^2\epsilon^2}\right) \leq \frac{2L_g T}{\mu} \log\left(\frac{32L_g^3 \sum_{t=0}^{T-1} \delta_t}{\mu\sigma^2\epsilon^2 T}\right).$$

And the total number of stochastic oracle calls is  $B$  times the above bound for iterations. ■