
BiSLS/SPS: Auto-tune Step Sizes for Stable Bi-level Optimization

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 The popularity of bi-level optimization (BO) in deep learning has spurred a growing
2 interest in studying gradient-based BO algorithms. However, existing algorithms
3 involve two coupled learning rates that can be affected by approximation errors
4 when computing hypergradients, making careful fine-tuning necessary to ensure
5 fast convergence. To alleviate this issue, we investigate the use of recently proposed
6 adaptive step-size methods, namely stochastic line search (SLS) and stochastic
7 Polyak step size (SPS), for computing both the upper and lower-level learning rates.
8 First, we revisit the use of SLS and SPS in single-level optimization without the
9 additional interpolation condition that is typically assumed in prior works. For such
10 settings, we investigate new variants of SLS and SPS that improve upon existing
11 suggestions in the literature and are simpler to implement. Importantly, these two
12 variants can be seen as special instances of general family of methods with an
13 envelope-type step-size. This unified envelope strategy allows for the extension
14 of the algorithms and their convergence guarantees to BO settings. Finally, our
15 extensive experiments demonstrate that the new algorithms, which are available in
16 both SGD and Adam versions, can find large learning rates with minimal tuning
17 and converge faster than corresponding vanilla SGD or Adam BO algorithms that
18 require fine-tuning.

19 1 Introduction

20 Bi-level optimization has found its applications in various fields of machine learning, such as
21 hyperparameter optimization [14, 17, 30, 40], adversarial training [51], data distillation [2, 53],
22 neural architecture search [28, 39], neural-network pruning [52], and meta-learning [13, 37, 11].
23 Specifically, it is used widely for problems that exhibit a hierarchical structure of the following form:

$$\min_{x \in X} F(x) = \mathbb{E}_{\phi}[f(x, y^*(x); \phi)] \quad \text{s.t.} \quad y^*(x) = \operatorname{argmin}_{y \in Y} \mathbb{E}_{\psi}[g(x, y; \psi)]. \quad (1)$$

24 Here, the solution to the lower-level objective g becomes the input to the upper-level objective f , and
25 in (1) the upper-level variable x is fixed when optimizing the lower-level variable y . To solve such
26 bi-level problems using gradient-based methods requires computing the hypergradient of F , which
27 based on the chain rule is given as [15]:

$$\nabla F(x) = \nabla_x f(x, y^*(x)) + \nabla_{xy}^2 g(x, y^*(x)) [\nabla_{yy}^2 g(x, y^*(x))]^{-1} \nabla_y f(x, y^*(x)). \quad (2)$$

28 In practice, the closed-form solution $y^*(x)$ is difficult to obtain, and one strategy is to run a few steps
29 of (stochastic) gradient descent on g w.r.t. y to get an approximation \bar{y} , and use \bar{y} in places of $y^*(x)$.
30 We denote the stochastic hypergradient based on \bar{y} as $h_f(x, \bar{y})$ and the stochastic gradient of g w.r.t. y
31 as h_g . This leads to a general gradient-based framework for solving bi-level optimization [15, 19, 4].

At each iteration k , run T (can be one or more) steps of SGD on y , i.e. $y^{k,t+1} = y^{k,t} - \beta h_g^{k,t}$, then run one step on x using the approximated hypergradient:

$$x^{k+1} = x^k - \alpha h_f(x^k, y^{k+1}), \quad \text{where } y^{k+1} = y^{k,T}. \quad (3)$$

Based on this framework, a series of stochastic algorithms have been developed to achieve the optimal or near-optimal rate of their deterministic counterparts [7, 8]. These algorithms can be broadly divided into single-loop ($T = 1$) or double-loop ($T > 1$) categories [23].

Unlike minimizing the single-level finite-sum (convex) problem

$$F(x) := \min_{x \in \mathcal{C}} \frac{1}{N} \sum_{i=1}^N f_i(x), \quad (4)$$

where only one learning rate is involved when using SGD, bi-level optimization involves tuning both the lower and upper-level learning rates (β and α respectively). This poses a significant challenge due to the potential correlation between these learning rates [19]. Thus, as observed in Figure 1, algorithm divergence can occur when either α or β is large. While there is considerable literature on achieving faster rates in bi-level optimization [24, 5, 7, 8], only a few studies have focused on stabilizing its training and automating the tuning of α and β . This work addresses the question: **Is it possible to utilize large α and β without manual tuning?** In doing so, we explore the use of stochastic adaptive-step size methods, namely stochastic Polyak step size (SPS) and stochastic line search (SLS), which utilize gradient information to adjust the learning rate at each iteration [44, 29]. These methods have been demonstrated to perform well in interpolation settings with strong convergence guarantees [44, 29]. However, applying them to bi-level optimization (BO) introduces significant challenges, as follows. ① BO requires tuning two correlated learning rates (for lower and upper-level). ② The bias in the stochastic approximation of the hypergradient complicates the practical performance and convergence analysis of SLS and SPS. ③ Other algorithmic challenges arise for both algorithms: For SLS, verifying the stochastic Armijo condition at the upper-level involves evaluating the objective at a new $(x, y^*(x))$ pair, while $y^*(x)$ is only approximately known; For SPS, most existing variants guarantee good performance only in interpolating settings, which are typically not satisfied for the upper-level objective in BO [22]. Before presenting our solutions to the challenges above in Sec 2, we first review the most closely related literature.




Figure 1: Results based on hyper-representation learning task (see Sec 4 for details). Validation loss against upper-level iterations for different values of β (left, $\alpha = 0.005$) and α (right, $\beta = 0.01$). Unless carefully tuned, vanilla SGD-based methods for BO are very unstable.

1.1 Related Work

Gradient-Based Bi-level Optimization Penalty or gradient-based approaches have been used for solving bi-level optimization problems [10, 45, 21]. Here we focus our discussions on stochastic gradient-based methods as they are closely related to this work. For double-loop algorithms, an early work (BSA) by Ghadimi and Wang [15] has derived the sample complexity of ϕ in achieving an ϵ -stationary point to be $\mathcal{O}(\epsilon^{-2})$, but require the number of lower-level steps to satisfy $T \sim \mathcal{O}(\epsilon^{-1})$. Using a warm start strategy (stocBiO), Ji et al. [22] removed this requirement on T . However, to achieve the same sample complexity, the batch size of stocBiO grows as $\mathcal{O}(\epsilon^{-1})$. Chen et al. [4] removed both requirements on T and batch size by using the smoothness properties of $y^*(x)$ and setting the step sizes α and β at the same scale. For single-loop algorithms, a pioneering work by Hong et al. [19] gave a sample complexity of $\mathcal{O}(\epsilon^{-2.5})$, provided α and β are on two different scales (TTSA). By making corrections to the y variable update (STABLE), Chen et al. [5] improved the rate to $\mathcal{O}(\epsilon^{-2})$. However, extra matrix projections required by STABLE can incur high computation cost [5, 4]. By incorporating momentum into the updates of x and y (SUSTAIN), Khanduri et al. [24] further improved the rate to $\mathcal{O}(\epsilon^{-1.5})$ [6]. Besides these single or double-loop algorithms, a series of works have drawn ideas from variance reduction to achieve faster convergence rates for BO. For example, Yang et al. [49] designed the VRBO algorithm based on SPIDER [12]. Dagr  ou et al. [7, 8] designed the SABA and SRBA algorithms based on SAGA and SARAH respectively, and demonstrate that they can achieve the optimal rate of $\mathcal{O}(\epsilon^{-1})$ [9, 35]. Huang et al. [20] proposes to use Adam-type step sizes in BO. However, it introduces three sequences of learning rates $(\alpha_k, \beta_k, \eta_k)$




Figure 2: Experiments on quadratic functions adapted from [29]. The objective is the sum of two-dimensional functions $f_i = \frac{1}{2}(x - x_i^*)^T H_i (x - x_i^*)$, where H_i is positive definite and $i = 1, 2$ (see Appendix B for more details). From left to right, we show: the objective value, distance to optimum, step size, and iterate trajectories.

that require tuning, which limits its practical usage. To our knowledge, none of these works have explicitly addressed the fundamental problem of how to select α and β in bi-level optimization. In this work, we focus on the alternating SGD framework (T can be 1 or larger), and design efficient algorithms that find large α and β without tuning, while ensuring the stability of training.

Adaptive Step Size Adaptive step-size such as Adam has found great success in modern machine learning, and different variants have been proposed [25, 38, 47, 31, 32]. Here, we limit our discussions on two adaptive step sizes that are most relevant to this work. The Armijo line search is a classic way for finding step sizes for gradient descent [48]. Vaswani et al. [44] extends it to the stochastic setting (SLS) and demonstrates that the algorithm works well with minimal tuning required under interpolation, where model fits the data perfectly. Hence, the method is adaptive to local smoothness of the objective, which is typically difficult to predict a priori. However, the theoretical guarantee of SLS in the non-interpolating regime is lacking. In fact, the results in Figure 3 suggest that SLS can perform poorly for convex losses when interpolation is not satisfied. Besides SLS, another adaptive method derived from Polyak step size is proposed by Loizou et al. [29] with the name stochastic Polyak step size (SPS). Loizou et al. [29] further places an upper bound on the step size resulting in the SPS_{\max} variant. Similar to SLS, the algorithm performs well when the model is over-parametrized. Without interpolation, the algorithm converges to a neighborhood of the solution whose size depends on this upper bound.

In a later work, Orvieto et al. [36] make the SPS converge to the exact solution by ensuring the step size and its upper bound are both non-increasing (DecSPS). However, enforcing monotonicity may result in the step size being smaller than decaying-step SGD and losing the adaptive features of SPS (see Figure 2, 3). In this work, we propose new versions of SLS and SPS that do not require monotonicity and extend them into the alternating SGD bi-level optimization framework (3).

2 Summary of Contributions

We discuss our main contributions in this section, which is organized as follows. First, we discuss our variants of SPS and SLS, and unify them under the umbrella of “envelope-type step-size”. Then, we extend the envelope-type step size to the bi-level setting. Finally, we discuss our bi-level line-search algorithms based on Adam and SGD.

Converging SPSB and SLSB by Envelope Approach

We first propose simple variants of SLS and SPS that converge in the non-interpolating setting while not requiring the step size to be monotonic. To this end, we introduce a new stochastic

Algorithm 1 BiSLS-Adam/SGD

Input: $x^0, y^0, K, T, \delta, \alpha_{b,0}, \beta_{b,0}, w, \eta$

Output: x

```

1: for  $k = 0, 1, \dots, K - 1$  do
2:    $y^{k,0} = y^k$ 
3:   for  $t = 0, 1, \dots, T - 1$  do
4:      $\beta_{b,k}^t \leftarrow \text{reset}(\beta, \beta_{b,0}, \eta, \text{opt})$   $\triangleright$  see Algorithm 2
5:      $\beta \leftarrow \text{linesearch based on (8) starting from } \beta_{b,k}^t$ 
6:      $y^{k,t+1} = y^{k,t} - \beta h_g^{k,t}$ 
7:   end for
8:    $y^{k+1} = y^{k,T-1}; \hat{x}^k = x^k; \hat{y}^{k+1} = y^{k+1}$ 
9:    $\alpha \leftarrow \text{reset}(\alpha, \alpha_{b,0}, \eta, \text{opt})$ 
10:  while (14) based on  $(\hat{x}^k, \hat{y}^{k+1}, \alpha, \delta)$  does not hold. do
11:     $\alpha = \alpha * w$ 
12:     $\hat{x}^k = x^k - \alpha h_f(x^k, y^{k+1})$  or
13:     $\hat{x}^k = x^k - \alpha A_k^{-1} h_f(x^k, y^{k+1})$ 
14:     $\hat{y}^{k+1} = y^{k+1} - \beta h_g(\hat{x}^k, y^{k+1})$ 
15:  end while
16:   $x^{k+1} = x^k - \alpha h_f^k$ 
17: end for

```

123 Polyak step size (SPSB). For comparison, we also recall the step-sizes of SPS_{\max} and DecSPS .
 124 For all methods, the iterate updates are given as $x_{k+1} = x_k - \gamma_k \nabla f_{i_k}(x^k)$ where i_k is sampled
 125 uniformly from $[n] = \{1, \dots, n\}$ at each iteration k . The step-sizes γ_k are then defined as follows:

$$\text{SPS}_{\max} [29]: \quad \gamma_k = \min\left\{\frac{f_{i_k}(x^k) - f_{i_k}^*}{c \|\nabla f_{i_k}(x^k)\|^2}, \gamma_{b,0}\right\} \quad (5)$$

$$\text{DecSPS} [36]: \quad \gamma_0 = \bar{\gamma} \quad \gamma_k = \frac{1}{c_k} \min\left\{\frac{f_{i_k}(x^k) - l_{i_k}^*}{\|\nabla f_{i_k}(x^k)\|^2}, c_{k-1} \gamma_{k-1}\right\} \quad \forall k \geq 1 \quad (6)$$

$$\text{SPSB (ours)}: \quad \gamma_k = \min\left\{\frac{f_{i_k}(x^k) - l_{i_k}^*}{c_k \|\nabla f_{i_k}(x^k)\|^2}, \gamma_{b,k}\right\}, \quad (7)$$

126 where $f_i^* = \inf_x f_i(x)$, $\bar{\gamma} = \frac{1}{c_0} \min\left\{\frac{f_{i_0}(x^0) - l_{i_0}^*}{\|\nabla f_{i_0}(x^0)\|^2}, c_0 \gamma_{b,0}\right\}$, c_k is non-decreasing, $\gamma_{b,k}$ is non-
 127 increasing, and $l_i^* \leq f_i^*$ is any lower bound.

128 Unlike SPS_{\max} in which $\gamma_{b,0}$ is a
 129 constant, our upper bound $\gamma_{b,k}$ is
 130 non-increasing. Also, unlike Dec-
 131 SPS in which both the step size and
 132 the upper bound are non-increasing
 133 (this is because $\gamma_k \leq \frac{c_k-1}{c_k} \gamma_{k-1}$ and
 134 $\min\left\{\frac{1}{2cL_{\max}}, \frac{c_0 \gamma_{b,0}}{c_k}\right\} \leq \gamma_k \leq \frac{c_0 \gamma_{b,0}}{c_k}$
 135 [36, Lemma 1]), we simplify the re-
 136 cursive structure and do not require
 137 the step-size to be monotonic. As
 138 we empirically observe in Figure 3,
 139 the step size of DecSPS is similar to
 140 that of decaying SGD and in fact can
 141 be much smaller. Interestingly, the
 142 resulting performance of DecSPS is
 143 worse than SPS_{\max} despite SPS_{\max} eventually becoming unstable once iterates get closer to the
 144 neighborhood of a solution and the step-size naturally behaves erratically. This is not unexpected due
 145 to small gradient norms (note division by gradient-norm in (5)) and dissimilarity between samples in
 146 the non-interpolating scenario. Moreover, note that the adaptivity of SPS in the early stage seems
 147 to be lost in DecSPS due to monotonicity of the latter. On the other hand, SPSB not only takes
 148 advantage of the large SPS steps that leads to fast convergence, but also stays regularized due to the
 149 non-increasing upper bound $\gamma_{b,k}$ in (19). These observations are further supported by the experiments
 150 on quadratic functions given in Figure 2, where we observe the fast convergence of SPSB and the
 151 instability of SPS_{\max} . Motivated by the good practical performance of SPSB, we take a similar
 152 approach for SLS. The SLS proposed and analyzed by Vaswani et al. [44] starts with $\gamma_{b,0}$ and in each
 153 iteration k finds the largest $\gamma_k \leq \gamma_{b,0}$ that satisfies:

$$f_{i_k}(x_k - \gamma_k \nabla f_{i_k}(x_k)) \leq f_{i_k}(x_k) - \bar{c} \cdot \gamma_k \|\nabla f_{i_k}(x_k)\|^2, \quad 0 < \bar{c} < 1. \quad (8)$$

154 To ensure its convergence without interpolation, we replace $\gamma_{b,0}$ with appropriate non-increasing
 155 sequence $\gamma_{b,k}$. We name this variant of SLS as SLSB. Interestingly, the empirical performance and
 156 step size of SLSB are similar to those of SPSB (see Figure 3). This can be explained by observing
 157 that the step sizes of SPSB and SLSB share similar envelope structures, as follows (see Lemma 1 in
 158 Appendix A):

$$\begin{aligned} \text{SPSB:} \quad & \min\left\{\frac{1}{2cL_{\max}}, \gamma_{b,k}\right\} \leq \gamma_k = \min\left\{\frac{f_{i_k}(x^k) - l_{i_k}^*}{c \|\nabla f_{i_k}(x^k)\|^2}, \gamma_{b,k}\right\}, \quad 0 < c, \\ \text{SLSB:} \quad & \min\left\{\frac{2(1-\bar{c})}{L_{\max}}, \gamma_{b,k}\right\} \leq \gamma_k \leq \min\left\{\frac{f_{i_k}(x^k) - l_{i_k}^*}{\bar{c} \|\nabla f_{i_k}(x^k)\|^2}, \gamma_{b,k}\right\}, \quad 0 < \bar{c} < 1. \end{aligned}$$

159 Therefore, we unify their analysis based on the following generic *envelope-type step size*:

$$\gamma_k = \min\{\max\{\gamma_{l,k}, \tilde{\gamma}_k\}, \gamma_{b,k}\}, \quad \gamma_{l,k} = \min\{\omega, \gamma_{b,k}\}, \quad (9)$$

160 where $\omega > 0$, $\gamma_{b,k}$ is non-increasing, and $\tilde{\gamma}_k$ satisfies $\gamma_{l,k} := \min\{\omega, \gamma_{b,k}\} \leq \tilde{\gamma}_k \leq \gamma_{b,k}$. We show
 161 that this envelope-type step size converges at a rate $\mathcal{O}(\frac{1}{\sqrt{K}})$ and $\mathcal{O}(\frac{1}{K})$ for convex and strongly-convex
 162 losses respectively.




Figure 3: Binary linear classification on w8a dataset using logistic loss [3]. We choose $\gamma_{b,0} = 1000$ for all algorithms; $c = 1$ and $\bar{c} = 1$ for SPS_{\max} and SLS respectively; $c_k = \sqrt{k+1}$ for DecSPS; $c_k = 1$ and $\gamma_{b,k} = \frac{\gamma_{b,0}}{\sqrt{k+1}}$ for SPSB; $\bar{c} = 0.1$ and $\gamma_{b,k} = \frac{\gamma_{b,0}}{\sqrt{k+1}}$ for SLSB; $\gamma_{b,k} = \frac{\gamma_{b,0}}{\sqrt{k+1}}$ for decaying-step SGD.

163 **Envelope Step Size for Bi-level Optimization (BiSPS)** We extend the analysis of envelope-type
 164 step size to the bi-level setting. The step sizes for upper and lower-level objectives of our general
 165 envelope-type method are:

$$\text{Upper: } \alpha_k = \min\{\max\{\alpha_{l,k}, \tilde{\alpha}_k\}, \alpha_{b,k}\} \quad \text{hence} \quad \alpha_{l,k} \leq \tilde{\alpha}_k \leq \alpha_{b,k} \quad (10)$$

$$\text{Lower: } \beta_{k,t} = \min\left\{\frac{g(x^k, y^{k,t}; \psi) - g(x^k, y_{x^k, \psi}^*; \psi)}{p \|\nabla_y g(x^k, y^{k,t}; \psi)\|^2}, \beta_{b,k}\right\} \quad \forall t, \quad (11)$$

166 where $y_{x^k, \psi}^*$ is the minimizer of the function $g(x^k, \cdot; \psi)$, and $\alpha_{l,k}$, $\alpha_{b,k}$, and $\beta_{b,k}$ are three non-
 167 increasing sequences. Note that $\beta_{b,k}$ is fixed over the lower-level iterations for a given k , therefore,
 168 this is equivalent to running T steps of SPS_{\max} to minimize the function g at each upper iteration k .
 169 However, the decrease in the upper bound $\beta_{b,k}$ with k is crucial to guarantee the overall convergence
 170 of the algorithm (see Theorem 3). Starting from the general step-size rules in (10), (11), our bi-level
 171 extension of SPS, which we call BiSPS, follow by setting α_k in the form of SPS computed using
 172 stochastic hypergradient h_f^k . That is,

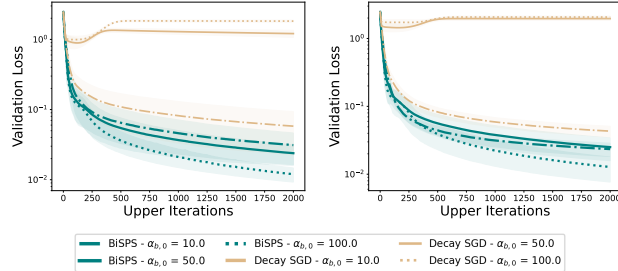
$$\tilde{\alpha}_k = \frac{f(x^k, y^{k+1}; \phi) - l_{f(\cdot, y^{k+1}; \phi)}^*}{p \|h_f^k\|^2}, \quad \alpha_{l,k} = \frac{\alpha_{l,0}}{\sqrt{k+1}}, \quad \alpha_{b,k} = \frac{\alpha_{b,0}}{\sqrt{k+1}}, \quad (12)$$

173 where $\alpha_{l,0} \leq \alpha_{b,0}$ and $l_{f(\cdot, y^{k+1}; \phi)}^*$ is a lower bound for $\inf_x f(x, y^{k+1}; \phi)$. For computing h_f^k , we
 174 can take a similar approach as previous works [15, 19, 4] that use Neumann series setting

$$h_f^k = \nabla_x f(x^k, y^{k+1}; \phi) - \nabla_{xy} g(x^k, y^{k+1}; \psi_0) \left[\frac{N}{L_g} \prod_{j=1}^{\bar{N}} (I - \nabla_{yy}^2 g(x^k, y^{k+1}; \psi_j)) \right] \nabla_y f(x^k, y^{k+1}; \phi), \quad (13)$$

175 where \bar{N} is sampled uniformly from $[N]$ and N is the total number of samples. For BiSPS, we use
 176 the same sample for $f(x^k, y^{k+1}; \phi)$ and $\nabla f(x^k, y^{k+1}; \phi)$ when evaluating $\tilde{\alpha}_k$ in (12). Interestingly,
 177 we also empirically observe that using independent samples for computing $\tilde{\alpha}_k$ and h_f^k resulting
 178 in similar performance as using the same sample. The optimal rate of SGD for non-convex bi-
 179 level optimization is $\mathcal{O}(\frac{1}{\sqrt{K}})$ without a growing batch size [4]. We show that BiSPS can obtain
 180 the same rate (see Theorem 3) by taking the envelope-type step-size of the form (10) and (11).

181 We implement BiSPS according to
 182 (12) and observe that it has better per-
 183 formances over decaying-step SGD
 184 with less variations across different
 185 values of $\alpha_{b,0}$ (see Figure 4 and note
 186 that decaying-step SGD is of the form
 187 $\frac{\alpha_{b,0}}{\sqrt{k+1}}$).



188 Stochastic Line-Search Algorithms 189 for Bi-level Optimization

190 The challenge of extending SLS to bi-
 191 level optimization is rooted in the
 192 term $y^*(x)$. In fact, we realize that
 193 some of the bi-level objectives are of
 194 the form $F(x) = f(y^*(x))$. That
 195 is, f does not have an explicit de-
 196 pendence on x , e.g. the data hyper-
 197 cleaning task [22]. This implies that when SLS takes a potential step on x , the approximation of $y^*(x)$
 198 (i.e. $\bar{y}(x)$) also needs to be updated, otherwise there is no change in function values. Moreover, the
 199 use of approximation $\bar{y}(x)$ and the stochastic estimation error in hypergradient would not guarantee a
 200 step size can be always found. To this end, we modify the Armijo line-search rule to be:

$$\begin{aligned} \text{BiSLS-SGD:} \quad & f(x^k - \alpha_k h_f^k, \hat{y}^{k+1}(x^k - \alpha_k h_f^k)) \leq f(x^k, y^{k+1}) - p \alpha_k \|h_f^k\|^2 + \delta, \\ \text{BiSLS-Adam:} \quad & f(x^k - \alpha_k A_k^{-1} h_f^k, \hat{y}^{k+1}(x^k - \alpha_k A_k^{-1} h_f^k)) \leq f(x^k, y^{k+1}) - p \alpha_k \|h_f^k\|_{A_k^{-1}}^2 + \delta, \end{aligned} \quad (14)$$

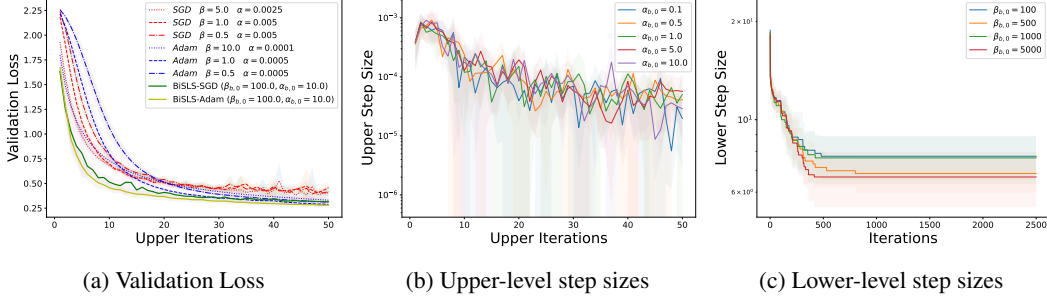


Figure 5: Results on hyper-representation learning task (see Sec 4 for details). (a) Validation loss against upper-level iterations for comparing BiSLS-Adam/SGD to fine-tuned Adam/SGD. (b)(c) Upper (left) and lower-level (right) learning rates found by BiSLS-Adam. For the tuned Adam, the optimal lower and upper-level learning rates are $\mathcal{O}(1)$ and $\mathcal{O}(10^{-4})$, respectively. BiSLS-Adam outperforms tuned Adam/SGD with a starting point that is 5 orders of magnitude larger than the optimal step size.

where $p, \delta > 0$ and A_k is a positive definite matrix such that $A_k^2 = G_k$. Similar to the single-level Adam case, the matrix G_k in the bi-level setting is defined as $G_k = (\beta_2 G_{k-1} + (1 - \beta_2) \text{diag}(h_f^k h_f^{kT})) / (1 - \beta_2^2)$ [25, 43]. Moreover, BiSLS-Adam takes the following steps for updating the variable x : $x^{k+1} = x^k - \alpha_k A_k^{-1} m_k$ where $m^{k+1} = \beta_1 m^k - (1 - \beta_1) h_f^k$. The details are given in Algorithms 1 and 2. We denote the search starting point for the upper-level as $\alpha_{b,k}$ at iteration k , and denote it as $\beta_{b,k}^t$ at step t within iteration k for the lower-level. We remark the following key benefits of resetting $\alpha_{b,k}$ and $\beta_{b,k}^t$ (by using Algorithm 2) to larger values with reference to α_k and β_k^t (respectively) at each step: (1) Avoid always searching from $\alpha_{b,0}$ or $\beta_{b,0}^0$, thus, reducing computation cost, and, (2) preserving an overall non-increasing (not necessarily monotonic) trend for $\alpha_{b,k}$ and $\beta_{b,k}^t$, thus, improving training stability. We found different values of η all work well (see Appendix B). The key algorithmic challenge we are facing is that during the backtracking process, for any candidate α_k , we need to compute $\hat{x}^k := x^k - \alpha_k h_f^k$ and approximate $y^*(\hat{x}^k)$ with \hat{y}^{k+1} (see Algorithm 1). To limit the cost induced by this nested loop, we limit the number of steps to obtain \hat{y}^{k+1} to be 1. Moreover, δ in (14) plays the role of a safeguard that ensures a step size can be found.

We set it to be small to avoid finding unrealistically large learning rates while tolerating some error in the hypergradient estimation (see Appendix B for experiments on the sensitivity of δ). In practice, we empirically find that simply setting $\delta = 0$ works well. In Figure 5a, we observe that BiSLS-Adam outperforms fine-tuned Adam or SGD. Surprisingly, its training is stable even when the search starting point $\alpha_{b,0}$ is 5 orders of magnitude larger than a fine-tuned learning rate ($\mathcal{O}(10^{-4})$). Importantly, BiSLS-Adam finds large upper and lower-level learning rates in early phase (see Figure 5b, 5c) for different values of $\alpha_{b,0}$ and $\beta_{b,0}$ that span 3 orders of magnitudes. Interestingly, the learning rates naturally decay with training (also see Figure 6c and 6d). In essence, BiSLS is a **truly adaptive (no knowledge of initialization required) and robust (different initialization works) method that finds large α and β without tuning**. In the next section, we give the convergence results of the envelope-type step size.

Algorithm 2 reset

Input: $p, q, \eta \geq 1$,
 opt
Output: p
1: **if** $\text{opt} = 1$ **then**
2: $p \leftarrow q$
3: **else if** $\text{opt} = 2$ **then**
4: $p \leftarrow p$
5: **else if** $\text{opt} = 3$ **then**
6: $p \leftarrow \eta \cdot p$
7: **end if**

3 Convergence Results

3.1 Envelope-type step size for single-level optimization

We first state the assumptions, which are standard in the literature, that will be used for analyzing single-level problems. Assumption 1 is on the Lipschitz continuity of f and f_i in Problem 4.

Assumption 1. The individual function f_i is convex and L_i -smooth such that $\|\nabla f_i(x) - \nabla f_i(x')\| \leq L_i \|x - x'\|, \forall i, \forall x \in \text{dom } f$ and the overall function f is L -smooth. We denote $L_{\max} \triangleq \max_i L_i$.

Furthermore, we assume there exists l_i^* such that $l_i^* \leq f_i^* := \inf_x f_i(x)$, $\forall i$, and f is lower bounded by f^* obtained by some x^* such that $f^* = f(x^*)$.

The following bounded gradient assumption is also used in the analysis of convex problems [41, 33].
Assumption 2. There exists $G > 0$ such that $\|\nabla f_i(x)\|^2 \leq G$, $\forall i$.

We first state the theorem for the envelop-type step size defined in (9) for convex functions.

Theorem 1. Suppose Assumption 1, 2 hold, each f_i is convex, $C = \text{dom } f$, γ_k is independent of the sample $\nabla f_k(x^k)$, and choose $\gamma_{b,k} = \frac{\gamma_{b,0}}{\sqrt{k+1}}$. Then, the envelope-type step size in (9) achieves the following rate,

$$\mathbb{E}[f(\bar{x}^K) - f(x^*)] \leq \frac{\|x^0 - x^*\|^2}{2\gamma_{l,K-1}K} + \frac{\gamma_{b,0}^2 G^2 \log(K)}{2\gamma_{l,K-1}K},$$

where $\gamma_{l,K-1} = \min\{\omega, \frac{\gamma_{b,0}}{\sqrt{K}}\}$ and $\bar{x}^K = \frac{1}{K} \sum_{k=0}^K x^k$.

We were not able to give a convergence result that uses the same sample for computing the step size and the gradient. However, we empirically observe that the performance is very similar when using either one or two independent samples per iteration (see Figure 2 and Appendix B). When two independent samples i_k and j_k are used per iteration, the first computes the gradient sample $\nabla f_{i_k}(x^k)$, and the other computes the step-size γ_k . For example, for SPSB this gives $\gamma_k = \min\{\frac{f_{j_k}(x^k) - l_{j_k}^*}{c_k \|\nabla f_{j_k}(x^k)\|^2}, \gamma_{b,k}\}$. This type of assumption has been used in several other works for analyzing adaptive step sizes [27, 44, 29]. Under this assumption, we specialize the results of Theorem 1 to SPSB and SLSB, where $\gamma_{l,K-1} = \min\{\frac{1}{2cL_{\max}}, \frac{\gamma_{b,0}}{\sqrt{K}}\}$ and $\gamma_{l,K-1} = \min\{\frac{2(1-\bar{c})}{L_{\max}}, \frac{\gamma_{b,0}}{\sqrt{K}}\}$ respectively. Concretely, for $K \geq \gamma_{b,0}^2 L_{\max}^2$, SLSB and SPSB with $\gamma_{b,k} = \frac{\gamma_{b,0}}{\sqrt{k+1}}$ and $c = \bar{c} = \frac{1}{2}$ achieve the following rate:
 $\mathbb{E}[f(\bar{x}^K) - f(x^*)] \leq \frac{\|x^0 - x^*\|^2}{2\gamma_{b,0}\sqrt{K}} + \frac{\gamma_{b,0}G^2 \log(K)}{2\sqrt{K}}$. Next, we state the result for the envelop-type step size when f is μ -strongly convex.

Theorem 2. Suppose a μ -strongly convex function f satisfying Assumptions 1 and 2, assume C is a closed and convex set, and γ_k is independent of the sample $\nabla f_k(x^k)$. Then an envelope-type step size as in (9) with $\gamma_{b,k} = \frac{\gamma_{b,0}}{k+1}$, $\gamma_{b,0} \geq \frac{1}{\mu}$, and $\omega\mu < 1$ achieves the following rate

$$\mathbb{E}[f(\bar{x}_K) - f(x^*)] \leq \frac{\mu k_0}{2(K - k_0)} (e^{-k_0\mu\omega} \|x_0 - x^*\|^2 + \gamma_{b,0}^2 G^2) + \frac{\gamma_{b,0} G^2 \log K}{2(K - k_0)},$$

where $\bar{x}_K = \frac{1}{K - k_0} \sum_{k=k_0}^{K-1} x^k$ and $k_0 = \max\{1, \lceil \gamma_{b,0}/\omega \rceil - 1\}$.

We can again apply the result of Theorem 2 to SPSB and SLSB with $\gamma_{b,k} = \frac{\gamma_{b,0}}{k+1}$, $\gamma_{b,0} \geq \frac{1}{\mu}$, $\omega = 1/L_{\max}$, and $c = \bar{c} = \frac{1}{2}$ to get an explicit rate: $\mathbb{E}[f(\bar{x}_K) - f(x^*)] \leq \frac{\mu k_0}{2(K - k_0)} (e^{-\frac{k_0\mu}{L_{\max}}} \|x_0 - x^*\|^2 + \gamma_{b,0}^2 G^2) + \frac{\gamma_{b,0} G^2 \log K}{2(K - k_0)}$, where $k_0 = \max\{1, \lceil \gamma_{b,0} L_{\max} \rceil - 1\}$.

Remark 1. Under the envelop-type step size framework and the assumption of two independent samples, SLSB and SPSB share the same convergence rates of $\mathcal{O}(\frac{1}{\sqrt{K}})$ and $\mathcal{O}(\frac{1}{K})$ as SGD with decaying step-size for convex and strongly-convex losses respectively. This is not surprising because of the structure of the envelope step-size in (9). Indeed, the proof is similar to the standard proof of analogous rate for SGD with decaying step-size. Nonetheless, we include it here for completeness.

3.2 Envelope-type step size for bi-level optimization

We start with recalling standard assumptions in BO [22, 19, 15, 4]. We denote $z = [x; y]$ and recall the bi-level problem in (1). The first assumption is on the lower-level objective g .

Assumption 3. The function $g(x, y)$ is μ_g strongly convex in y for any given x . Moreover, ∇g is Lipschitz continuous: $\|\nabla g(x_1, y_1) - \nabla g(x_2, y_2)\| \leq L_g \|z_1 - z_2\|$ (also assume that this holds true for each sampled function $g(x, y; \psi)$), and $\nabla^2 g$ is Lipschitz continuous: $\|\nabla^2 g(x_1, y_1) - \nabla^2 g(x_2, y_2)\| \leq L_G \|z_1 - z_2\|$. We further assume that $\|\nabla_{xy}^2 g(x, y)\| \leq C_g$, and the condition number is defined as $\kappa = \frac{L_g}{\mu_g}$.

Next, we state the assumptions on the upper objective f .

Assumption 4. The function f and its gradients are Lipschitz continuous. That is: $\|f(x_1, y_1) - f(x_2, y_2)\| \leq L_1 \|z_1 - z_2\|$ and $\|\nabla f(x_1, y_1) - \nabla f(x_2, y_2)\| \leq L_{f,1} \|z_1 - z_2\|$. We also assume that $\|\nabla_y f(x, y)\| \leq C_f$.

Furthermore, we make the following standard assumptions on the estimates of ∇f , ∇g , and $\nabla^2 g$.

Assumption 5. The stochastic gradients are unbiased: $\mathbb{E}_\phi[\nabla f(x, y; \phi)] = \nabla f(x, y)$, $\mathbb{E}_\psi[\nabla g(x, y; \psi)] = \nabla g(x, y)$, and $\mathbb{E}_\psi[\nabla^2 g(x, y; \psi)] = \nabla^2 g(x, y)$. The variances of $\nabla f(x, y; \phi)$ and $\nabla^2 g(x, y; \psi)$ are bounded: $\mathbb{E}_\phi[\|\nabla f(x, y; \phi) - \nabla f(x, y)\|^2] \leq \sigma_f^2$ and $\mathbb{E}_\psi[\|\nabla^2 g(x, y; \psi) - \nabla^2 g(x, y)\|^2] \leq \sigma_G^2$.

Finally, we introduce the bounded optimal function value assumption in (15), which is used specifically for analyzing step size of the form (11) in the bi-level setting:

$$\mathbb{E}_\psi[g(x, y^*(x); \psi) - g(x, y_{x,\psi}^*; \psi)] \leq \sigma_g^2, \forall x, \quad (15)$$

$$\mathbb{E}[\|\nabla_y g(x, y) - \nabla_y g(x, y; \phi)\|^2] \leq \sigma_g^2, \forall x, y, \quad (16)$$

where $y^*(x) = \min_y g(x, y)$ and $y_{x,\psi}^* = \min_y g(x, y; \psi)$ for a given x (recall that at any iteration k , the lower-level steps in BiSPS are SPS_{max} with an upper bound $\beta_{b,k}$; furthermore, $\beta_{b,k}$ is non-increasing w.r.t. upper iteration k). The one-variable analogous assumption of (15) has been used in the analysis of SPS_{max} [29]. Here, we extend it to a two-variable function. Unlike the bounded variance assumption (16), which needs to hold true for all x and y , we require (15) to hold at $y^*(x)$ for any given x . As mentioned previously, the closed form solution $y^*(x)$ is difficult to obtain. Thus, we define the following expression by replacing $y^*(x)$ with y in (2):

$$\bar{\nabla} f(x, y) = \nabla_x f(x, y) + \nabla_{xy}^2 g(x, y) [\nabla_{yy}^2 g(x, y)]^{-1} \nabla_y f(x, y). \quad (17)$$

A stochastic Neumann series in (13) approximates (17) with x and y being x^k and y^{k+1} (respectively), also recall that y^{k+1} is an approximation of $y^*(x^k)$ by running T lower-level SGD steps to minimize g w.r.t. y for a fixed x^k . Based on Assumptions 3, 4, and 5, we have the following results to be used in the analysis [15]: $\|\nabla F(x_1) - \nabla F(x_2)\| \leq L_F \|x_1 - x_2\|$, $\|y^*(x_1) - y^*(x_2)\| \leq L_y \|x_1 - x_2\|$, and $\|\bar{\nabla} f(x, y^*(x)) - \bar{\nabla} f(x, y)\| \leq L_f \|y^*(x) - y\|$. Furthermore, the bias in the stochastic hypergradient in (13) (denoted as B) decays exponentially with N and its variance is bounded, i.e. $\mathbb{E}[\|h_f^k - \mathbb{E}[h_f^k]\|^2] \leq \tilde{\sigma}_f^2$ (see Appendix A for details) [19].

Now, we state our main theorem based on step size of the form (10) and (11).

Theorem 3. Suppose f and g satisfy Assumptions 3, 4, and 5, learning-rate upper bounds $\alpha_{b,k} = \frac{\alpha_{b,0}}{\sqrt{k+1}}$ and $\alpha_{l,k} = \frac{\alpha_{l,0}}{\sqrt{k+1}}$ with $\alpha_{b,0}$ and $\alpha_{l,0}$ satisfying $\frac{1}{L_F + 4L_g^2} \geq \frac{\alpha_{b,0}^2}{\alpha_{l,0}}$ and $\alpha_{l,0} \leq \alpha_{b,0}$. Further assume that α_k is independent of the stochastic hypergradient h_f^k , and each sampled function $g(x, y; \psi)$ is convex. Then under the Assumption (15) with $p \geq \frac{1}{2}$, $C_k = \min\{\frac{1}{2pL_g}, \beta_{b,k}\}$, $T \geq \frac{\log(\alpha_{b,0} L_f^2 + 2)}{-\log(1 - \mu_g C_{K-1})}$, and $\beta_{b,k} = \frac{\beta_{b,0}}{k+1}$, BiSPS achieves the rate:

$$\frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E}[\|\nabla F(x^k)\|^2] \leq \tilde{O}\left(\frac{\kappa^3}{\sqrt{K}} + \frac{\kappa^2 \log K}{\sqrt{K}}\right). \quad (18)$$

Remark 2. We further give the convergence result under the bounded variance assumption (16) in Appendix A. Theorem 3 shows that BiSPS matches the optimal rate of SGD up to a logarithmic factor without a growing batch size. We notice that the assumption (15) largely simplifies the expression on T and does not require an explicit upper bound on $\beta_{b,0}$. As in the single-level case, whether using one sample or two samples (which makes upper-level step-size independent of gradient) gives similar empirical performances (see Appendix B). Note that the independence assumption is only needed for the upper-level. Thus, the two-sample requirement of theorem does not apply to the lower-level problem. This is useful from computational standpoint as typical bi-level algorithms run multiple lower-level updates for each upper-level iteration.

4 Additional Hyper-Representation and Data Distillation Experiments

Hyper-representation learning: The experiments are performed on MNIST dataset using LeNet [26, 42]. We use conjugate gradient method for solving system of equations when computing the




Figure 6: Validation loss (a) and accuracy (b) against iterations. (a) Comparisons between whether to use or not use line-search at the upper or lower level; (b) Generalization performance of BiSLS-Adam/SGD and fine-tuned Adam/SGD; (c) Validation loss against iterations for different values of $\alpha_{b,0}$ ($\beta_{b,0}$ fixed at 100). (d) Same plot as (c) but for different values of $\beta_{b,0}$ ($\alpha_{b,0}$ fixed at 10).

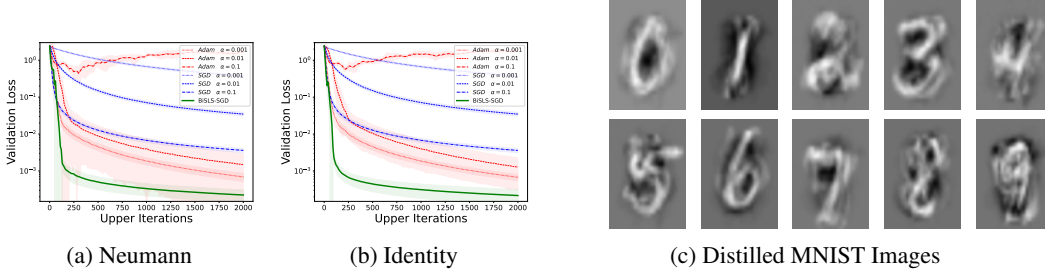


Figure 7: (a)(b): Comparison between BiSLS-SGD and Adam/SGD for Data Distillation on MNIST dataset. Validation loss plotted against iterations. (a) Hypergradient computed using Neumann series; (b) Inverse Hessian in (2) treated as the Identity [30] when computing the hypergradient; (c) Distilled MNIST images after 3000 iterations of BiSLS-SGD.

hypergradient [17]. The upper and lower-level objectives are to optimize the embedding layers and the classifier (i.e. the last layer of the neural net), respectively (see Appendix B for details). For constant-step SGD and Adam, we tune the lower-level learning rate $\beta \in \{10.0, 5.0, 1.0, 0.5, 0.1, 0.05, 0.01\}$. For the upper-level learning rate, we tune $\alpha \in \{0.001, 0.0025, 0.005, 0.01, 0.05, 0.1\}$ for SGD, and $\alpha \in \{10^{-5}, 5 \cdot 10^{-5}, 10^{-4}, 5 \cdot 10^{-4}, 0.001, 0.01\}$ for Adam (recall that δ in (14) is set to 0). Based on the results of Figure 6, we make the following key observations: ① **line-search at the upper-level is essential for achieving the optimal performance** (Figure 6a); ② **BiSLS-Adam/SGD not only converges fast but also generalizes well** (Figure 6b); ③ **BiSLS-Adam/SGD is highly robust to search starting points $\alpha_{b,0}$ and $\beta_{b,0}$** (Figure 6c, 6d). It addresses the fundamental question of how to tune α and β in bi-level optimization (see Appendix B for additional results on search cost). **Data distillation:** The goal of data distillation is to generate a small set of synthetic data from an original dataset that preserves the performance of a model when trained using the generated data [46, 50]. We adapted the experiment set up from Lorraine et al. [30] to distill MNIST digits. We present the results in Figure 7, where we observe that BiSLS-SGD converges significantly faster than fine-tuned Adam or SGD, and generate realistic MNIST images (see Appendix B for more results).

5 Conclusion

In this work, we have given simple alternatives to SLS and SPS that show good empirical performance in non-interpolating scenario without requiring the step size to be monotonic. We unify their analysis based on a simplified envelope-type step size, and extend the analysis to the bi-level setting while designing a SPS-based bi-level algorithm. In the end, we propose bi-level line-search algorithm BiSLS-Adam/SGD that is empirically truly robust and adaptive to learning rate initialization. Our work opens several possible future directions. Given the superior performance of BiSLS, we prioritize an analysis of its convergence rates. The difficulty stems from: (a) the bias in hypergradient estimation; (b) the dual updates in x and $y^*(x)$ (incurring nested loop structures); (c) the error in estimating $y^*(x)$. On single-level optimization, we remark as an important direction to relax the two-sample assumption on SPSB/SLSB. Ultimately, we hope to promote further research on bi-level optimization algorithms with minimal tuning.

References

- [1] Amir Beck. *First-order methods in optimization*. SIAM, 2017.
- [2] Zalán Borsos, Mojmir Mutny, and Andreas Krause. Coresets via bilevel optimization for continual learning and streaming. *Advances in Neural Information Processing Systems*, 33: 14879–14890, 2020.
- [3] Chih-Chung Chang and Chih-Jen Lin. Libsvm: a library for support vector machines. *ACM transactions on intelligent systems and technology (TIST)*, 2(3):1–27, 2011.
- [4] Tianyi Chen, Yuejiao Sun, and Wotao Yin. Tighter analysis of alternating stochastic gradient method for stochastic nested problems, 2021.
- [5] Tianyi Chen, Yuejiao Sun, Quan Xiao, and Wotao Yin. A single-timescale method for stochastic bilevel optimization, 2022.
- [6] Ashok Cutkosky and Francesco Orabona. Momentum-based variance reduction in non-convex sgd, 2020.
- [7] Mathieu Dagr  ou, Pierre Ablin, Samuel Vaiter, and Thomas Moreau. A framework for bilevel optimization that enables stochastic and global variance reduction algorithms, 2022.
- [8] Mathieu Dagr  ou, Thomas Moreau, Samuel Vaiter, and Pierre Ablin. A lower bound and a near-optimal algorithm for bilevel empirical risk minimization, 2023.
- [9] Aaron Defazio, Francis Bach, and Simon Lacoste-Julien. Saga: A fast incremental gradient method with support for non-strongly convex composite objectives, 2014.
- [10] James E Falk and Jiming Liu. On bilevel programming, part i: general nonlinear cases. *Mathematical Programming*, 70:47–72, 1995.
- [11] Chen Fan, Parikshit Ram, and Sijia Liu. Sign-maml: Efficient model-agnostic meta-learning by signsgd, 2021.
- [12] Cong Fang, Chris Junchi Li, Zhouchen Lin, and Tong Zhang. Spider: Near-optimal non-convex optimization via stochastic path integrated differential estimator, 2018.
- [13] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks, 2017.
- [14] Luca Franceschi, Michele Donini, Paolo Frasconi, and Massimiliano Pontil. Forward and reverse gradient-based hyperparameter optimization, 2017.
- [15] Saeed Ghadimi and Mengdi Wang. Approximation methods for bilevel programming, 2018.
- [16] Robert Mansel Gower, Nicolas Loizou, Xun Qian, Alibek Sailanbayev, Egor Shulgin, and Peter Richt  rik. Sgd: General analysis and improved rates. In *International conference on machine learning*, pages 5200–5209. PMLR, 2019.
- [17] Riccardo Grazzi, Luca Franceschi, Massimiliano Pontil, and Saverio Salzo. On the iteration complexity of hypergradient computation, 2020.
- [18] Riccardo Grazzi, Massimiliano Pontil, and Saverio Salzo. Convergence properties of stochastic hypergradients, 2021.
- [19] Mingyi Hong, Hoi-To Wai, Zhaoran Wang, and Zhuoran Yang. A two-timescale framework for bilevel optimization: Complexity analysis and application to actor-critic, 2022.
- [20] Feihu Huang, Junyi Li, and Shangqian Gao. Biadam: Fast adaptive bilevel optimization methods, 2023.
- [21] Yo Ishizuka and Eitaro Aiyoshi. Double penalty method for bilevel optimization problems. *Annals of Operations Research*, 34(1):73–88, 1992.
- [22] Kaiyi Ji, Junjie Yang, and Yingbin Liang. Bilevel optimization: Convergence analysis and enhanced design, 2021.
- [23] Kaiyi Ji, Mingrui Liu, Yingbin Liang, and Lei Ying. Will bilevel optimizers benefit from loops, 2022.
- [24] Prashant Khanduri, Siliang Zeng, Mingyi Hong, Hoi-To Wai, Zhaoran Wang, and Zhuoran Yang. A near-optimal algorithm for stochastic bilevel optimization via double-momentum, 2021.

- [25] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [26] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [27] Xiaoyu Li and Francesco Orabona. On the convergence of stochastic gradient descent with adaptive stepsizes, 2019.
- [28] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search, 2019.
- [29] Nicolas Loizou, Sharan Vaswani, Issam Laradji, and Simon Lacoste-Julien. Stochastic polyak step-size for sgd: An adaptive learning rate for fast convergence, 2021.
- [30] Jonathan Lorraine, Paul Vicol, and David Duvenaud. Optimizing millions of hyperparameters by implicit differentiation. In *International Conference on Artificial Intelligence and Statistics*, pages 1540–1552. PMLR, 2020.
- [31] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization, 2019.
- [32] Liangchen Luo, Yuanhao Xiong, Yan Liu, and Xu Sun. Adaptive gradient methods with dynamic bound of learning rate, 2019.
- [33] Arkadi Nemirovski, Anatoli Juditsky, Guanghui Lan, and Alexander Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on optimization*, 19(4): 1574–1609, 2009.
- [34] Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2003.
- [35] Lam M Nguyen, Jie Liu, Katya Scheinberg, and Martin Takáč. Sarah: A novel method for machine learning problems using stochastic recursive gradient. In *International Conference on Machine Learning*, pages 2613–2621. PMLR, 2017.
- [36] Antonio Orvieto, Simon Lacoste-Julien, and Nicolas Loizou. Dynamics of sgd with stochastic polyak stepsizes: Truly adaptive variants and convergence to exact solution, 2022.
- [37] Aravind Rajeswaran, Chelsea Finn, Sham Kakade, and Sergey Levine. Meta-learning with implicit gradients, 2019.
- [38] Sashank J Reddi, Satyen Kale, and Sanjiv Kumar. On the convergence of adam and beyond. *arXiv preprint arXiv:1904.09237*, 2019.
- [39] Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zhihui Li, Xiaojiang Chen, and Xin Wang. A comprehensive survey of neural architecture search: Challenges and solutions. *ACM Computing Surveys (CSUR)*, 54(4):1–34, 2021.
- [40] Amirreza Shaban, Ching-An Cheng, Nathan Hatch, and Byron Boots. Truncated back-propagation for bilevel optimization. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1723–1732. PMLR, 2019.
- [41] Shai Shalev-Shwartz, Yoram Singer, and Nathan Srebro. Pegasos: Primal estimated sub-gradient solver for svm. In *Proceedings of the 24th international conference on Machine learning*, pages 807–814, 2007.
- [42] Daouda Sow, Kaiyi Ji, and Yingbin Liang. On the convergence theory for hessian-free bilevel algorithms, 2022.
- [43] Sharan Vaswani, Issam Laradji, Frederik Kunstner, Si Yi Meng, Mark Schmidt, and Simon Lacoste-Julien. Adaptive gradient methods converge faster with over-parameterization (but you should do a line-search). *arXiv preprint arXiv:2006.06835*, 2020.
- [44] Sharan Vaswani, Aaron Mishkin, Issam Laradji, Mark Schmidt, Gauthier Gidel, and Simon Lacoste-Julien. Painless stochastic gradient: Interpolation, line-search, and convergence rates, 2021.
- [45] Luis Vicente, Gilles Savard, and Joaquim Júdice. Descent approaches for quadratic bilevel programming. *Journal of optimization theory and applications*, 81(2):379–399, 1994.
- [46] Tongzhou Wang, Jun-Yan Zhu, Antonio Torralba, and Alexei A. Efros. Dataset distillation, 2020.

- 445 [47] Rachel Ward, Xiaoxia Wu, and Leon Bottou. Adagrad stepsizes: Sharp convergence over
446 nonconvex landscapes, 2021.
- 447 [48] Jorge Nocedal Stephen J Wright. Numerical optimization, 2006.
- 448 [49] Junjie Yang, Kaiyi Ji, and Yingbin Liang. Provably faster algorithms for bilevel optimization,
449 2021.
- 450 [50] Ruonan Yu, Songhua Liu, and Xinchao Wang. Dataset distillation: A comprehensive review,
451 2023.
- 452 [51] Yihua Zhang, Guanhua Zhang, Prashant Khanduri, Mingyi Hong, Shiyu Chang, and Sijia Liu.
453 Revisiting and advancing fast adversarial training through the lens of bi-level optimization. In
454 *International Conference on Machine Learning*, pages 26693–26712. PMLR, 2022.
- 455 [52] Yihua Zhang, Yuguang Yao, Parikshit Ram, Pu Zhao, Tianlong Chen, Mingyi Hong, Yanzhi
456 Wang, and Sijia Liu. Advancing model pruning via bi-level optimization, 2023.
- 457 [53] Yongchao Zhou, Ehsan Nezhadarya, and Jimmy Ba. Dataset distillation using neural feature
458 regression. *arXiv preprint arXiv:2206.00719*, 2022.

Notes: A few small typos have been revised in the main body. The $>$ sign in (14) is corrected to \leq , and the missing A_k^{-1} for BiSLS-Adam in (14) is added. In Theorem 3 and the remark that follows, the letter γ is corrected to letter β . In Assumption 3, L_f is changed to $L_{f,1}$.

A Proofs of Theorems and Additional Convergence Results

A.1 Useful Lemmas

Lemma 1 provides more details on the envelope structure of SPSB and SLSB given in (10) and (11). The lower bound in (19) will also be used in Lemma 5 (for bounding the term $\|y^{k+1} - y^*(x^k)\|^2$).

Lemma 1. *Under the Assumption 1, we have the following:*

$$\text{SPSB : } \min\left\{\frac{1}{2cL_{\max}}, \gamma_{b,k}\right\} \leq \gamma_k = \min\left\{\frac{f_{i_k}(x^k) - l_{i_k}^*}{c\|\nabla f_{i_k}(x^k)\|^2}, \gamma_{b,k}\right\}, \quad 0 < c, \quad (19)$$

$$\text{SLSB : } \min\left\{\frac{2(1-\bar{c})}{L_{\max}}, \gamma_{b,k}\right\} \leq \gamma_k \leq \min\left\{\frac{f_{i_k}(x^k) - l_{i_k}^*}{\bar{c}\|\nabla f_{i_k}(x^k)\|^2}, \gamma_{b,k}\right\}, \quad 0 < \bar{c} < 1. \quad (20)$$

Proof. The bounds in (19) have been shown in [29, 36]. For (20), the first part of the inequality has been shown in [44]. For the second part, recall the Armijo condition (14):

$$f_{i_k}(x_k - \gamma_k \nabla f_{i_k}(x_k)) \leq f_{i_k}(x_k) - \bar{c} \cdot \gamma_k \|\nabla f_{i_k}(x_k)\|^2, \quad 0 < \bar{c} < 1.$$

We can then rearrange this to obtain

$$\gamma_k \leq \frac{f_{i_k}(x_k) - f_{i_k}(x_k - \gamma_k \nabla f_{i_k}(x_k))}{\bar{c}\|\nabla f_{i_k}(x_k)\|^2} \leq \frac{f_{i_k}(x_k) - f_{i_k}^*}{\bar{c}\|\nabla f_{i_k}(x_k)\|^2} \leq \frac{f_{i_k}(x_k) - l_{i_k}^*}{\bar{c}\|\nabla f_{i_k}(x_k)\|^2}, \quad (21)$$

where $l_{i_k}^*$ is any lower bound for $f_{i_k}^*$. Also recall that $\gamma_{b,k}$ is the search starting point at iteration k , hence (20) holds for SLSB. \square

Lemma 2 gives the expressions for the constants L_F , L_y , and L_f . The proof can be found in [15, Lem 2.2].

Lemma 2. *Under Assumptions 3 and 4, we have the following:*

$$\begin{aligned} \|\nabla F(x_1) - \nabla F(x_2)\| &\leq L_F \|x_1 - x_2\|, \\ \|y^*(x_1) - y^*(x_2)\| &\leq L_y \|x_1 - x_2\|, \\ \|\bar{\nabla} f(x, y^*(x)) - \bar{\nabla} f(x, y)\| &\leq L_f \|y^*(x) - y\|, \end{aligned}$$

where

$$\begin{aligned} L_f &= L_{f,1} + \frac{L_{f,1}L_g}{\mu_g} + \frac{L_1}{\mu_g}(L_G + \frac{L_gL_G}{\mu_g}) \sim \mathcal{O}(\kappa^2) \\ L_y &= \frac{L_g}{\mu_g} \sim \mathcal{O}(\kappa) \\ L_F &= L_{f,1} + \frac{L_g(L_{f,1} + L_f)}{\mu_g} + \frac{L_1}{\mu_g}(L_G + \frac{L_gL_G}{\mu_g}) \sim \mathcal{O}(\kappa^3) \end{aligned}$$

Lemma 3 is on the bias and variance of the stochastic hypergradient in (13), which has the following form [19]

$$h_f^k = \nabla_x f(x^k, y^{k+1}; \phi) - \nabla_{xy} g(x^k, y^{k+1}; \psi_0) \left[\frac{N}{L_g} \prod_{j=1}^{\bar{N}} (I - \nabla_{yy}^2 g(x^k, y^{k+1}; \psi_j)) \right] \nabla_y f(x^k, y^{k+1}; \phi). \quad (22)$$

Recall that the hypergradient surrogate defined in (17) based on (x^k, y^{k+1}) is

$$\bar{\nabla} f(x^k, y^{k+1}) = \nabla_x f(x^k, y^{k+1}) + \nabla_{xy}^2 g(x^k, y^{k+1}) [\nabla_{yy}^2 g(x^k, y^{k+1})]^{-1} \nabla_y f(x^k, y^{k+1}). \quad (23)$$

Given a filtration \mathcal{F}_k' up to and including x^k and y^{k+1} , the bias of the stochastic hypergradient is defined as $B = \bar{\nabla} f(x^k, y^{k+1}) - \mathbb{E}[h_f^k | \mathcal{F}_k']$, and the variance is defined as $\mathbb{E}[\|\bar{\nabla} f(x^k, y^{k+1}) - \mathbb{E}[h_f^k | \mathcal{F}_k']\|^2]$. Lemma 3 has been proven in [19, Lem 1.] (also see [4, Lem 5.]). Lemmas 1, 2, and 3 will be used in the proofs of Theorems 3 and 4.

483 **Lemma 3.** Under Assumptions 3, 4, and 5, the bias and variance of the stochastic hypergradient h_f^k
 484 satisfy the following

$$\text{Bias: } \|\bar{\nabla} f(x^k, y^{k+1}) - \mathbb{E}[h_f^k | \mathcal{F}_k']\| \leq \frac{C_g C_f}{\mu_g} (1 - \frac{\mu_g}{L_g})^N, \forall k$$

$$\text{Variance: } \mathbb{E}[\|\bar{\nabla} f(x^k, y^{k+1}) - \mathbb{E}[h_f^k | \mathcal{F}_k']\|^2] \leq \tilde{\sigma}_f^2, \forall k,$$

485 where N is the total number of samples, and $\tilde{\sigma}_f^2 = \sigma_f^2 + [(\sigma_f^2 + L_1^2)(\sigma_G^2 + 2L_g^2) + \sigma_f^2 L_g^2] \frac{3}{\mu_g^2} \sim \mathcal{O}(\kappa^2)$.

486 Lemma 4 is on the descent of the quantity $V^k := F(x^k) - F(x^*) + \|y^k - y^*(x^k)\|^2$. It will be used
 487 in the proofs of Theorem 3 and 4.

488 **Lemma 4.** Suppose F satisfies Assumptions 3, 4, and 5, sequences $\alpha_{b,k} = \frac{\alpha_{b,0}}{\sqrt{k+1}}$ and $\alpha_{l,k} = \frac{\alpha_{l,0}}{\sqrt{k+1}}$

489 with $\alpha_{b,0}$ and $\alpha_{l,0}$ satisfying $\frac{1}{L_F + 4L_y^2} \geq \frac{\alpha_{b,0}^2}{\alpha_{l,0}}$ and $\alpha_{l,0} \leq \alpha_{b,0}$. Further assume that α_k is independent
 490 of h_f^k . Then step sizes of the form (10) and (11) achieve the following:

$$\begin{aligned} E[V^{k+1}] &\leq \mathbb{E}[V^k] - \frac{\alpha_{l,k}}{2} \mathbb{E}[\|\nabla F(x^k)\|^2] + (\alpha_{b,k} L_f^2 + 2) \mathbb{E}[\|y^{k+1} - y^*(x^k)\|^2] + \\ &\quad \alpha_{b,k} B^2 + (2L_y^2 \alpha_{b,k}^2 + \frac{L_F \alpha_{b,k}^2}{2}) \tilde{\sigma}_f^2 - \mathbb{E}[\|y^k - y^*(x^k)\|^2], \forall k, \end{aligned} \quad (24)$$

491 where $V^k = F(x^k) - F(x^*) + \|y^{k+1} - y^*(x^{k+1})\|^2$ and recall that F is the upper-level loss defined
 492 in (1).

493 *Proof.* We denote $\mathbb{E}[h_f^k | \mathcal{F}_k'] = \bar{h}_f^k$. By the L_F -smoothness of the objective F :

$$F(x^{k+1}) \leq F(x^k) + \langle \nabla F(x^k), x^{k+1} - x^k \rangle + \frac{L_F}{2} \|x^{k+1} - x^k\|^2.$$

494 Take expectation conditioned on a filtration of past iterates \mathcal{F}_k' (up to and include x^k, y^{k+1}):

$$\begin{aligned} \mathbb{E}[F(x^{k+1}) | \mathcal{F}_k'] &\leq F(x^k) + \mathbb{E}[\langle \nabla F(x^k), x^{k+1} - x^k \rangle | \mathcal{F}_k'] + \frac{L_F}{2} \mathbb{E}[\|x^{k+1} - x^k\|^2 | \mathcal{F}_k'] \\ &= F(x^k) - \mathbb{E}[\alpha_k \langle \nabla F(x^k), h_f^k \rangle | \mathcal{F}_k'] + \frac{L_F \mathbb{E}[\alpha_k^2 | \mathcal{F}_k']}{2} \mathbb{E}[\|h_f^k - \bar{h}_f^k + \bar{h}_f^k\|^2 | \mathcal{F}_k'] \\ &\stackrel{(a)}{=} F(x^k) - \mathbb{E}[\alpha_k | \mathcal{F}_k'] \langle \nabla F(x^k), \bar{h}_f^k \rangle + \frac{L_F \mathbb{E}[\alpha_k^2 | \mathcal{F}_k']}{2} \|\bar{h}_f^k\|^2 + \frac{L_F \mathbb{E}[\alpha_k^2 | \mathcal{F}_k']}{2} \tilde{\sigma}_f^2 \\ &= F(x^k) - \frac{\mathbb{E}[\alpha_k | \mathcal{F}_k']}{2} \|\nabla F(x^k)\|^2 - \frac{\mathbb{E}[\alpha_k | \mathcal{F}_k']}{2} \|\bar{h}_f^k\|^2 + \frac{\mathbb{E}[\alpha_k | \mathcal{F}_k']}{2} \|\nabla F(x^k) - \bar{h}_f^k\|^2 + \\ &\quad \frac{L_F \mathbb{E}[\alpha_k^2 | \mathcal{F}_k']}{2} \|\bar{h}_f^k\|^2 + \frac{L_F \mathbb{E}[\alpha_k^2 | \mathcal{F}_k']}{2} \tilde{\sigma}_f^2, \end{aligned}$$

495 where (a) is by the assumption that α_k is independent of h_f^k . Then expand the term $\|\nabla F(x^k) - \bar{h}_f^k\|^2$
 496 as follows:

$$\begin{aligned} \|\nabla F(x^k) - \bar{h}_f^k\|^2 &= \|\nabla F(x^k) - \bar{\nabla} f(x^k, y^{k+1}) + \bar{\nabla} f(x^k, y^{k+1}) - \bar{h}_f^k\|^2 \\ &\leq 2\|\nabla F(x^k) - \bar{\nabla} f(x^k, y^{k+1})\|^2 + 2\|\bar{\nabla} f(x^k, y^{k+1}) - \bar{h}_f^k\|^2 \\ &\stackrel{(b)}{\leq} 2L_f^2 \|y^{k+1} - y^*(x^k)\|^2 + 2B^2, \end{aligned}$$

497 where (b) is by Lemma 2 and 3. Substituting this into the above:

$$\begin{aligned} \mathbb{E}[F(x^{k+1}) | \mathcal{F}_k'] &\leq F(x^k) - \frac{\mathbb{E}[\alpha_k | \mathcal{F}_k']}{2} \|\nabla F(x^k)\|^2 - \frac{\mathbb{E}[\alpha_k | \mathcal{F}_k']}{2} \|\bar{h}_f^k\|^2 + \frac{L_F \mathbb{E}[\alpha_k^2 | \mathcal{F}_k']}{2} \|\bar{h}_f^k\|^2 + \\ &\quad E[\alpha_k | \mathcal{F}_k'] L_f^2 \|y^{k+1} - y^*(x^k)\|^2 + E[\alpha_k | \mathcal{F}_k'] B^2 + \frac{L_F \mathbb{E}[\alpha_k^2 | \mathcal{F}_k']}{2} \tilde{\sigma}_f^2 \\ &\stackrel{(c)}{\leq} F(x^k) - \frac{\alpha_{l,k}}{2} \|\nabla F(x^k)\|^2 - \frac{\alpha_{l,k}}{2} \|\bar{h}_f^k\|^2 + \frac{L_F \alpha_{b,k}^2}{2} \|\bar{h}_f^k\|^2 + \\ &\quad \alpha_{b,k} L_f^2 \|y^{k+1} - y^*(x^k)\|^2 + \alpha_{b,k} B^2 + \frac{L_F \alpha_{b,k}^2}{2} \tilde{\sigma}_f^2, \end{aligned}$$

where (c) is by $\alpha_{l,k} \leq \alpha_k$ and $\alpha_{b,k} \geq \alpha_k$. Then take total expectations and subtract $F(x^*)$:

$$\begin{aligned} \mathbb{E}[F(x^{k+1}) - F(x^*)] &\leq \mathbb{E}[F(x^k) - F(x^*)] - \frac{\alpha_{l,k}}{2} \mathbb{E}[\|\nabla F(x^k)\|^2] - \frac{\alpha_{l,k}}{2} \mathbb{E}[\|\bar{h}_f^k\|^2] + \frac{L_F \alpha_{b,k}^2}{2} \mathbb{E}[\|\bar{h}_f^k\|^2] + \\ &\quad \alpha_{b,k} L_f^2 \mathbb{E}[\|y^{k+1} - y^*(x^k)\|^2] + \alpha_{b,k} B^2 + \frac{L_F \alpha_{b,k}^2}{2} \tilde{\sigma}_f^2 \end{aligned} \quad (25)$$

Now define the potential function $V^k := F(x^k) - F(x^*) + \|y^{k+1} - y^*(x^{k+1})\|^2$ and expand the term $\|y^{k+1} - y^*(x^{k+1})\|^2$ as follows:

$$\begin{aligned} \|y^{k+1} - y^*(x^{k+1})\|^2 &= \|y^{k+1} - y^*(x^k) + y^*(x^k) - y^*(x^{k+1})\|^2 \\ &\leq 2\|y^{k+1} - y^*(x^k)\|^2 + 2\|y^*(x^k) - y^*(x^{k+1})\|^2 \\ &\stackrel{(d)}{\leq} 2\|y^{k+1} - y^*(x^k)\|^2 + 2L_y^2 \|x^{k+1} - x^k\|^2 \\ &= 2\|y^{k+1} - y^*(x^k)\|^2 + 2L_y^2 \alpha_k^2 \|h_f^k\|^2 \\ &= 2\|y^{k+1} - y^*(x^k)\|^2 + 2L_y^2 \alpha_k^2 \|h_f^k - \bar{h}_f^k + \bar{h}_f^k\|^2, \end{aligned}$$

where (d) is by Lemma 2. Take expectation conditioned on \mathcal{F}'_k :

$$\begin{aligned} \mathbb{E}[\|y^{k+1} - y^*(x^{k+1})\|^2 | \mathcal{F}'_k] &\leq 2\|y^{k+1} - y^*(x^k)\|^2 + 2L_y^2 \alpha_k^2 \|\bar{h}_f^k\|^2 + 2L_y^2 \alpha_k^2 \tilde{\sigma}_f^2 \\ &\leq 2\|y^{k+1} - y^*(x^k)\|^2 + 2L_y^2 \alpha_{b,k}^2 \|\bar{h}_f^k\|^2 + 2L_y^2 \alpha_{b,k}^2 \tilde{\sigma}_f^2. \end{aligned}$$

Then, take total expectations:

$$\mathbb{E}[\|y^{k+1} - y^*(x^{k+1})\|^2] \leq 2\mathbb{E}[\|y^{k+1} - y^*(x^k)\|^2] + 2L_y^2 \alpha_{b,k}^2 \mathbb{E}[\|\bar{h}_f^k\|^2] + 2L_y^2 \alpha_{b,k}^2 \tilde{\sigma}_f^2. \quad (26)$$

Now, based on the definition of V^k and combining (25) and (26):

$$\begin{aligned} \mathbb{E}[V^{k+1}] &\leq \mathbb{E}[F(x^k) - F(x^*)] - \frac{\alpha_{l,k}}{2} \mathbb{E}[\|\nabla F(x^k)\|^2] \\ &\quad - \frac{\mathbb{E}[\|\bar{h}_f^k\|^2]}{2} (\alpha_{l,k} - L_F \alpha_{b,k}^2 - 4L_y^2 \alpha_{b,k}^2) + (\alpha_{b,k} L_f^2 + 2) \mathbb{E}[\|y^{k+1} - y^*(x^k)\|^2] + \\ &\quad \alpha_{b,k} B^2 + (2L_y^2 \alpha_{b,k}^2 + \frac{L_F \alpha_{b,k}^2}{2}) \tilde{\sigma}_f^2 \\ &\stackrel{(e)}{\leq} \mathbb{E}[F(x^k) - F(x^*)] - \frac{\alpha_{l,k}}{2} \mathbb{E}[\|\nabla F(x^k)\|^2] + (\alpha_{b,k} L_f^2 + 2) \mathbb{E}[\|y^{k+1} - y^*(x^k)\|^2] + \\ &\quad \alpha_{b,k} B^2 + (2L_y^2 \alpha_{b,k}^2 + \frac{L_F \alpha_{b,k}^2}{2}) \tilde{\sigma}_f^2 \\ &= \mathbb{E}[V^k] - \frac{\alpha_{l,k}}{2} \mathbb{E}[\|\nabla F(x^k)\|^2] + (\alpha_{b,k} L_f^2 + 2) \mathbb{E}[\|y^{k+1} - y^*(x^k)\|^2] + \\ &\quad \alpha_{b,k} B^2 + (2L_y^2 \alpha_{b,k}^2 + \frac{L_F \alpha_{b,k}^2}{2}) \tilde{\sigma}_f^2 - \mathbb{E}[\|y^k - y^*(x^k)\|^2], \end{aligned}$$

where (e) is because $\frac{1}{L_F + 4L_y^2} \geq \frac{\alpha_{b,0}^2}{\alpha_{l,0}}$, which guarantees that $\alpha_{l,k} = \frac{\alpha_{l,0}}{\sqrt{k+1}} \geq (L_F + 4L_y^2) \alpha_{b,k} = (L_F + 4L_y^2) \frac{\alpha_{b,0}^2}{k+1}$. \square

Lemma 5 and Lemma 6 give two alternatives for bounding the term $\|y^{k+1} - y^*(x^k)\|^2$. Lemma 5 is based on the assumption $\mathbb{E}_\psi[g(x, y^*(x); \psi) - g(x, y_{x,\psi}^*; \psi)] \leq \sigma_g^2, \forall x$. The proof for its one-variable analogous assumption is given in Loizou et al. [29]. Here we follow a similar approach for the two-variable function $g(x, y)$. Lemma 5 will be used in the proof of Theorem 3.

Lemma 5. Suppose Assumptions 3, 5, and the bounded optimal function value assumption (15) hold. Further assume that each sampled function $g(x, y; \psi)$ is convex, then step size of the form 11 achieves the following:

$$\mathbb{E}[\|y^{k+1} - y^*(x^k)\|^2] \leq (1 - \mu_g C_k)^T \mathbb{E}[\|y^k - y^*(x^k)\|^2] + 2\beta_{b,k} T \sigma_g^2,$$

where $C_k = \min\{\frac{1}{2pL_g}, \beta_{b,k}\}$.

514 *Proof.* We denote $h_g^{k,t} = \nabla_y g(x^k, y^{k,t}; \psi)$, and $\mathcal{F}'_{k,t}$ be a filtration up to and including x^k and $y^{k,t}$.
 515 We have,

$$\begin{aligned}
 \|y^{k,t+1} - y^*(x^k)\|^2 &= \|y^{k,t} - \beta_{k,t} h_g^{k,t} - y^*(x^k)\|^2 \\
 &= \|y^{k,t} - y^*(x^k)\|^2 - 2\beta_{k,t} \langle y^{k,t} - y^*(x^k), h_g^{k,t} \rangle + \beta_{k,t}^2 \|h_g^{k,t}\|^2 \\
 &\stackrel{(a)}{\leq} \|y^{k,t} - y^*(x^k)\|^2 - 2\beta_{k,t} \langle y^{k,t} - y^*(x^k), h_g^{k,t} \rangle + \frac{\beta_{k,t}}{p} [g(x^k, y^{k,t}; \psi) - g(x^k, y_{x^k, \psi}^*; \psi)] \\
 &\stackrel{(b)}{\leq} \|y^{k,t} - y^*(x^k)\|^2 - 2\beta_{k,t} \langle y^{k,t} - y^*(x^k), h_g^{k,t} \rangle + 2\beta_{k,t} [g(x^k, y^{k,t}; \psi) - g(x^k, y_{x^k, \psi}^*; \psi)] \\
 &= \|y^{k,t} - y^*(x^k)\|^2 - 2\beta_{k,t} \langle y^{k,t} - y^*(x^k), h_g^{k,t} \rangle + \\
 &\quad 2\beta_{k,t} [g(x^k, y^{k,t}; \psi) - g(x^k, y^*(x^k); \psi) + g(x^k, y^*(x^k); \psi) - g(x^k, y_{x^k, \psi}^*; \psi)] \\
 &= \|y^{k,t} - y^*(x^k)\|^2 + 2\beta_{k,t} [-\langle y^{k,t} - y^*(x^k), h_g^{k,t} \rangle + g(x^k, y^{k,t}; \psi) - g(x^k, y^*(x^k); \psi)] + \\
 &\quad 2\beta_{k,t} [g(x^k, y^*(x^k); \psi) - g(x^k, y_{x^k, \psi}^*; \psi)] \\
 &\stackrel{(c)}{\leq} \|y^{k,t} - y^*(x^k)\|^2 + 2C_k [-\langle y^{k,t} - y^*(x^k), h_g^{k,t} \rangle + g(x^k, y^{k,t}; \psi) - g(x^k, y^*(x^k); \psi)] + \\
 &\quad 2\beta_{b,k} [g(x^k, y^*(x^k); \psi) - g(x^k, y_{x^k, \psi}^*; \psi)],
 \end{aligned}$$

516 where (a) is by Lemma 1, (b) is by choosing $p \geq \frac{1}{2}$, and, (c) is by individual convexity of
 517 $g(x, y; \psi)$ such that $-\langle y^{k,t} - y^*(x^k), h_g^{k,t} \rangle + g(x^k, y^{k,t}; \psi) - g(x^k, y^*(x^k); \psi) \leq 0$ and recalling
 518 that $C_k = \min\{\frac{1}{2pL_g}, \beta_{b,k}\} \leq \beta_{k,t}$ by Lemma 1. Take expectation conditioned on $\mathcal{F}'_{k,t}$ and note that
 519 $\mathbb{E}[h_g^{k,t} | \mathcal{F}'_{k,t}] = \nabla_y g(x^k, y^{k+1})$, $\mathbb{E}[g(x^k, y^{k,t}; \psi) | \mathcal{F}'_{k,t}] = g(x^k, y^{k,t})$, and $\mathbb{E}[g(x^k, y^*(x^k); \psi)] =$
 520 $g(x^k, y^*(x^k))$:

$$\begin{aligned}
 \mathbb{E}[\|y^{k,t+1} - y^*(x^k)\|^2 | \mathcal{F}'_{k,t}] &\leq \|y^{k,t} - y^*(x^k)\|^2 + 2C_k [-\langle y^{k,t} - y^*(x^k), \nabla_y g(x^k, y^{k+1}) \rangle + \\
 &\quad g(x^k, y^{k,t}) - g(x^k, y^*(x^k))] + 2\beta_{b,k} \sigma_g^2.
 \end{aligned} \tag{27}$$

521 Now, based on the strong convexity of g w.r.t. y ,

$$-\langle y^{k,t} - y^*(x^k), \nabla_y g(x^k, y^{k+1}) \rangle + g(x^k, y^{k,t}) - g(x^k, y^*(x^k)) \leq \frac{\mu_g}{2} \|y^{k,t} - y^*(x^k)\|^2,$$

522 we can further obtain (by taking total expectations of (27) and using strong-convexity):

$$\mathbb{E}[\|y^{k,t+1} - y^*(x^k)\|^2] \leq (1 - \mu_g C_k) \mathbb{E}[\|y^{k,t} - y^*(x^k)\|^2] + 2\beta_{b,k} \sigma_g^2.$$

523 Solve this recursively from $t = 0$ to $t = T - 1$ (recall T is the total number of lower-level steps,
 524 $y^{k,0} = y^k$ and $y^{k,T} = y^{k,T}$):

$$\begin{aligned}
 \mathbb{E}[\|y^{k+1} - y^*(x^k)\|^2] &\leq (1 - \mu_g C_k)^T \mathbb{E}[\|y^k - y^*(x^k)\|^2] + 2\beta_{b,k} \sigma_g^2 \sum_{j=0}^{T-1} (1 - \mu_g C_k)^j \\
 &\leq (1 - \mu_g C_k)^T \mathbb{E}[\|y^k - y^*(x^k)\|^2] + 2\beta_{b,k} T \sigma_g^2.
 \end{aligned}$$

525 □

526 Lemma 6 is based on the standard bounded variance assumption $\mathbb{E}_\psi[\|\nabla_y g(x, y^*(x); \psi) -$
 527 $\nabla_y g(x, y^*(x))\|^2] \leq \sigma_g^2, \forall x, y$ in the bi-level optimization literature [19, 4]. Lemma 6 will be
 528 used in the proof of Theorem 4.

529 **Lemma 6.** Suppose Assumptions 3, 5 and the bounded variance assumption (16) hold. Suppose that
 530 $p \geq \max\{\frac{\mu_g}{\mu_g + L_g}, \frac{\mu_g + L_g}{4L_g}\}$, $\beta_{b,0} \leq \min\{\frac{2}{\mu_g + L_g}, \frac{\mu_g + L_g}{2\mu_g L_g}, \frac{1}{2pL_g - \frac{2\mu_g L_g}{\mu_g + L_g}}\}$. Then step size of the form
 531 11 achieves the following:

$$\mathbb{E}[\|y^{k+1} - y^*(x^k)\|^2] \leq (\frac{\beta_{b,k}}{C_k} - \frac{2\mu_g L_g}{\mu_g + L_g} \beta_{b,k})^T \mathbb{E}[\|y^k - y^*(x^k)\|^2] + T \beta_{b,k}^2 \sigma_g^2,$$

532 where $C_k = \min\{\frac{1}{2pL_g}, \beta_{b,k}\}$.

533 *Proof.* Similar to the proof of Lemma 5, we can start with

$$\|y^{k,t+1} - y^*(x^k)\|^2 = \|y^{k,t} - y^*(x^k)\|^2 - 2\beta_{k,t}\langle y^{k,t} - y^*(x^k), h_g^{k,t} \rangle + \beta_{k,t}^2 \|h_g^{k,t}\|^2.$$

534 Divide both sides by $\beta_{k,t}$

$$\frac{\|y^{k,t+1} - y^*(x^k)\|^2}{\beta_{k,t}} = \frac{\|y^{k,t} - y^*(x^k)\|^2}{\beta_{k,t}} - 2\langle y^{k,t} - y^*(x^k), h_g^{k,t} \rangle + \beta_{k,t} \|h_g^{k,t}\|^2.$$

535 Then use the facts that $\beta_{k,t} \leq \beta_{b,k}$ and $\beta_{k,t} \geq C_k$,

$$\frac{\|y^{k,t+1} - y^*(x^k)\|^2}{\beta_{b,k}} \leq \frac{\|y^{k,t} - y^*(x^k)\|^2}{C_k} - 2\langle y^{k,t} - y^*(x^k), h_g^{k,t} \rangle + \beta_{b,k} \|h_g^{k,t}\|^2.$$

536 Next, take expectation conditioned on the Filtration $\mathcal{F}'_{k,t}$ up to and including x^k and $y^{k,t}$

$$\begin{aligned} \frac{1}{\beta_{b,k}} \mathbb{E}[\|y^{k,t+1} - y^*(x^k)\|^2 | \mathcal{F}'_{k,t}] &\leq \frac{1}{C_k} \mathbb{E}[\|y^{k,t} - y^*(x^k)\|^2 | \mathcal{F}'_{k,t}] - 2\langle y^{k,t} - y^*(x^k), \nabla g(x^k, y^{k,t}) \rangle \\ &\quad + \beta_{b,k} \mathbb{E}[\|h_g^{k,t}\|^2 | \mathcal{F}'_{k,t}] \\ &= \frac{1}{C_k} \mathbb{E}[\|y^{k,t} - y^*(x^k)\|^2 | \mathcal{F}'_{k,t}] - 2\langle y^{k,t} - y^*(x^k), \nabla g(x^k, y^{k,t}) \rangle \\ &\quad + \beta_{b,k} \mathbb{E}[\|h_g^{k,t} - \nabla g(x^k, y^{k,t}) + \nabla g(x^k, y^{k,t})\|^2 | \mathcal{F}'_{k,t}] \\ &\stackrel{(a)}{\leq} \frac{1}{C_k} \mathbb{E}[\|y^{k,t} - y^*(x^k)\|^2 | \mathcal{F}'_{k,t}] - 2\langle y^{k,t} - y^*(x^k), \nabla g(x^k, y^{k,t}) \rangle \\ &\quad + \beta_{b,k} \sigma_g^2 + \beta_{b,k} \|\nabla g(x^k, y^{k,t})\|^2, \end{aligned}$$

537 where (a) is by the bounded variance assumption (16). Based on strong-convexity of g [34, Theorem
538 2.1.11], we have

$$\begin{aligned} \frac{1}{\beta_{b,k}} \mathbb{E}[\|y^{k,t+1} - y^*(x^k)\|^2 | \mathcal{F}'_{k,t}] &\leq \frac{1}{C_k} \|y^{k,t} - y^*(x^k)\|^2 - \frac{2\mu_g L_g}{\mu_g + L_g} \|y^{k,t} - y^*(x^k)\|^2 \\ &\quad - \frac{2}{\mu_g + L_g} \|\nabla g(x^k, y^{k,t})\|^2 + \beta_{b,k} \sigma_g^2 + \beta_{b,k} \|\nabla g(x^k, y^{k,t})\|^2 \end{aligned}$$

539 Multiply by $\beta_{b,k}$ in both sides, group terms, and take total expectations to reach

$$\begin{aligned} \mathbb{E}[\|y^{k,t+1} - y^*(x^k)\|^2] &\leq \left(\frac{\beta_{b,k}}{C_k} - \frac{2\mu_g L_g}{\mu_g + L_g} \beta_{b,k} \right) \mathbb{E}[\|y^{k,t} - y^*(x^k)\|^2] + \beta_{b,k}^2 \sigma_g^2 \\ &\quad + \beta_{b,k} \left(\beta_{b,k} - \frac{2}{\mu_g + L_g} \right) \|\nabla g(x^k, y^{k,t})\|^2 \\ &\stackrel{(b)}{\leq} \left(\frac{\beta_{b,k}}{C_k} - \frac{2\mu_g L_g}{\mu_g + L_g} \beta_{b,k} \right) \mathbb{E}[\|y^{k,t} - y^*(x^k)\|^2] + \beta_{b,k}^2 \sigma_g^2, \end{aligned}$$

540 where in (b) we have chosen $\beta_{b,k} \leq \beta_{b,0} \leq \frac{2}{\mu_g + L_g}, \forall k$. Solving this recursion similar to Lemma 5,
541 we obtain:

$$\mathbb{E}[\|y^{k+1} - y^*(x^k)\|^2] \leq \left(\frac{\beta_{b,k}}{C_k} - \frac{2\mu_g L_g}{\mu_g + L_g} \beta_{b,k} \right)^T \mathbb{E}[\|y^k - y^*(x^k)\|^2] + T \beta_{b,k}^2 \sigma_g^2. \quad (28)$$

542 In (28), we require $\frac{\beta_{b,k}}{C_k} - \frac{2\mu_g L_g}{\mu_g + L_g} \beta_{b,k} \geq 0$ (recall $C_k = \min\{\frac{1}{2pL_g}, \beta_{b,k}\} \leq \beta_{k,t} \leq \beta_{b,k}$), this is
543 equivalent to $C_k \leq \frac{\mu_g + L_g}{2\mu_g L_g}$. In case of $C_k = \frac{1}{2pL_g}$, we choose $p \geq \frac{\mu_g}{\mu_g + L_g}$ (to avoid contradictions,
544 we also choose $p \geq \frac{\mu_g + L_g}{4L_g}$). In case of $C_k = \beta_{b,k}$, we choose $\beta_{b,k} \leq \beta_{b,0} \leq \frac{\mu_g + L_g}{2\mu_g L_g}$.

545 We also require $\frac{\beta_{b,k}}{C_k} - \frac{2\mu_g L_g}{\mu_g + L_g} \beta_{b,k} \leq 1$. In case of $C_k = \beta_{b,k}$, this is equivalent to
546 $\frac{2\mu_g L_g}{\mu_g + L_g} \beta_{b,k} \geq 0$, which is satisfied by all $\beta_{b,k}$. In case of $C_k = \frac{1}{2pL_g}$, we choose $\beta_{b,k} \leq$
547 $\beta_{b,0} \leq \frac{1}{2pL_g - \frac{2\mu_g L_g}{\mu_g + L_g}}$. Putting everything together, we have $p \geq \max\{\frac{\mu_g}{\mu_g + L_g}, \frac{\mu_g + L_g}{4L_g}\}$, and
548 $\beta_{b,0} \leq \min\{\frac{2}{\mu_g + L_g}, \frac{\mu_g + L_g}{2\mu_g L_g}, \frac{1}{2pL_g - \frac{2\mu_g L_g}{\mu_g + L_g}}\}$. \square

549 A.2 Single-level Convex Proofs

550 A.2.1 Proof of Theorem 1

551 The proof of Theorem 1 is similar to the standard proof of decaying-step SGD (GD) that can be found
552 in e.g. Beck [1]. Here, we give the proof for completeness.

$$\begin{aligned} \|x^{k+1} - x^*\|^2 &= \|x^k - x^*\|^2 - 2\gamma_k \langle x^k - x^*, \nabla f_i(x^k) \rangle + \gamma_k^2 \|\nabla f_i(x^k)\|^2 \\ &\stackrel{(a)}{\leq} \|x^k - x^*\|^2 - 2\gamma_k (f_i(x^k) - f_i(x^*)) + \gamma_k^2 \|\nabla f_i(x^k)\|^2 \\ &\stackrel{(b)}{\leq} \|x^k - x^*\|^2 - 2\gamma_k (f_i(x^k) - f_i(x^*)) + \gamma_{b,k}^2 G^2, \end{aligned}$$

553 where (a) is by individual convexity of f_i , and (b) is by Assumption 2. Take conditional expectation
554 and assume that γ_k is independent of sample k

$$\begin{aligned} \mathbb{E}[\|x^{k+1} - x^*\|^2 | x^k] &\stackrel{(c)}{\leq} \|x^k - x^*\|^2 - 2E[\gamma_k | x^k] (f(x^k) - f(x^*)) + \gamma_{b,k}^2 G^2 \\ &\stackrel{(d)}{\leq} \|x^k - x^*\|^2 - 2\gamma_{l,k} (f(x^k) - f(x^*)) + \gamma_{b,k}^2 G^2, \end{aligned}$$

555 where (c) is by independence of γ_k and x^k , and (d) is because $\gamma_{l,k} \leq \gamma_k, \forall k$. Take total expectation
556 and rearrange

$$2\gamma_{l,k} \mathbb{E}[f(x^k) - f(x^*)] \leq \mathbb{E}[\|x^k - x^*\|^2] - E[\|x^{k+1} - x^*\|^2] + \gamma_{b,k}^2 G^2$$

557 Using the fact that $\gamma_{l,K-1} \leq \gamma_{l,k} \quad \forall k \in [K-1]$ and set $\gamma_{b,k} = \frac{\gamma_0}{\sqrt{k+1}}$, we obtain

$$2\gamma_{l,K-1} \mathbb{E}[f(x^k) - f(x^*)] \leq \mathbb{E}[\|x^k - x^*\|^2] - E[\|x^{k+1} - x^*\|^2] + \frac{\gamma_0^2}{k+1} G^2$$

558 Summing over $k = 0$ to $k = K-1$ we obtain

$$\begin{aligned} 2\gamma_{l,K-1} \frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E}[f(x^k) - f(x^*)] &\leq \frac{\|x^0 - x^*\|^2 - \mathbb{E}[\|x^K - x^*\|^2]}{K} + \gamma_0^2 G^2 \frac{1}{K} \sum_{k=0}^{K-1} \frac{1}{k+1} \\ &\leq \frac{\|x^0 - x^*\|^2}{K} + \frac{\gamma_0^2 G^2 \log(K)}{K} \end{aligned}$$

559 Define $\bar{x} = \frac{1}{K} \sum_{k=0}^{K-1} x^k$, apply Jensen's inequality and rearrange

$$\mathbb{E}[f(\bar{x}^K) - f(x^*)] \leq \frac{\|x^0 - x^*\|^2}{2\gamma_{l,K-1}K} + \frac{\gamma_0^2 G^2 \log(K)}{2\gamma_{l,K-1}K}.$$

560 A.3 Proof of Theorem 2

561 The approach of Theorem 2 is similar to [16, Theorem 3.2]. The crucial difference is that the step
562 size in [16, Theorem 3.2] is constant ($\gamma_k = \frac{1}{L_{\max}}$) for $k \leq 4\lceil \kappa \rceil$, whereas for envelope-type step size
563 it is of the form:

$$\gamma_k = \min\{\max\{\gamma_{l,k}, \tilde{\gamma}_k\}, \gamma_{b,k}\}, \quad \gamma_{l,k} = \min\{w, \gamma_{b,k}\},$$

564 where $\tilde{\gamma}_k$ can be (e.g. in the case of SPSB):

$$\tilde{\gamma}_k = \min\left\{\frac{f_{i_k}(x^k) - l_{i_k}^*}{c_k \|\nabla f_{i_k}(x^k)\|^2}, \frac{\gamma_{b,0}}{k+1}\right\}, \quad w = \frac{1}{2c_k L_{\max}}.$$

565 The proof of Theorem 2 suggests that the step size can be either $\frac{f_{i_k}(x^k) - l_{i_k}^*}{c_k \|\nabla f_{i_k}(x^k)\|^2}$ or $\frac{\gamma_{b,0}}{k+1}$ depending on
566 their magnitudes for $k \leq k_0 - 1$ ($k_0 = \max\{1, \lceil \gamma_0/\omega \rceil - 1\}$). After k_0 iterations, the step size is
567 $\gamma_k = \frac{\gamma_{b,0}}{k+1}$. This finding is numerically confirmed by the experimental results in Section B.

568 To proceed with the proof, we have:

$$\begin{aligned}\|x^{k+1} - x^*\|^2 &= \|x^k - \gamma_k \nabla f_i(x^k) - x^*\|^2 \\ &= \|x^k - x^*\|^2 - 2\gamma_k \langle \nabla f_i(x^k), x^k - x^* \rangle + \gamma_k^2 \|\nabla f_i(x^k)\|^2 \\ &\stackrel{(a)}{\leq} \|x^k - x^*\|^2 - 2\gamma_k \langle \nabla f_i(x^k), x^k - x^* \rangle + \gamma_{b,k}^2 \|\nabla f_i(x^k)\|^2,\end{aligned}$$

569 where (a) is because $\gamma_k \leq \gamma_{b,k}, \forall k$. Take conditional expectations

$$\begin{aligned}\mathbb{E}[\|x^{k+1} - x^*\|^2 | x^k] &\leq \|x^k - x^*\|^2 - 2\mathbb{E}[\gamma_k | x^k] \langle \nabla f(x^k), x^k - x^* \rangle + \gamma_{b,k}^2 \mathbb{E}[\|\nabla f_i(x^k)\|^2 | x^k] \\ &\stackrel{(b)}{\leq} \|x^k - x^*\|^2 - \mu \mathbb{E}[\gamma_k | x^k] \|x^k - x^*\|^2 - 2\mathbb{E}[\gamma_k | x^k] [f(x^k) - f(x^*)] + \gamma_{b,k}^2 G^2 \\ &\stackrel{(c)}{\leq} \|x^k - x^*\|^2 - \mu \gamma_{l,k} \|x^k - x^*\|^2 - 2\gamma_{l,k} [f(x^k) - f(x^*)] + \gamma_{b,k}^2 G^2\end{aligned}$$

570 where (b) is by bounded gradients assumption and strong convexity of f . Take total expectation and
571 rearrange

$$2\mathbb{E}[f(x^k) - f(x^*)] \leq \frac{(1 - \mu \gamma_{l,k}) \mathbb{E}[\|x^k - x^*\|^2] - \mathbb{E}[\|x^{k+1} - x^*\|^2]}{\gamma_{l,k}} + \frac{\gamma_{b,k}^2 G^2}{\gamma_{l,k}}$$

572 Choose $k_0 = \max\{1, \lceil \gamma_0 / \omega \rceil - 1\}$, then for $\forall k$ s.t. $k \geq k_0$, we have $\gamma_{l,k} = \min\{\omega, \gamma_{b,k}\} =$
573 $\gamma_{b,k} = \frac{\gamma_0}{k+1}$, which means $\gamma_k = \frac{\gamma_0}{k+1}$ after k_0 steps. Within the first k_0 steps, the step size is
574 $\gamma_k = \min\{\omega, \gamma_{b,k}\}$. Hence, for $k \geq k_0$ we have

$$2\mathbb{E}[f(x^k) - f(x^*)] \leq \frac{(1 - \mu \gamma_{l,k}) \mathbb{E}[\|x^k - x^*\|^2] - \mathbb{E}[\|x^{k+1} - x^*\|^2]}{\gamma_{l,k}} + \frac{\gamma_0 G^2}{k+1}$$

575 Now, sum from $k = k_0$ to $K - 1$

$$2 \sum_{k=k_0}^{K-1} \mathbb{E}[f(x^k) - f(x^*)] \leq \sum_{k=k_0}^{K-1} \frac{(1 - \mu \gamma_{l,k}) \mathbb{E}[\|x^k - x^*\|^2] - \mathbb{E}[\|x^{k+1} - x^*\|^2]}{\gamma_{l,k}} + \sum_{k=k_0}^{K-1} \frac{\gamma_0 G^2}{k+1} \quad (29)$$

576 For the first term in (29), call it A , we expand it as

$$\begin{aligned}A &= \sum_{k=k_0+1}^{K-1} \mathbb{E}[\|x^k - x^*\|^2] \left(\frac{1}{\gamma_{l,k}} - \frac{1}{\gamma_{l,k-1}} - \mu \right) + \mathbb{E}[\|x_{k_0} - x^*\|^2] \left(\frac{1}{\gamma_{l,k_0}} - \mu \right) - \frac{\mathbb{E}[\|x_K - x^*\|^2]}{\gamma_{l,K-1}} \\ &\leq \sum_{k=k_0+1}^{K-1} \mathbb{E}[\|x^k - x^*\|^2] \left(\frac{k+1}{\gamma_0} - \frac{k}{\gamma_0} - \mu \right) + \mathbb{E}[\|x_{k_0} - x^*\|^2] \left(\frac{k_0+1}{\gamma_0} - \mu \right) \\ &\stackrel{(d)}{\leq} \mathbb{E}[\|x_{k_0} - x^*\|^2] \mu k_0.\end{aligned}$$

577 where (d) is because $\gamma_0 \geq \frac{1}{\mu}$, we have $\frac{k+1}{\gamma_0} - \frac{k}{\gamma_0} - \mu \leq 0, \forall k$ and $\frac{k_0+1}{\gamma_0} - \mu \leq \mu k_0$. For the second
578 term in (29), call it B , we have

$$B = \sum_{k=k_0}^{K-1} \frac{\gamma_0 G^2}{k+1} \leq \gamma_0 G^2 \int_{k_0}^{K-1} \frac{1}{x+1} dx = \gamma_0 G^2 [\log(K) - \log(k_0+1)] \leq \gamma_0 G^2 \log K$$

579 Putting A and B together, we obtain

$$\frac{1}{K - k_0} \sum_{k=k_0}^{K-1} \mathbb{E}[f(x^k) - f(x^*)] \leq \frac{\mathbb{E}[\|x_{k_0} - x^*\|^2] \mu k_0}{2(K - k_0)} + \frac{\gamma_0 G^2 \log(K)}{2(K - k_0)} \quad (30)$$

580 Within the first $k_0 - 1$ iterations, similarly to the above, we have

$$\begin{aligned}\mathbb{E}[\|x^{k+1} - x^*\|] &\leq (1 - \mu \gamma_{l,k}) \mathbb{E}[\|x^k - x^*\|^2] - 2\gamma_{l,k} \mathbb{E}[f(x^k) - f(x^*)] + \gamma_{b,k}^2 G^2 \\ &\leq (1 - \mu \gamma_{l,k}) \mathbb{E}[\|x^k - x^*\|^2] + \gamma_{b,k}^2 G^2.\end{aligned}$$

581 For the first $k \leq k_0 - 1$ iterations, $\gamma_{l,k} = \omega$ where $\omega\mu < 1$; thus, we obtain the following

$$\mathbb{E}[\|x^{k+1} - x^*\|] \leq (1 - \mu\omega)E[\|x^k - x^*\|^2] + \gamma_{b,k}^2 G^2.$$

582 Solve this recursively,

$$\begin{aligned} \mathbb{E}[\|x^{k_0} - x^*\|^2] &\leq (1 - \mu\omega)^{k_0} \|x^0 - x^*\|^2 + \sum_{k=0}^{k_0-1} (1 - \mu\omega)^{k_0-k-1} \frac{\gamma_0^2 G^2}{(k+1)^2} \\ &\leq (1 - \mu\omega)^{k_0} \|x^0 - x^*\|^2 + \sum_{k=0}^{k_0-1} \frac{\gamma_0^2 G^2}{(k+1)^2} \\ &\leq (1 - \mu\omega)^{k_0} \|x^0 - x^*\|^2 + \gamma_0^2 G^2 \int_0^{k_0-1} \frac{1}{(1+x)^2} dx \\ &= (1 - \mu\omega)^{k_0} \|x^0 - x^*\|^2 + \gamma_0^2 G^2 (1 - \frac{1}{k_0}) \\ &\leq (1 - \mu\omega)^{k_0} \|x^0 - x^*\|^2 + \gamma_0^2 G^2. \end{aligned} \quad (31)$$

583 Putting this into (30), we obtain

$$\frac{1}{K - k_0} \sum_{k=k_0}^{K-1} \mathbb{E}[f(x^k) - f(x^*)] \leq \frac{\mu k_0}{2(K - k_0)} \{\exp(-k_0\mu\omega) \|x_0 - x^*\|^2 + \gamma_0^2 G^2\} + \frac{\gamma_0 G^2 \log K}{2(K - k_0)}.$$

584 Define $\bar{x}_K = \frac{1}{K - k_0} \sum_{k=k_0}^{K-1} x^k$, then by Jensen's inequality we have

$$\mathbb{E}[f(\bar{x}_K) - f(x^*)] \leq \frac{\mu k_0}{2(K - k_0)} \{\exp(-k_0\mu\omega) \|x_0 - x^*\|^2 + \gamma_0^2 G^2\} + \frac{\gamma_0 G^2 \log K}{2(K - k_0)},$$

585 where $k_0 = \max\{1, \lceil \gamma_0/\omega \rceil - 1\}$ and $\gamma_0 \geq \frac{1}{\mu}$.

586 A.4 Bi-level Proofs

587 A.4.1 Proof of Theorem 3

588 Start with Lemma 4:

$$\begin{aligned} E[V^{k+1}] &\leq \mathbb{E}[V^k] - \frac{\alpha_{l,k}}{2} \mathbb{E}[\|\nabla F(x^k)\|^2] + (\alpha_{b,k} L_f^2 + 2) \mathbb{E}[\|y^{k+1} - y^*(x^k)\|^2] + \\ &\quad \alpha_{b,k} B^2 + (2L_y^2 \alpha_{b,k}^2 + \frac{L_F \alpha_{b,k}^2}{2}) \tilde{\sigma}_f^2 - \mathbb{E}[\|y^k - y^*(x^k)\|^2]. \end{aligned}$$

589 We substitute the result of Lemma 5 for the expression $\mathbb{E}[\|y^{k+1} - y^*(x^k)\|^2]$,

$$\begin{aligned} E[V^{k+1}] &\leq \mathbb{E}[V^k] - \frac{\alpha_{l,k}}{2} \mathbb{E}[\|\nabla F(x^k)\|^2] + [(\alpha_{b,k} L_f^2 + 2)(1 - \mu_g C_k)^T - 1] \mathbb{E}[\|y^k - y^*(x^k)\|^2] + \\ &\quad 2\alpha_{b,k} \beta_{b,k} T L_f^2 \sigma_g^2 + 4\beta_{b,k} T \sigma_g^2 + \alpha_{b,k} B^2 + (2L_y^2 + \frac{L_F}{2}) \alpha_{b,k}^2 \tilde{\sigma}_f^2 \\ &\stackrel{(a)}{\leq} \mathbb{E}[V^k] - \frac{\alpha_{l,k}}{2} \mathbb{E}[\|\nabla F(x^k)\|^2] + [(\alpha_{b,0} L_f^2 + 2)(1 - \mu_g C_{K-1})^T - 1] \mathbb{E}[\|y^k - y^*(x^k)\|^2] + \\ &\quad 2\alpha_{b,k} \beta_{b,k} T L_f^2 \sigma_g^2 + 4\beta_{b,k} T \sigma_g^2 + \alpha_{b,k} B^2 + (2L_y^2 + \frac{L_F}{2}) \alpha_{b,k}^2 \tilde{\sigma}_f^2 \\ &\stackrel{(b)}{\leq} E[V^k] - \frac{\alpha_{l,k}}{2} E[\|\nabla F(x^k)\|^2] + 2\alpha_{b,k} \beta_{b,k} T L_f^2 \sigma_g^2 + 4\beta_{b,k} T \sigma_g^2 + \alpha_{b,k} B^2 + (2L_y^2 + \frac{L_F}{2}) \alpha_{b,k}^2 \tilde{\sigma}_f^2, \end{aligned}$$

590 where (a) is by $\alpha_{b,0} \geq \alpha_{b,k}$, $C_{K-1} \leq C_k, \forall k$, hence $(\alpha_{b,0} L_f^2)(1 - \mu_g C_{K-1})^T \geq (\alpha_{b,k} L_f^2)(1 -$

591 $\mu_g C_k)^T$ (recall that $C_k = \min\{\frac{1}{2pL_g}, \beta_{b,k}\}$ in Lemma 1); (b) is by $T \geq \frac{\log[\alpha_{b,0} L_f^2 + 2]}{-\log(1 - \mu_g C_{K-1})}$, which

592 implies that $(\alpha_{b,0} L_f^2 + 2)(1 - \mu_g C_{K-1})^T \leq 1$. Now, rearrange and use the fact that $\alpha_{l,K-1} \leq \alpha_{l,k}$,

$$\alpha_{l,K-1} \mathbb{E}[\|\nabla F(x^k)\|^2] \leq 2\mathbb{E}[V^k] - 2\mathbb{E}[V^{k+1}] + 4\alpha_{b,k} \beta_{b,k} T L_f^2 \sigma_g^2 + 8\beta_{b,k} T \sigma_g^2 + 2\alpha_{b,k} B^2 + (4L_y^2 + L_F) \alpha_{b,k}^2 \tilde{\sigma}_f^2.$$

593 Then sum over $k = 0$ to $K - 1$:

$$\begin{aligned}
\frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E}[\|\nabla F(x^k)\|^2] &\leq \frac{2V^0}{\alpha_{l,K-1}K} + \frac{4TL_f^2\sigma_g^2}{\alpha_{l,K-1}K} \sum_{k=0}^{K-1} \alpha_{b,k}\beta_{b,k} + \frac{8T\sigma_g^2}{\alpha_{l,K-1}K} \sum_{k=0}^{K-1} \beta_{b,k} + \frac{2B^2}{\alpha_{l,K-1}K} \sum_{k=0}^{K-1} \alpha_{b,k} + \\
&\quad \frac{(4L_y^2 + L_F)\tilde{\sigma}_f^2}{\alpha_{l,K-1}K} \sum_{k=0}^{K-1} \alpha_{b,k}^2 \\
&\stackrel{(c)}{=} \frac{2V^0}{\alpha_{l,K-1}K} + \frac{4TL_f^2\sigma_g^2\alpha_{b,0}\beta_{b,0}}{\alpha_{l,K-1}K} \sum_{k=0}^{K-1} \frac{1}{(k+1)^{1.5}} + \frac{8T\sigma_g^2\beta_{b,0}}{\alpha_{l,K-1}K} \sum_{k=0}^{K-1} \frac{1}{k+1} + \\
&\quad + \frac{2B^2\alpha_{b,0}}{\alpha_{l,K-1}K} \sum_{k=0}^{K-1} \frac{1}{\sqrt{k+1}} + \frac{(4L_y^2 + L_F)\tilde{\sigma}_f^2\alpha_{b,0}^2}{\alpha_{l,K-1}K} \sum_{k=0}^{K-1} \frac{1}{k+1} \\
&\stackrel{(d)}{\leq} \frac{2V^0}{\alpha_{l,0}\sqrt{K}} + \frac{8TL_f^2\sigma_g^2\alpha_{b,0}\beta_{b,0}}{\alpha_{l,0}\sqrt{K}} + \frac{8T\sigma_g^2\beta_{b,0}\log(K)}{\alpha_{l,0}\sqrt{K}} + \\
&\quad \frac{2B^2\alpha_{b,0}}{\alpha_{l,0}} + \frac{(4L_y^2 + L_F)\tilde{\sigma}_f^2\alpha_{b,0}^2\log(K)}{\alpha_{l,0}\sqrt{K}},
\end{aligned}$$

594 where we substituted $\alpha_{b,k} = \frac{\alpha_{b,0}}{\sqrt{k+1}}$ and $\beta_{b,k} = \frac{\beta_{b,0}}{k+1}$ in (c); (d) is based on $\sum_{k=0}^{K-1} \frac{1}{(k+1)^{1.5}} \leq 2$,
595 $\sum_{k=0}^{K-1} \frac{1}{k+1} \leq \log(K)$, and $\sum_{k=0}^{K-1} \frac{1}{\sqrt{k+1}} \leq \sqrt{K}$. Recall that in Lemma 2, we have $L_f \sim \mathcal{O}(\kappa^2)$,
596 $L_y \sim \mathcal{O}(\kappa)$, and $L_F \sim \mathcal{O}(\kappa^3)$. Also recall that $\alpha_{b,k} = \frac{\alpha_{b,0}}{\sqrt{k+1}}$ and $\alpha_{l,k} = \frac{\alpha_{l,0}}{\sqrt{k+1}}$, hence
597 we choose $\alpha_{l,0} \sim \alpha_{b,0} \sim \mathcal{O}(\kappa^{-3})$, $T \sim \mathcal{O}(\kappa)$, and $\beta_{b,0} \sim \mathcal{O}(\kappa^{-2})$. Then we can obtain
598 $\frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E}[\|\nabla F(x^k)\|^2] \leq \mathcal{O}(\frac{\kappa^3}{\sqrt{K}} + \frac{\kappa^2 \log(K)}{\sqrt{K}})$.

599 A.4.2 Theorem 4 and its proof

600 **Theorem 4.** Suppose f and g satisfy Assumptions 3, 4, and 5, and, learning-rate upper bounds $\alpha_{b,k} =$
601 $\frac{\alpha_{b,0}}{\sqrt{k+1}}$ and $\alpha_{l,k} = \frac{\alpha_{l,0}}{\sqrt{k+1}}$ with $\alpha_{b,0}$ and $\alpha_{l,0}$ satisfying $\frac{1}{L_F + 4L_y^2} \geq \frac{\alpha_{b,0}^2}{\alpha_{l,0}}$ and $\alpha_{l,0} \leq \alpha_{b,0}$. Further
602 assume that α_k is independent of the stochastic hypergradient h_f^k . Then, under the Bounded-variance
603 assumption in (16) with $p \geq \max\{\frac{\mu_g}{\mu_g + L_g}, \frac{\mu_g + L_g}{4L_g}\}$, $\beta_{b,0} \leq \min\{\frac{2}{\mu_g + L_g}, \frac{\mu_g + L_g}{2\mu_g L_g}, \frac{1}{2pL_g - \frac{2\mu_g L_g}{\mu_g + L_g}}\}$, and
604 $T \geq \frac{\log(\frac{3}{2}\alpha_{b,0}L_f^2 + 2)}{\min\{-\log(1 - \frac{2\mu_g L_g}{\mu_g + L_g}\beta_{b,K-1}), -\log((2pL_g - \frac{2\mu_g L_g}{\mu_g + L_g})\beta_{b,0})\}}$, BiSPS achieves the following rate:

$$\frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E}[\|\nabla F(x^k)\|^2] \leq \tilde{\mathcal{O}}(\frac{\kappa^3}{\sqrt{K}} + \frac{\kappa^2 \log K}{\sqrt{K}}).$$

605 *Proof.* Start with Lemma 4:

$$\begin{aligned}
E[V^{k+1}] &\leq \mathbb{E}[V^k] - \frac{\alpha_{l,k}}{2} \mathbb{E}[\|\nabla F(x^k)\|^2] + (\alpha_{b,k}L_f^2 + 2)\mathbb{E}[\|y^{k+1} - y^*(x^k)\|^2] + \\
&\quad \alpha_{b,k}B^2 + (2L_y^2\alpha_{b,k}^2 + \frac{L_F\alpha_{b,k}^2}{2})\tilde{\sigma}_f^2 - \mathbb{E}[\|y^k - y^*(x^k)\|^2].
\end{aligned}$$

606 We substitute the result of Lemma 6 for the expression $\mathbb{E}[\|y^{k+1} - y^*(x^k)\|^2]$,

$$\begin{aligned}
E[V^{k+1}] &\leq \mathbb{E}[V^k] - \frac{\alpha_{l,k}}{2} \mathbb{E}[\|\nabla F(x^k)\|^2] + [(\alpha_{b,k}L_f^2 + 2)(\frac{\gamma_{b,k}}{C_k} - \frac{2\mu_g L_g}{\mu_g + L_g}\beta_{b,k})^T - 1]\mathbb{E}[\|y^k - y^*(x^k)\|^2] + \\
&\quad + L_f^2 T \sigma_g^2 \alpha_{b,k} \beta_{b,k}^2 + 2T \sigma_g^2 \beta_{b,k}^2 + \alpha_{b,k} B^2 + [2L_y^2 + \frac{L_F}{2}]\alpha_{b,k}^2 \tilde{\sigma}_f^2 \\
&\stackrel{(a)}{\leq} E[V^k] - \frac{\alpha_{l,k}}{2} \mathbb{E}[\|\nabla F(x^k)\|^2] + L_f^2 T \sigma_g^2 \alpha_{b,k} \beta_{b,k}^2 + 2T \sigma_g^2 \beta_{b,k}^2 + \alpha_{b,k} B^2 + [2L_y^2 + \frac{L_F}{2}]\alpha_{b,k}^2 \tilde{\sigma}_f^2,
\end{aligned}$$

607 where in (a) we have chosen $T \geq \max\{\frac{\log(\alpha_{b,0}L_f^2+2)}{-\log(1-\frac{2\mu_gL_g}{\mu_g+L_g}\beta_{b,K-1})}, \frac{\log(\alpha_{b,0}L_f^2+2)}{-\log(2pL_g-\frac{2\mu_gL_g}{\mu_g+L_g})\beta_{b,0}}\}$. This ensures
 608 that $T \geq \frac{\log(\alpha_{b,0}L_f^2+2)}{-\log(\frac{\beta_{b,k}}{C_k}-\frac{2\mu_gL_g}{\mu_g+L_g}\beta_{b,k})}, \forall k$, which guarantees that $(\alpha_{b,k}L_f^2+2)(\frac{\beta_{b,k}}{C_k}-\frac{2\mu_gL_g}{\mu_g+L_g}\beta_{b,k})^T - 1 \leq 0$.
 609 Now, rearrange and use the fact that $\alpha_{l,K-1} \leq \alpha_{l,k}$,

$$\alpha_{l,K-1}\mathbb{E}[\|\nabla F(x^k)\|^2] \leq 2\mathbb{E}[V^k] - 2\mathbb{E}[V^{k+1}] + 2\alpha_{b,k}\beta_{b,k}^2TL_f^2\sigma_g^2 + 4\beta_{b,k}^2T\sigma_g^2 + 2\alpha_{b,k}B^2 + (4L_y^2 + L_F)\alpha_{b,k}^2\tilde{\sigma}_f^2.$$

610 Sum over $k = 0$ to $k = K - 1$:

$$\begin{aligned} \frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E}[\|\nabla F(x^k)\|^2] &\leq \frac{2V^0}{\alpha_{l,K-1}K} + \frac{2TL_f^2\sigma_g^2}{\alpha_{l,K-1}K} \sum_{k=0}^{K-1} \alpha_{b,k}\beta_{b,k}^2 + \frac{4T\sigma_g^2}{\alpha_{l,K-1}K} \sum_{k=0}^{K-1} \beta_{b,k}^2 + \\ &\quad \frac{2B^2}{\alpha_{l,K-1}K} \sum_{k=0}^{K-1} \alpha_{b,k} + \frac{(4L_y^2 + L_F)\tilde{\sigma}_f^2}{\alpha_{l,K-1}K} \sum_{k=0}^{K-1} \alpha_{b,k}^2 \\ &\stackrel{(b)}{=} \frac{2V^0}{\alpha_{l,K-1}K} + \frac{2TL_f^2\sigma_g^2\alpha_{b,0}\beta_{b,0}^2}{\alpha_{l,K-1}K} \sum_{k=0}^{K-1} \frac{1}{(k+1)^{3/2}} + \frac{4T\sigma_g^2\beta_{b,0}^2}{\alpha_{l,K-1}K} \sum_{k=0}^{K-1} \frac{1}{k+1} + \\ &\quad \frac{2B^2\alpha_{b,0}}{\alpha_{l,K-1}K} \sum_{k=0}^{K-1} \frac{1}{(k+1)^{0.5}} + \frac{(4L_y^2 + L_F)\tilde{\sigma}_f^2\alpha_{b,0}^2}{\alpha_{l,K-1}K} \sum_{k=0}^{K-1} \frac{1}{k+1} \\ &\stackrel{(c)}{\leq} \frac{2V^0}{\alpha_{l,K-1}K} + \frac{2TL_f^2\sigma_g^2\alpha_{b,0}\beta_{b,0}^2}{\alpha_{b,0}\sqrt{K}} + \frac{4T\sigma_g^2\beta_{b,0}^2\log(K)}{\alpha_{b,0}\sqrt{K}} + \\ &\quad + \frac{2B^2\alpha_{b,0}}{\alpha_{l,0}} + \frac{(4L_y^2 + L_F)\tilde{\sigma}_f^2\alpha_{b,0}^2\log(K)}{\alpha_{l,0}\sqrt{K}}, \end{aligned}$$

611 where in (b) we have substituted $\alpha_{b,k} = \frac{\alpha_{b,0}}{\sqrt{k+1}}$ and $\beta_{b,k} = \frac{\beta_{b,0}}{\sqrt{k+1}}$, and (c) is by $\sum_{k=0}^{K-1} \frac{1}{(k+1)^{3/2}} \leq 2$,
 612 $\sum_{k=0}^{K-1} \frac{1}{k+1} \leq \log(K)$, and $\sum_{k=0}^{K-1} \frac{1}{(k+1)^{1/2}} \leq \sqrt{K}$. Similar to the proof of Theorem 3, we choose
 613 $\alpha_{l,0} \sim \alpha_{b,0} \sim \mathcal{O}(\kappa^{-3})$, $T \sim \mathcal{O}(\kappa)$, and $\beta_{b,0} \sim \mathcal{O}(\kappa^{-1})$ to obtain $\frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E}[\|\nabla F(x^k)\|^2] \leq$
 614 $\mathcal{O}(\frac{\kappa^3}{\sqrt{K}} + \frac{\kappa^2\log(K)}{\sqrt{K}})$. \square

615 B Additional Experiment Results

616 This section is organized as follows. First, we discuss synthetic quadratics experiments. Second,
 617 we provide more details on the sensitivity of the algorithms to the choices of δ in (14), on the reset
 618 procedure, and on the search cost of BiSLS. Third, we compare the empirical performance of 1-sample
 619 vs 2-samples implementations of our algorithms for single-level convex and bi-level optimization.
 620 Some additional results for hyper-representation learning and data distillation experiments are also
 621 presented. We run 5 independent runs for all our experiments.

622 B.1 Synthetic Quadratics

The experiments on quadratic functions are adapted from Loizou et al. [29]. The training objective is as follows:

$$f(x) = \frac{1}{2}(x - x_1^*)^T H_1(x - x_1^*) + \frac{1}{2}(x - x_2^*)^T H_2(x - x_2^*),$$

where H_i ($i = 1, 2$) are positive definite. The optimal solutions x_i^* ($i = 1, 2$) are generated randomly from a standard normal distribution. Specifically, H_i is defined as follows:

$$H_i = O^T \cdot \text{Diag}(\log(1 + \lambda_i)) \cdot O, \quad i = 1, 2,$$

623 where O and λ_i are taken from the spectral decomposition of $P^T P$, and P is generated from the
 624 standard normal distribution. Figure 8 shows the convergence of various algorithms with different
 625 starting points. Interestingly, both SPSB with either 1 sample or 2 samples (1 sample for computing
 626 the gradient and the other for computing the step size) converge to the optimal solution (labelled with
 627 a star).




Figure 8: Iterate trajectories of different starting points for the synthetic quadratic experiments.

B.2 Sensitivity of δ , reset, and search cost

In this section, we discuss the effects of δ in (14), η in Algorithm 2, and comment on the search cost of the options in the Reset Algorithm 2. Recall that the line-search condition (8) assumes that we can find a largest $\gamma_k \leq \gamma_{b,0}$ to satisfy it. However, in practice, we apply a backtracking procedure, i.e. $\gamma_k = \gamma_k * w, 0 < w < 1$, until γ_k satisfies (8). Therefore, the found learning rate is not guaranteed to be the largest. Nonetheless, we assume that γ_k is the largest to simplify our analysis given above (similar arguments apply to line-search at both upper and lower-level in the bi-level optimization). The experiments in this section are based on hyper-representation learning [42]. In this case, the objective of the induced bilevel-optimization problem can be written as:

$$\begin{aligned} \min_w F(w) &= \frac{1}{2D_{X_1}} \|\tilde{f}(X_1; w)c^*(w) - Y_1\|^2 \\ \text{s.t. } c^*(w) &= \operatorname{argmin}_c \frac{1}{2D_{X_2}} \|\tilde{f}(X_2; w)c - Y_2\|^2 + \frac{\lambda}{2} \|c\|^2, \end{aligned}$$

where (X_1, Y_1) and (X_2, Y_2) are validation and training data sets with sizes D_{X_1} and D_{X_2} , respectively; $\tilde{f}(\cdot; w)$ are the embedding layers of the model parameterized by w ; and, c is the classification layer. Moreover, we use conjugate gradient methods (CG) [17, 18] to solve the linear system when computing the hypergradient for hyper-representation learning experiments.

Reset While Algorithm 2 (reset) can be applied to both upper and lower-level problems, we focus our discussions here on the upper-level learning rate (α_k). This is because we empirically find it to be more critical for the convergence performance (see Figure 6a). As shown in Algorithm 2, reset has 3 options. Options 1, 2, and 3 search starting from $\alpha_{b,0}$, α_{k-1} , and $\eta\alpha_{k-1}$, at iteration k respectively. Option 1 has the highest search cost as it always starts from the same initial upper bound ($\alpha_{b,0}$). Option 2 ensures the monotonicity of the learning rate due to $\alpha_k \leq \alpha_{b,k} = \alpha_{k-1}$. Option 3 chooses the search starting point at iteration k ($\alpha_{b,k}$) by multiplying the previous learning rate (α_{k-1}) by a factor $\eta \geq 1$. As in the single-level convex case where monotonicity in the step size can potentially lead to slow convergence (see Figure 3), we again observe that monotonicity in the upper learning rate (i.e. option 2) leads to poorer performance when compared against options 1 or 3 as shown in Figure 9. Finally, we compare the performance of different η s in option 3 (note that $\eta = 1$ in option 3 is equivalent to option 2). We observe in Figure 10 that different η s perform equally well. This shows the robustness of our algorithm to the choice of η . As mentioned previously, the choice of option 3 over option 1 are due to 2 reasons: (a) reduced search cost; (b) provides an overall non-increasing and non-monotonic trend of upper bound $\alpha_{b,k}$. We discuss search cost of different η s in option 3 below.

Search Cost We investigate the line-search cost based on options 1, 2, and 3 in reset (Algorithm 2) for the upper-level problem. For option 1, the search cost in terms of number of search rounds (i.e. number of evaluations using (14)) per iteration for BiSLS-SGD and BiSLS-Adam is 89 ± 15 and 115 ± 16 , respectively. The search cost of BiSLS-Adam is higher than that of BiSLS-SGD because the feasible learning rate range for Adam is typically smaller than SGD at the upper-level. Based on the




Figure 9: Validation loss against iterations with search options 1, 2, and 3 for the upper-level learning rate. The results for BiSLS-SGD and BiSLS-Adam are in (a) and (b), respectively. For the lower-level search, we fix it option 1 with $\beta_{b,0} = 100$. Results are based on hyper-representation learning.

661 results in Figure 11, we observe the use of option 3 in reset can significantly reduce the search cost for
 662 both BiSLS-Adam and BiSLS-SGD. For example, choosing $\eta = 2$ in option 3 results in an average
 663 upper-level search cost of only ~ 9 rounds per iteration, which is much smaller than that of option 1.
 664 As we have shown in Figure 9, options 1 and 3 have nearly the same performance. Therefore, option
 665 3 is an efficient algorithm that maintains good performance while reducing computation cost. Option
 666 2 has the lowest search cost (~ 4 rounds per iteration). However, its performance is not as good as
 667 option 1 or 3 as observed in Figure 9. Moreover, the average lower-level search cost is only 1 round
 668 per iteration when option 1 is used (see Figure 11b).

669 **Sensitivity on δ** As mentioned in Sec 2, due to the stochastic error in hypergradient computation,
 670 further complicated by the approximation error of $y^*(x)$ (see (14)), a learning rate is not guaranteed
 671 to be found in the bi-level case. Specifically, this is in contrast to the single-level convex problems.
 672 To avoid this, we introduce in (14) a δ slack to give some tolerance to such errors. Here, we give a
 673 thorough investigation of the effects of δ on performance. We vary its magnitude across 6 orders
 674 for both reset options 1 and 3 (see Algorithm 2 and discussions on reset above). We observe that
 675 despite a large difference on the magnitudes of δ , they all share very similar performance for both
 676 BiSLS-SGD and BiSLS-Adam: see Figures 12 and 13. We summarize the key findings in this section as
 677 follows: ① **The option 3 in reset has good empirical performance (outperforms option 2) and**
 678 **is an effective way to reduce search cost (Figure 9, 11);** ② **BiSLS is highly robust to different**
 679 **choices of η in option 3 and δ in (14) (Figure 10, 12, 13).**

680 B.3 Data distillation objective and additional results

We let $\mathcal{L}_S(w)$ denote the loss evaluated on dataset S with model weights w . The objective of *data distillation* can be expressed as a BO problem as follows:

$$D^* = \operatorname{argmin}_D \mathcal{L}_{\tilde{V}}(w^*(D)) \quad \text{s.t.} \quad w^*(D) = \operatorname{argmin}_w \mathcal{L}_D(w),$$

681 where \tilde{V} is of the same size as D and subsampled from the entire (original) dataset V . The solution
 682 D^* is the distilled data, e.g. 9 MNIST digits each corresponding to a different label. In figure 14a,
 683 we show the performance of BiSPS for different values of $\alpha_{b,0}$ in comparison with BiSLS-SGD, and
 684 observe that BiSLS-SGD has better performance. In 14b, we show the results when we increase
 685 the number of lower-level iterations (T) from 20 to 50. As observed for $T = 20$ (in Figure 7),
 686 BiSLS-SGD here also outperforms a fine-tuned Adam or SGD.

687 B.4 1-sample or 2-samples versions of algorithms for convex and bi-level optimization

688 We provide additional results to compare the performance of 1-sample and 2-samples (one for
 689 computing the gradient and the other for computing the step size) versions of our algorithms for

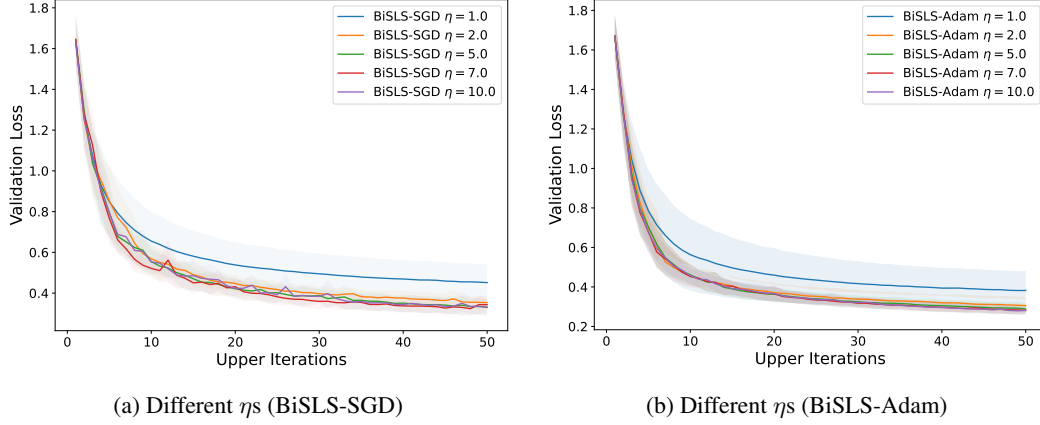


Figure 10: Validation loss against iterations for different η s based on reset option 3. Results for BiSLS-SGD are given in (a) and for BiSLS-Adam are given in (b). Note that $\eta = 1$ in reset option 3 is equivalent to reset option 2. For the lower-level search, we fix it option 1 with $\beta_{b,0} = 100$. Results are based on hyper-representation learning.




Figure 11: (a) Upper-level search cost measured in terms of average number of search rounds per iteration against different η s in reset (Algorithm 2). (b) Lower-level search cost measured in the same way as upper-level against different lower-level search starting points ($\beta_{b,0}$). The lower-level search is done with option 1 (see above for discussions about these options).

SPSB and BiSPS used for single-level and bi-level optimization, respectively. In the single-level case (Figure 15), we observe that 2-samples SPSB performs just as well as 1-sample SPSB. Interestingly, we observe that their step sizes also follow a similar pattern. That is: an initial increase followed by a regime where $\gamma_k = \frac{f_{i_k}(x^k) - l_{i_k}^*}{c \|\nabla f_{i_k}(x^k)\|^2}$ is frequently used, and eventually changes to decaying-step SGD. This seems to also match with Theorem 2 where a transition point for SPSB ($k_0 = \max\{1, \lceil \gamma_0 / \omega \rceil - 1\}$, $w = \frac{1}{2cL_{\max}}$, $\gamma_0 \geq \frac{1}{\mu}$) is predicted. At the same time, we also note that (perhaps, unsurprisingly) the 2-samples version seems to have a slightly more oscillatory behavior than the 1-sample version as shown in Figure 15. SLSB with either 1-sample or 2-samples also result in a similar performance and step size. Overall, despite the requirements of Theorems 1 and 2 for a 2-samples assumption, the empirical performance of 1-sample and 2-samples for either SPSB or SLSB appears to be very similar. Moving on to the bi-level case, recall that Theorems 3 and 4 require the 2-samples assumption (i.e., α_k independent of h_f^k) for the upper-level learning rate. We empirically verify this assumption with both hyper-representation learning and data distillation experiments. For hyper-representation learning experiments in Figure 16, BiSPS with either 1-sample or 2-samples for different values of $\alpha_{b,0}$ show similar performance. In fact, for $\alpha_{b,0} = 0.1$ we even observe that the 2-samples variant outperforms the 1-sample BiSPS. For data distillation experiments

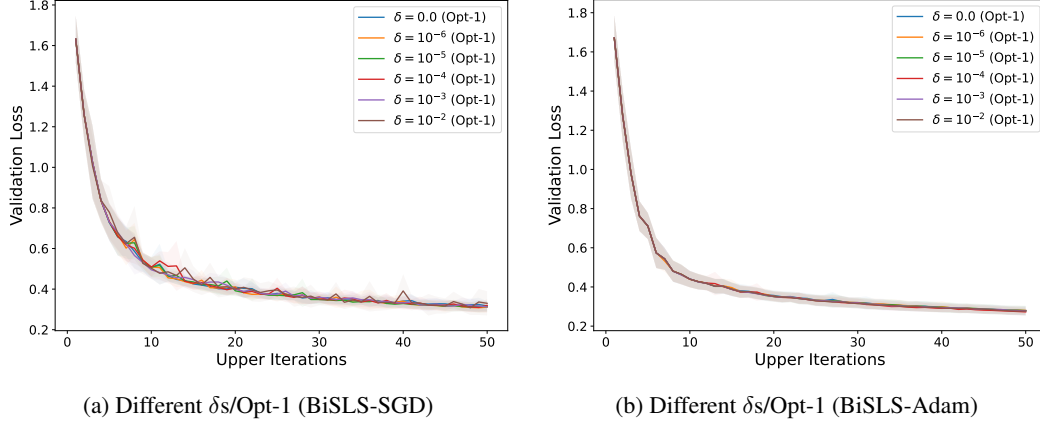


Figure 12: Validation loss against iterations for different δ s based on reset option 1. Results for BiSLS-SGD (a) and for BiSLS-Adam (b). For the lower-level search, we fix it option 1 with $\beta_{b,0} = 100$. Results are based on hyper-representation learning.

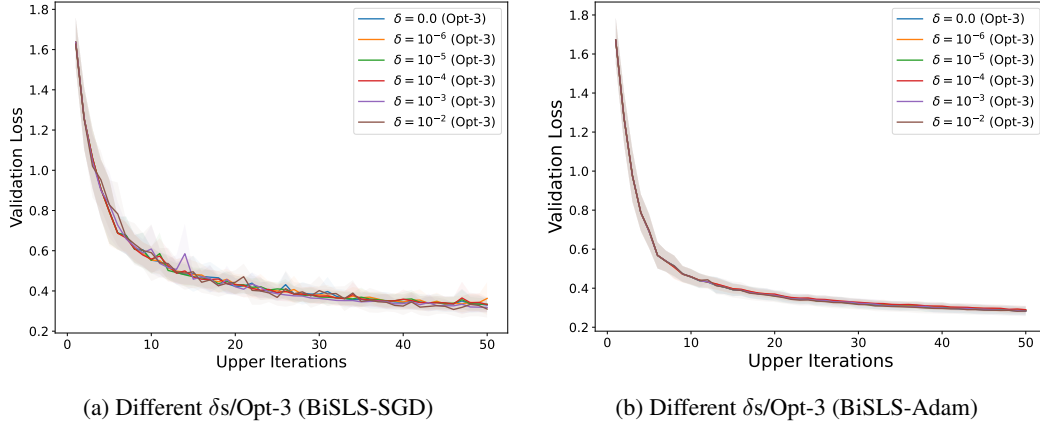


Figure 13: Validation loss against iterations for different δ s based on reset option 3 ($\eta = 10$). Results for BiSLS-SGD (a) and for BiSLS-Adam (b). For the lower-level search, we fix it option 1 with $\beta_{b,0} = 100$. Results are based on hyper-representation learning.

706 in Figure 17, the performances of 1-sample and 2-samples BiSPS are similar to each other when
 707 $\alpha_{b,0} = 10.0$ or $\alpha_{b,0} = 50.0$. In general, the performance difference between 1-sample and 2-samples
 708 in the single-level or bi-level settings is small.

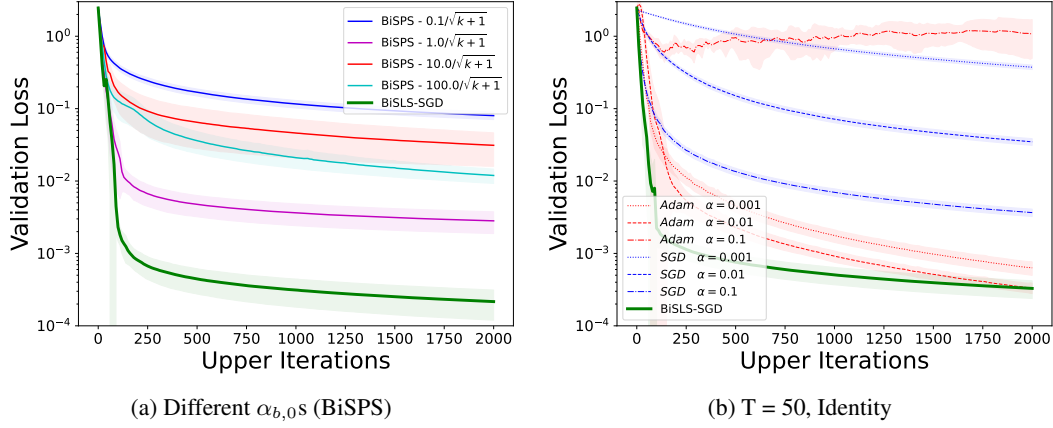


Figure 14: Validation loss against iterations. (a) Comparison between BiSPS with different $\alpha_{b,0}$ s and BiSLS-SGD. (b) Comparison between BiSLS-SGD to fine-tuned Adam/SGD (β_k fixed at 10^{-4}). Inverse Hessian in (2) treated as the Identity [30] when computing the hypergradient. Recall that T is the total number of lower-level iterations and we have shown the results for $T = 20$ in Figure 7b.




Figure 15: Binary linear classification on w8a dataset using logistic loss [3]. Train loss (left) and step size (right) against iterations. We choose $\gamma_{b,0} = 1000$ for all algorithms. The upper bound for either SPSB or SLSB decays as $\gamma_{b,k} = \frac{\gamma_{b,0}}{\sqrt{k+1}}$. For decaying-step SGD, the learning rate schedule is $\frac{\gamma_{b,0}}{\sqrt{k+1}}$.




Figure 16: Comparison between BiSPS (2-samples), BiSPS (1-sample) and decaying-step SGD. Experiments are based on hyper-representation learning. For either version of BiSPS, the lower-level learning rate (β_k) is fixed at 10. The hypergradient is computed using conjugate gradient [17].




Figure 17: Comparison between BiSPS (2-samples), BiSPS (1-sample) and decaying-step SGD. Experiments are based on data distillation. For either version of BiSPS, the lower-level learning rate (β_k) is fixed at 10^{-4} . The Inverse Hessian in (2) is treated as the Identity when computing the hypergradient [30].