



Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Salehi, MS;Mukherjee, S;Roberts, L;Ehrhardt, MJ

Title:

An Adaptively Inexact First-Order Method for Bilevel Optimization with Application to Hyperparameter Learning

Date:

2025-01-01

Citation:

Salehi, M. S., Mukherjee, S., Roberts, L. & Ehrhardt, M. J. (2025). An Adaptively Inexact First-Order Method for Bilevel Optimization with Application to Hyperparameter Learning. *SIAM Journal on Mathematics of Data Science*, 7 (3), pp.906-936. <https://doi.org/10.1137/24M1653513>.

Persistent Link:

<https://hdl.handle.net/11343/362888>

License:

[CC BY](#)

An adaptively inexact first-order method for bilevel optimization with application to hyperparameter learning *

Mohammad Sadegh Salehi[†], Subhadip Mukherjee[‡], Lindon Roberts[§], and Matthias J. Ehrhardt[¶]

Abstract. Various tasks in data science are modeled utilizing the variational regularization approach, where manually selecting regularization parameters presents a challenge. The difficulty gets exacerbated when employing regularizers involving a large number of hyperparameters. To overcome this challenge, bilevel learning can be employed to learn such parameters from data. However, neither exact function values nor exact gradients with respect to the hyperparameters are attainable, necessitating methods that only rely on inexact evaluation of such quantities. State-of-the-art inexact gradient-based methods a priori select a sequence of the required accuracies and cannot identify an appropriate step size since the Lipschitz constant of the hypergradient is unknown. In this work, we propose an algorithm with backtracking line search that only relies on inexact function evaluations and hypergradients and show convergence to a stationary point. Furthermore, the proposed algorithm determines the required accuracy dynamically rather than manually selected before running it. Our numerical experiments demonstrate the efficiency and feasibility of our approach for hyperparameter estimation on a range of relevant problems in imaging and data science such as total variation and field of experts denoising and multinomial logistic regression. Particularly, the results show that the algorithm is robust to its own hyperparameters such as the initial accuracies and step size.

Key words. Bilevel Optimization, Bilevel Learning, Hyperparameter Optimization, Backtracking Line Search, Machine Learning.

MSC codes. 65K10, 90C25, 90C26, 90C06, 90C31, 94A08

1. Introduction. In this work, we consider bilevel learning, where the hyperparameter learning problem is formulated as a bilevel optimization problem [17, 47] where our goal is to solve

$$(1.1a) \quad \min_{\theta \in \mathbb{R}^d} \left\{ f(\theta) := \frac{1}{m} \sum_{i=1}^m g_i(\hat{x}_i(\theta)) + r(\theta) \right\}$$

$$(1.1b) \quad s.t. \quad \hat{x}_i(\theta) := \arg \min_{x \in \mathbb{R}^n} h_i(x, \theta), \quad i = 1, \dots, m.$$

The upper-level functions g_i are a measure of quality for the solution of the lower-level problem (1.1b). For example, in supervised learning, they may take the form $g_i(x) = \|x - x_i^*\|^2$

*Submitted to the editors DATE.

Funding: The work of Mohammad Sadegh Salehi was supported by a scholarship from the EPSRC Centre for Doctoral Training in Statistical Applied Mathematics at Bath (SAMBa), under the project EP/S022945/1. Matthias J. Ehrhardt acknowledges support from the EPSRC (EP/S026045/1, EP/T026693/1, EP/V026259/1). Lindon Roberts is supported by the Australian Research Council Discovery Early Career Award DE240100006.

[†]Department of Mathematical Sciences, University of Bath, Bath, BA2 7AY, UK (mss226@bath.ac.uk).

[‡]Department of Electronics & Electrical Communication Engineering, Indian Institute of Technology (IIT), Kharagpur, India (smukherjee@ece.iitkgp.ac.in).

[§]School of Mathematics and Statistics, University of Sydney, Camperdown NSW 2006, Australia (lindon.roberts@sydney.edu.au).

[¶]Department of Mathematical Sciences, University of Bath, Bath, BA2 7AY, UK (m.ehrhardt@bath.ac.uk).

where x_i^* is the desired solution of the lower-level problem.

Problems of the form (1.1) arise in many areas of data science, for example, when the task at hand is modeled using variational regularization approaches, which are commonly used in the fields of image reconstruction, image processing and machine vision. For instance, one can learn the weights of the regularizer in regression [45], or the parameters of the regularizer [33, 16, 44] for tasks such as image denoising, deblurring, inpainting, segmentation, and single-image super-resolution. Moreover, sophisticated data-adaptive regularizers can have millions of parameters; examples include input-convex neural networks [40, 1] and convex-ridge regularizer neural networks [27]. However, currently, neither of these data-adaptive regularizers are learned by bilevel learning due to computational difficulties. The relevance of bilevel learning in variational regularization methods is not restricted to regularizers. It can also be employed to learn the parameters in analytical deep priors [19, 2], and the sampling operator in magnetic resonance imaging (MRI) [49] and seismic imaging [20]. These tasks can also involve a large number of parameters; for example, the authors in [49] demonstrated the use of more than a million parameters. Furthermore, based on the theoretical approach introduced in [9], bilevel learning can be employed to learn the optimal structure of the regularizer.

To address the bilevel learning problem (1.1), one can employ model-free approaches like grid search and random search [6] when dealing with a limited set of hyperparameters and a small, bounded hyperparameter space. Moving beyond the capabilities of model-free approaches, model-based approaches such as Bayesian methods [24] and derivative-free optimization [13, 12, 21] prove to be efficient in solving (1.1) when a small number of hyperparameters are involved. However, here we are interested in computationally scalable algorithms that can solve bilevel learning problems potentially with millions of parameters, thus, in this work, we consider gradient-based methods. In the context of the gradient-based approach, firstly, a first-order or quasi-Newton algorithm [16, 38] is employed to solve the lower-level problem (1.1b). Then, gradient descent is applied to the upper-level problem (1.1a) after computing the gradient $\nabla f(\theta)$ (also referred to as the *hypergradient*) with respect to parameters θ .

Under certain regularity assumptions, when the problem is sufficiently smooth [21, 45], the hypergradient can be calculated by utilizing the implicit function theorem (IFT) [21] or by using automatic differentiation (AD) [39]. Another type of algorithm called piggyback [29] utilizes AD to compute the hypergradient, and they can differ in the required assumptions on the lower-level problem (1.1b). This approach is favorable for nonsmooth lower-level problems. The works [7, 8] provided an extensive analysis of this method, where [8] considered approximating the hypergradient under the usage of different splitting methods, and [7] studied the convergence of piggyback AD when the lower-level problem is a saddle point problem and employed primal-dual algorithms for solving it. As an application, the authors in [10] applied piggyback AD to learn discretization of total variation. Quasi-Newton algorithms [43] are used in [46] for solving the lower-level problem and utilized their approximation of the inverse of the Hessian of the lower-level objective to replace linear solvers like conjugate gradient (CG).

In practice, due to the usage of numerical solvers and the large-scale nature of the problems of interest, computing the exact hypergradient is not feasible [22]; hence, it should be approximated with an accuracy that leads to optimization with inexact gradients in the upper-level problem (1.1a). Computing an approximate hypergradient using the IFT approach requires solving a linear system and solving it up to a tolerance using CG. The first error-bound analysis

of the approximate gradient computed in this approach was done in [45] and improved in [55]. In [26], convergence rates for bilevel problems with approximate hypergradients were studied using a pre-specified number of lower-level iterations. Improved rates were later presented in [32] by fixing the number of lower-level iterations and using warm-start. Numerous other works (see [30, 15] and references therein) discuss convergence rates for solving bilevel problems in the stochastic setting. In [28], an iteration complexity analysis and a priori error bounds for the approximate hypergradient and all sub-problems were introduced by separately considering IFT + CG and AD approaches and reformulating the lower-level problem as a fixed-point problem. Subsequently, the error bound analysis of IFT + CG was further extended by providing computable a priori and a posteriori bounds in [22]. On the other hand, the error-bound analysis of approximating the hypergradient using AD was studied in [39]. Furthermore, the authors in [22] proposed a unified perspective in which AD can be seen as equivalent to IFT+CG.

1.1. Challenges.

Required accuracy for hypergradient approximation. Although much work has been done on different methods of finding the hypergradient and its error bound, convergence theory for hyperparameter optimization with an inexact hypergradient has often assumed the presence of an accurate approximate hypergradient. However, in practice, approximating the hypergradient to a high degree of accuracy results in prohibitively high computational costs. Moreover, empirical evidence, as demonstrated in studies such as [22], indicates that even when the accuracy is not very high or strictly increasing under a priori assumptions like summability, inexact gradient descent on the upper-level problem can still yield progress. As elevating accuracy necessitates additional computations, implementing a dynamic strategy for selecting the required accuracy, allowing it to adjust according to the demands of the optimization, opens the possibility of reducing the overall computational cost and increase the efficiency of the optimization process.

Selecting upper-level step sizes. The convergence of inexact gradient descent on the upper-level problem has primarily been studied under the assumption of a sufficiently small fixed step size [45]. However, due to the unavailability of a closed-form solution for the lower-level problem (1.1b), estimating the Lipschitz constant of the hypergradient and consequently employing a method with the largest possible fixed step size is unrealistic. While line search approaches such as backtracking may seem plausible to address this issue, they typically require evaluating the exact upper-level objective value, which is unattainable since only an approximate lower-level solution is available. Additionally, line search methods like the Armijo rule [43] necessitate the approximated hypergradient to be a descent direction, which is not obvious.

1.2. Existing work and contributions. The IFT+CG approach for approximating the hypergradient, given a decreasing sequence of accuracies, was initially utilized in the hyperparameter optimization approximate gradient (HOAG) algorithm [45]. In this algorithm, convergence of the inexact gradient descent in the gradient mapping was demonstrated under the conditions of having an inexact gradient with a summable error bound and a sufficiently small step size [45]. While HOAG offered improved computational efficiency compared to methods relying on highly accurate approximate hypergradients, the assumption of summability may require more lower-level computations than necessary. Moreover, as discussed regarding the challenge of finding a suitable small enough fixed step size, these assumptions could potentially limit the

practical performance of the method. In addressing the challenge of determining a suitable step size for inexact gradient descent in the upper-level, the author in [45] proposed a heuristic line search in numerical experiments. However, this line search lacked any convergence guarantee. Building on this, an accelerated version with a restarting scheme was later introduced in [53] to improve the performance. Additionally, a bilevel method with adaptive step sizes was presented in [31]. This method, which employs a fixed number of lower-level solver steps, falls within the category of stochastic bilevel methods.

In [51], the authors introduced an IFT-based algorithm that performs only a single step to solve the lower-level problem. The algorithm utilizes CG with high accuracy or analytic exact inversion of the Hessian of the lower-level problem. Another class of single-level methods, which replace the lower-level problem with a so-called *value function* [18], are the fully first-order methods [37, 34, 35]. These methods reformulate the bilevel problem as a constrained optimization problem, eliminating the need for second-order differentiation. They use a fixed number of iterations to approximate the lower-level solution and require careful tuning of step sizes and penalty multipliers. The majority of works in these approaches primarily consider stochastic bilevel problems.

Moreover, some recent work has developed line search methods for single-level problems where the objective function value is computed with errors and the estimated gradient is inexact [50] and potentially stochastic [5]. However, in these methods, the error is bounded but irreducible and not controllable. Additionally, two other stochastic line search methods were proposed in [52] and [25]. The method in [52] was later applied within a stochastic bilevel framework in [23]. Meanwhile, [25] addresses the issue of overly small step sizes in [52] by relaxing the monotonic decrease condition, thus providing a non-monotone stochastic line search method for single-level problems.

An existing adaptive and inexact framework for bilevel learning, known as DFO-LS [21], operates by solving the lower-level problem up to a specified tolerance. Subsequently, leveraging an approximate lower-level solution, the framework utilizes a derivative-free algorithm [12] tailored for solving nonlinear least-squares to update the parameters in the upper-level. This method dynamically determines the necessary accuracy for solving the lower-level problem by considering the trust-region radius of the derivative-free algorithm, ensuring that the upper-level solver can progress effectively. While this approach shows promise in tasks such as determining the parameters of total variation denoising, it faces scalability challenges as the number of parameters increases [17].

In this work, we propose a verifiable backtracking line search scheme in [subsection 3.1](#) that relies solely on inexact function evaluations and the inexact hypergradient. This scheme guarantees sufficient decrease in the exact upper-level objective function. We also establish the required conditions for finding a valid step size using this line search method in [subsection 3.1](#) and prove the convergence of inexact gradient descent in the upper-level objective function when employing our inexact line search in [subsection 3.2](#). As part of our line search scheme, we analyze the sufficient conditions to ensure that the hypergradient is a descent direction in [subsection 3.1](#). In practice, we use the the posteriori error bound for the approximate hypergradient from [22], which will be reviewed in [section 2](#), to verify these conditions in numerical settings.

We present a practical algorithm in [section 3](#) that connects all of our theoretical results,

determining the required accuracy for the inexact hypergradient to be considered a descent direction and the inexact line search (sufficient decrease condition) to be held, adaptively. As a result, our algorithm provides a robust and efficient method for bilevel learning. Furthermore, as the numerical experiments in [section 4](#) demonstrate, our approach outperforms state-of-the-art methods such as HOAG [45] and DFO-LS [21] in solving problems like multinomial logistic regression and variational image denoising, respectively.

1.3. Notation. We denote the parameters as $\theta \in \Theta \subseteq \mathbb{R}^d$, the k -th iterate of parameters as θ_k , and each element of the vector of parameters as $\theta[i]$, where $1 \leq i \leq d$. The notation $\|\cdot\|$ represents the Euclidean norm of vectors and the 2-norm of matrices. The gradient of the lower-level objective function $h_i : \mathbb{R}^n \times \mathbb{R}^d \rightarrow \mathbb{R}$ with respect to x is expressed by $\nabla_x h_i$. Throughout the paper, we denote the minimizer of h_i w.r.t. x for a fixed θ by \hat{x}_i and the approximation of it by \tilde{x}_i . Additionally, the Hessian of h_i with respect to x is represented as $\nabla_x^2 h_i$ and the Jacobian of the lower-level objective h_i is denoted as $\nabla_{x\theta}^2 h_i$. Moreover, $\partial \hat{x}_i : \mathbb{R}^d \rightarrow \mathbb{R}^n \times \mathbb{R}^d$ stands for the derivative of $\theta \mapsto \hat{x}_i(\theta)$ with respect to θ .

2. Background. For ease of notation, we will only consider one data point in (2.1) and the index of the data points i is omitted in the following discussion. Similarly, we do not consider the regularizer r in (1.1a). Thus, our general bilevel learning problem we will study in the main part of the paper takes the form:

$$(2.1a) \quad \min_{\theta \in \mathbb{R}^d} \{f(\theta) := g(\hat{x}(\theta))\}$$

$$(2.1b) \quad s.t. \quad \hat{x}(\theta) := \arg \min_{x \in \mathbb{R}^n} h(x, \theta).$$

Now, we consider the following assumptions.

Assumption 2.1. We make the following assumptions on the lower level loss h :

- I. There exist $\mu(\theta)$, $L(\theta) \in \mathbb{R}$, $0 < \mu(\theta) \leq L(\theta)$ such that $\mu(\theta)I \preceq \nabla_x^2 h(x, \theta) \preceq L(\theta)I$. Moreover, $\nabla_x h$ and $\nabla_x^2 h$ are continuous in θ .
- II. $\nabla_x^2 h(x, \theta)$ is L_H Lipschitz continuous in x and $\nabla_{x\theta}^2 h(x, \theta)$ is L_J Lipschitz in x uniformly for all θ .

Assumption 2.2. The function f is $L_{\nabla f}$ -smooth and g is $L_{\nabla g}$ -smooth, which means f and g are continuously differentiable with $L_{\nabla f}$ and $L_{\nabla g}$ Lipschitz gradients, respectively. Moreover, $L_{\nabla f} > 0$ and $L_{\nabla g} > 0$, and the upper-level loss g is bounded below.

Remark 2.3. [Assumption 2.1 I.](#) implies that the lower-level objective function is strongly convex. [Assumption 2.1 II.](#) will be needed solely in [Theorem 2.5](#), which we utilize in the practical part of our work and not in our analysis. Our assumptions on the upper-level loss g in [Assumption 2.2](#) encompass commonly used loss functions such as the *mean-squared error* (MSE) and *multinomial logistic loss* [41], as well as popular non-convex losses, including the *peak signal-to-noise ratio* (PSNR), which is related to MSE, the *structural similarity index measure* (SSIM) [17], and the *biweight loss* [3, 11]. Our method and theory considers non-convex g , but remain applicable—and can be improved with tighter bounds—for a convex g , with slight modifications to the proofs and the line search conditions we derive.

Remark 2.4. Note that having $\mu(\theta)$ and $L(\theta)$ uniformly bounded in [Assumption 2.1 I.](#) (i.e., there exist $0 < \mu \leq L < \infty$ such that $\mu \leq \mu(\theta) \leq L(\theta) \leq L$) is sufficient to ensure that f is

$L_{\nabla f}$ -smooth, but it is not necessary (e.g., when $\nabla g(x) = 0$). Moreover, the existence of $L_{\nabla f}$ is required for our convergence analysis, but we do not use its value in our algorithm or practical applications.

The gradient of f takes the form below

$$(2.2) \quad \nabla f(\theta) = \nabla(g \circ \hat{x})(\theta) = \partial \hat{x}(\theta)^T \nabla g(\hat{x}(\theta)).$$

For all $x \in \mathbb{R}^n$, under [Assumption 2.1](#) and with the use of the implicit function theorem [4], we can state

$$\partial \hat{x}(\theta)^T := -\nabla_{x\theta}^2 h(x, \theta)^T \nabla_x^2 h(x, \theta)^{-1}.$$

Based on [Assumption 2.1](#) and [Assumption 2.2](#), we have the following useful inequalities, which we will use in later sections. As stated in [42, Lemma 1.2.3], if $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is a $L_{\nabla f}$ -smooth function, then for any $x, y \in \mathbb{R}^d$ we have

$$(2.3) \quad f(y) \leq f(x) + \nabla f(x)^T (y - x) + \frac{L_{\nabla f}}{2} \|y - x\|^2.$$

Note that (see e.g. [42]) if $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is convex and smooth, then for any $x, y \in \mathbb{R}^d$ we have

$$(2.4) \quad f(y) \geq f(x) + \nabla f(x)^T (y - x).$$

At each iteration $k \in \mathbb{N}_0$, as finding $\hat{x}(\theta_k)$ exactly is not feasible due to the usage of numerical solvers, we solve (2.1b) inexactly with tolerance ϵ_k to find $\tilde{x}(\theta_k)$ such that $\|\tilde{x}(\theta_k) - \hat{x}(\theta_k)\| \leq \epsilon_k$ using the a posteriori bound $\|\nabla_x h(\tilde{x}(\theta_k), \theta_k)\| \leq \epsilon_k \mu(\theta_k)$. This a posteriori bound is derived from the strong convexity of the lower-level objective function, as explained in [21]. Also, instead of computing $\nabla_x^2 h(\tilde{x}(\theta_k), \theta_k)^{-1} \nabla g(\tilde{x}(\theta_k))$, we solve the linear system $\nabla_x^2 h(\tilde{x}(\theta_k), \theta_k) q_k = \nabla g(\tilde{x}(\theta_k))$ with residual δ_k . Then, we denote the inexact gradient (hypergradient) of the upper-level problem (2.1a) at each iteration k by $z_k := -\nabla_{x\theta} h(\tilde{x}(\theta_k), \theta_k)^T q_k$ and the corresponding error by $e_k := z_k - \nabla f(\theta_k)$.

The following theorem provides a computable a posteriori bound for the error of the approximate hypergradient, which we utilize in the next section.

Theorem 2.5. (*[22, Theorem 10]*) *Suppose Assumption 2.1 hold. Let us define*

$$c(x(\theta)) := \frac{L_{\nabla g} \|\nabla_{x\theta}^2 h(x(\theta), \theta)\|}{\mu(\theta)} + L_{H^{-1}} \|\nabla g(x(\theta))\| \|\nabla_{x\theta}^2 h(x(\theta), \theta)\| + \frac{L_J \|\nabla g(x(\theta))\|}{\mu(\theta)},$$

where $L_{H^{-1}}$ is the Lipschitz constant of $\nabla_x^2 h(x, \theta)^{-1}$ in x uniformly for all θ .¹ Then, at each iteration $k = 1, 2, \dots$, we have the following a posteriori bound for the inexact gradient:

$$(2.5) \quad \|e_k\| = \|z_k - \nabla f(\theta_k)\| \leq c(\tilde{x}(\theta_k)) \epsilon_k + \frac{\|\nabla_{x\theta}^2 h(\tilde{x}(\theta_k), \theta_k)\|}{\mu(\theta)} \delta_k + \frac{L_J L_{\nabla g}}{\mu(\theta)} \epsilon_k^2.$$

¹From [Assumption 2.1](#), it follows that $\nabla_x^2 h^{-1}$ is Lipschitz continuous.

3. Method of Adaptive Inexact Descent (MAID) and Convergence Analysis. In this section, we propose the *Method of Adaptive Inexact Descent* (MAID) for solving the problem (2.1), stated in Algorithm 3.1. This algorithm uses the update

$$(3.1) \quad \theta_{k+1} = \theta_k - \alpha_k z_k, \quad k = 0, 1, \dots,$$

where z_k represents the inexact hypergradient at iteration k , as computed in Algorithm 3.2. An appropriate step size $\alpha_k > 0$ is computed using backtracking. However, conventional backtracking line search rules, such as the *Armijo rule* [43], require access to an exact function evaluation. This could be available up to machine precision in (2.1) but would be computationally prohibitively expensive. To address this challenge, we introduce a verifiable line search method that generalizes the Armijo rule to the inexact setting utilizes only an approximation of f . We control the accuracies to ensure that $-z_k$ is a descent direction. Our analysis begins with the Armijo rule in the exact setting, where we have access to f and we can utilize the negative of the exact gradient, that is $-\nabla f(\theta_k)$, as the descent direction.

The computation of the approximate hypergradient z_k in Algorithm 3.2 employs the implicit function theorem and numerical solvers, following the same procedure as outlined in the previous section, with tolerances ϵ_k and δ_k . It checks whether the sufficient condition in subsection 3.1 (Proposition 3.1 and Lemma 3.2) is met for $-z_k$ to be a descent direction for f . If not, it shrinks ϵ_k and δ_k , recalculates z_k , and repeats this process until $-z_k$ satisfies the requirement in Proposition 3.1 and Lemma 3.2. With the approximate hypergradient z_k and the tolerance ϵ_k which was used to compute it, Algorithm 3.1 employs the line search rule $\psi(\alpha_k) \leq 0$, where ψ is defined in (3.5), utilizing approximations and bounds for $f(\theta_k)$ and $\nabla f(\theta_k)$. If it fails to find a suitable step size and reaches the maximum number of backtracking iterations, it reduces ϵ_k , recalculates z_k , and retries the backtracking process with an increased maximum number of iterations until it successfully finds an appropriate step size and completes the descent update. This backtracking failure can occur due to insufficient iterations to capture a very small step size or ϵ_k that is not small enough. Note that in the process of increasing the number of backtracking iterations upon failure, as shown in line 3 of Algorithm 3.1, it may appear that the maximum number of backtracking steps can grow indefinitely. However, in our analysis in subsequent sections, particularly in Proposition 3.18, we demonstrate that backtracking failures are bounded, so the maximum number of backtracking steps remains finite.

3.1. Line search with approximate hypergradient and inexact function evaluations. In order to find a suitable step size when we do not have any information about the Lipschitz constant of the gradient, line search methods are usually employed [43]. In Armijo rule, at each iteration k , for a given direction z_k and a starting step size β_k , the task is to find the smallest $i_k \in \mathbb{N}_0$ such that

$$(3.2) \quad f(\theta_k - \beta_k \rho^{i_k} z_k) \leq f(\theta_k) - \zeta \beta_k \rho^{i_k} \nabla f(\theta_k)^T z_k$$

holds for $0 < \rho, \zeta < 1$. After finding i_k , we denote step size by $\alpha_k := \beta_k \rho^{i_k}$. In our setting, we do not have access to the function f nor its gradient; instead, we only have $g(\tilde{x}(\theta_k))$ and an inexact direction z_k . Therefore, we will introduce a line search condition with inexact gradient and function evaluations to find a suitable step size.

Algorithm 3.1 Method of Adaptive Inexact Descent (MAID). Hyperparameters: $\underline{\rho}, \bar{\rho} \in (0, 1)$ control the reduction and increase of the step size α_k , respectively; $\underline{\nu}, \bar{\nu} \in (0, 1)$ govern the reduction and increase of accuracies δ_k and ϵ_k ; $\max_{\text{BT}} \in \mathbb{N}$ is the maximum number of backtracking iterations.

```

1: Input  $\theta_0 \in \mathbb{R}^d$ , accuracies  $\epsilon_0, \delta_0 > 0$ , step size  $\alpha_0 > 0$ .
2: for  $k = 0, 1, \dots$  do
3:   for  $j = \max_{\text{BT}}, \max_{\text{BT}} + 1, \dots$  do
4:      $z_k, \epsilon_k, \delta_k \leftarrow \text{INEXACTGRADIENT}(\theta_k, \epsilon_k, \delta_k)$ 
5:     for  $i = 0, 1, \dots, j - 1$  do
6:       if inexact sufficient decrease  $\psi(\alpha_k) \leq 0$  holds then ▷ Lemma 3.5
7:         Go to line 11 ▷ Backtracking Successful
8:          $\alpha_k \leftarrow \underline{\rho}\alpha_k$  ▷ Adjust the starting step size
9:          $\epsilon_k \leftarrow \underline{\nu}\epsilon_k$  ▷ Backtracking Failed and needs higher accuracy
10:         $\delta_k \leftarrow \underline{\nu}\delta_k$ 
11:        $\theta_{k+1} \leftarrow \theta_k - \alpha_k z_k$  ▷ Gradient descent update
12:        $\epsilon_{k+1} \leftarrow \bar{\nu}\epsilon_k$  ▷ Increasing  $\epsilon_k$ 
13:        $\delta_{k+1} \leftarrow \bar{\nu}\delta_k$  ▷ Increasing  $\delta_k$ 
14:        $\alpha_{k+1} \leftarrow \bar{\rho}\alpha_k$  ▷ Increasing  $\alpha_k$ 

```

Algorithm 3.2 Calculating an inexact hypergradient, ensuring it is a descent direction. Hyperparameters: $\eta \in (0, 1)$ ensures that the computed inexact gradient z is a descent direction; $\underline{\nu} \in (0, 1)$ is used to reduce the accuracies δ and ϵ .

```

1: Input:  $\theta \in \mathbb{R}^d$ , accuracies  $\epsilon, \delta > 0$ .
2: function INEXACTGRADIENT( $\theta, \epsilon, \delta$ )
3:   while True do
4:     Solve lower-level problem to find  $\tilde{x}(\theta)$  such that  $\|\nabla_x h(\tilde{x}(\theta), \theta)\| \leq \epsilon\mu$ .
5:     Solve  $\nabla_x^2 h(\tilde{x}(\theta), \theta)q = \nabla g(\tilde{x}(\theta))$  with tolerance  $\delta$ .
6:     Calculate  $z = -(\nabla_{x\theta}^2 h(\tilde{x}(\theta), \theta))^T q$ .
7:      $\omega \leftarrow$  Calculate upper bound (2.5) for error  $\|e\|$  using Theorem 2.5.
8:     if  $\omega \leq (1 - \eta)\|z\|$  then
9:       return  $z, \epsilon, \delta$ 
10:     $\delta \leftarrow \underline{\nu}\delta$ ,
11:     $\epsilon \leftarrow \underline{\nu}\epsilon$ .

```

We commence by deriving a condition under which the approximate hypergradient is a descent direction. Subsequently, we establish computable upper and lower bounds for the exact upper-level function by leveraging the inexact components. Moving forward, we introduce our line search.

Descent direction. The following proposition provides a sufficient condition to ensure that $-z_k$ is a descent direction for f at θ_k , that is $z_k^T \nabla f(\theta_k) > 0$.

Proposition 3.1. *Let $e_k := z_k - \nabla f(\theta_k)$ be the error in the approximate hypergradient. If*

$\|e_k\| < \|z_k\|$, then $-z_k$ is a descent direction for f at θ_k .

Proof. Since

$$z_k^T \nabla f(\theta_k) = z_k^T (z_k - e_k) = \|z_k\|^2 - z_k^T e_k,$$

$-z_k$ is a descent direction if and only if $z_k^T e_k < \|z_k\|^2$.

Now, from the assumption, by multiplying both sides of the inequality $\|e_k\| < \|z_k\|$ by $\|z_k\| > 0$ we have

$$\|e_k\| \|z_k\| < \|z_k\|^2.$$

Moreover, using Cauchy–Schwarz, we have

$$z_k^T e_k \leq \|z_k\| \|e_k\| < \|z_k\|^2,$$

thus $-z_k$ is a descent direction. ■

Note that the condition $\|e_k\| < \|z_k\|$ in [Proposition 3.1](#) must hold uniformly to demonstrate the convergence of a descent method [[43](#), Theorem 3.2]. Therefore, instead, we will assume the stronger condition $\|e_k\| \leq (1 - \eta)\|z_k\|$ for $0 < \eta < 1$.

In addition, due to the unavailability of $\nabla f(\theta_k)$, we establish a computable bound for $z_k^T \nabla f(\theta_k)$ to facilitate our analysis.

Lemma 3.2. *Let $\|e_k\| \leq (1 - \eta)\|z_k\|$. Then, $\eta\|z_k\|^2 \leq \nabla f(\theta_k)^T z_k$.*

Proof. Combining Cauchy–Schwarz and $\|e_k\| \leq (1 - \eta)\|z_k\|$ yields $e_k^T z_k \leq \|z_k\|^2 - \eta\|z_k\|^2$. Since $\nabla f(\theta_k)^T z_k = \|z_k\|^2 - e_k^T z_k$, we find the desired bound. ■

We now derive computable estimates for $f(\theta_k)$ which will be used to establish our alternative sufficient decrease condition.

Verifiable line search with approximate hypergradient. To obtain an inexact and computable sufficient decrease condition, we first derive computable upper and lower bounds for f in [Lemma 3.3](#). Then, [Lemma 3.5](#) provides a backtracking line search condition for sufficient decrease. We adopt the inequality $\psi(\alpha_k) \leq 0$ from [\(3.5\)](#) as our inexact backtracking line search rule in [Algorithm 3.1](#), where, for $\alpha_k = \underline{\rho}^{i_k} \beta_k$, $i_k \geq 0$ is the smallest integer satisfying the inequality. Note that this is a practical line search as we can compute all of its components.

Lemma 3.3. *Let g be $L_{\nabla g}$ -smooth, with $\hat{x}(\theta_k)$ as defined in [\(2.1b\)](#), and let $\tilde{x}(\theta_k)$ be an approximation of $\hat{x}(\theta_k)$ such that $\|\hat{x}(\theta_k) - \tilde{x}(\theta_k)\| \leq \epsilon_k$. We have the following lower and upper bounds for $f(\theta_k) = g(\hat{x}(\theta_k))$:*

$$(3.3a) \quad g(\hat{x}(\theta_k)) \leq g(\tilde{x}(\theta_k)) + \|\nabla g(\tilde{x}(\theta_k))\| \epsilon_k + \frac{L_{\nabla g}}{2} \epsilon_k^2,$$

$$(3.3b) \quad g(\hat{x}(\theta_k)) \geq g(\tilde{x}(\theta_k)) - \|\nabla g(\tilde{x}(\theta_k))\| \epsilon_k - \frac{L_{\nabla g}}{2} \epsilon_k^2.$$

Proof. Since g is $L_{\nabla g}$ -smooth, using the inequality [\(2.3\)](#) we have

$$g(\hat{x}(\theta_k)) \leq g(\tilde{x}(\theta_k)) + \nabla g(\tilde{x}(\theta_k))^T (\hat{x}(\theta_k) - \tilde{x}(\theta_k)) + \frac{L_{\nabla g}}{2} \|\tilde{x}(\theta_k) - \hat{x}(\theta_k)\|^2.$$

Now, by rearranging and utilizing the Cauchy-Schwarz inequality, we derive

$$g(\hat{x}(\theta_k)) - g(\tilde{x}(\theta_k)) \leq \frac{L_{\nabla g}}{2} \|\tilde{x}(\theta_k) - \hat{x}(\theta_k)\|^2 + \|\nabla g(\tilde{x}(\theta_k))\| \|(\tilde{x}(\theta_k) - \hat{x}(\theta_k))\|.$$

Since $\|\tilde{x}(\theta_k) - \hat{x}(\theta_k)\| \leq \epsilon_k$, the inequality (3.3a) holds.

Furthermore, using (3.3a), we can write

$$g(\hat{x}(\theta_k)) - g(\tilde{x}(\theta_k)) \leq \|\nabla g(\tilde{x}(\theta_k))\| \epsilon_k + \frac{L_{\nabla g}}{2} \epsilon_k^2.$$

Since the right-hand-side of the inequality above is positive, we have

$$g(\hat{x}(\theta_k)) - g(\tilde{x}(\theta_k)) \geq -\|\nabla g(\tilde{x}(\theta_k))\| \epsilon_k - \frac{L_{\nabla g}}{2} \epsilon_k^2,$$

which yields the desired inequality (3.3b). ■

Remark 3.4. Note that if g is $L_{\nabla g}$ -smooth and convex, from (2.4), we can write

$$f(\theta_k) = g(\hat{x}(\theta_k)) \geq g(\tilde{x}(\theta_k)) + \nabla g(\tilde{x}(\theta_k))^T (\hat{x}(\theta_k) - \tilde{x}(\theta_k)).$$

Following steps similar to Lemma 3.3, we obtain

$$(3.4) \quad g(\hat{x}(\theta_k)) \geq g(\tilde{x}(\theta_k)) - \|\nabla g(\tilde{x}(\theta_k))\| \epsilon_k,$$

which provides a tighter bound than (3.3b).

The following lemma employs the bounds (3.3a) and (3.3b) to provide a line search rule such that all of its components are accessible.

Lemma 3.5. *Let $\lambda \in \mathbb{R}$ and suppose that g is $L_{\nabla g}$ -smooth. Denote*

$$\begin{aligned} \bar{U}(x, \epsilon) &:= g(x) + \|\nabla g(x)\| \epsilon + \frac{L_{\nabla g}}{2} \epsilon^2, \\ \underline{U}(x, \epsilon) &:= g(x) - \|\nabla g(x)\| \epsilon - \frac{L_{\nabla g}}{2} \epsilon^2, \end{aligned}$$

and

$$(3.5) \quad \psi(\alpha_k) := \bar{U}(\tilde{x}(\theta_{k+1}), \epsilon_{k+1}) - \underline{U}(\tilde{x}(\theta_k), \epsilon_k) + \lambda \alpha_k \|z_k\|^2.$$

If the backtracking line search condition $\psi(\alpha_k) \leq 0$ is satisfied, then the sufficient descent condition $f(\theta_{k+1}) - f(\theta_k) \leq -\lambda \alpha_k \|z_k\|^2$ holds.

Proof. Using (3.3a) for θ_{k+1} and (3.3b) for θ_k , we have

$$f(\theta_{k+1}) - f(\theta_k) \leq g(\tilde{x}(\theta_{k+1})) + \|\nabla g(\tilde{x}(\theta_{k+1}))\| \epsilon_{k+1} + \frac{L_{\nabla g}}{2} \epsilon_{k+1}^2 - g(\tilde{x}(\theta_k)) + \|\nabla g(\tilde{x}(\theta_k))\| \epsilon_k + \frac{L_{\nabla g}}{2} \epsilon_k^2.$$

The inequality above together with the definition of $\bar{U}(x, \epsilon)$, $\underline{U}(x, \epsilon)$, and $\psi(\alpha_k) \leq 0$ imply $f(\theta_{k+1}) - f(\theta_k) \leq -\lambda \alpha_k \|z_k\|^2$ as required. ■

Remark 3.6. If g is convex, then by using (3.4) instead of (3.3b) in Lemma 3.5 and defining

$$(3.6) \quad \tilde{\psi}(\alpha_k) := \psi(\alpha_k) - \frac{L_{\nabla g}}{2} \epsilon_k^2,$$

the sufficient descent condition $f(\theta_{k+1}) - f(\theta_k) \leq -\lambda \alpha_k \|z_k\|^2$ holds if $\tilde{\psi}(\alpha_k) \leq 0$ is satisfied.

Now, since the components of the line search rule $\psi(\alpha_k) \leq 0$ depend on ϵ_k , we investigate the accuracy ϵ_k for which a step size α_k exists that satisfies $\psi(\alpha_k) \leq 0$ a priori.

Existence of a step size in the line search. In the following discussion, our a priori analysis demonstrates that, given a non-stationary θ_k and an inexact hypergradient z_k serving as a descent direction, there exists a non-empty, non-singular interval of accuracies for which, for each accuracy, a corresponding interval of step sizes satisfies the line search condition $\psi(\alpha_k) \leq 0$ (3.5) introduced in Lemma 3.5. The following auxiliary lemma aids in achieving the aforementioned goals.

Lemma 3.7. *Let $\lambda \in \mathbb{R}$, $\eta \leq 1$, and $\|e_k\| \leq (1 - \eta)\|z_k\|$. Denote*

$$(3.7) \quad w_k := \|\nabla g(\tilde{x}(\theta_k))\| + \|\nabla g(\tilde{x}(\theta_{k+1}))\|,$$

$\bar{\epsilon}_k := \max\{\epsilon_k, \epsilon_{k+1}\}$, and $\phi(\alpha_k) := 2w_k \bar{\epsilon}_k + 2L_{\nabla g} \bar{\epsilon}_k^2 + (\alpha_k^2 \frac{L_{\nabla f}}{2} + (\lambda - \eta)\alpha_k)\|z_k\|^2$. Then,

$$(3.8) \quad \psi(\alpha_k) \leq \phi(\alpha_k),$$

where ψ is as defined in Lemma 3.5.

Proof. Since $f(\theta) = g(\hat{x}(\theta))$ is $L_{\nabla f}$ -smooth, applying (2.3) to $g(\hat{x}(\theta_{k+1})) = g(\hat{x}(\theta_k - \alpha_k z_k))$ in (3.3b) for θ_{k+1} gives

$$\begin{aligned} g(\tilde{x}(\theta_{k+1})) &\leq g(\hat{x}(\theta_{k+1})) + \|\nabla g(\tilde{x}(\theta_{k+1}))\| \epsilon_{k+1} + \frac{L_{\nabla g}}{2} \epsilon_{k+1}^2 \\ &\leq g(\hat{x}(\theta_k)) - \alpha_k \nabla f(\theta_k)^T z_k + \alpha_k^2 \frac{L_{\nabla f}}{2} \|z_k\|^2 + \|\nabla g(\tilde{x}(\theta_{k+1}))\| \epsilon_{k+1} + \frac{L_{\nabla g}}{2} \epsilon_{k+1}^2. \end{aligned}$$

Leveraging the bounds (3.3a) for $g(\hat{x}(\theta_k))$ and Lemma 3.2 for $\nabla f(\theta_k)^T z_k$, we derive

$$\begin{aligned} g(\tilde{x}(\theta_{k+1})) &\leq g(\tilde{x}(\theta_k)) + \|\nabla g(\tilde{x}(\theta_k))\| \epsilon_k + \frac{L_{\nabla g}}{2} \epsilon_k^2 + \left(\alpha_k^2 \frac{L_{\nabla f}}{2} - \alpha_k \eta \right) \|z_k\|^2 \\ &\quad + \|\nabla g(\tilde{x}(\theta_{k+1}))\| \epsilon_{k+1} + \frac{L_{\nabla g}}{2} \epsilon_{k+1}^2. \end{aligned}$$

Adding $\|\nabla g(\tilde{x}(\theta_{k+1}))\| \epsilon_{k+1} + \frac{L_{\nabla g}}{2} \epsilon_{k+1}^2 - g(\tilde{x}(\theta_k)) + \|\nabla g(\tilde{x}(\theta_k))\| \epsilon_k + \frac{L_{\nabla g}}{2} \epsilon_k^2 + \lambda \alpha_k \|z_k\|^2$ to both sides we obtain

$$\psi(\alpha_k) \leq 2\|\nabla g(\tilde{x}(\theta_k))\| \epsilon_k + 2\|\nabla g(\tilde{x}(\theta_{k+1}))\| \epsilon_{k+1} + L_{\nabla g} (\epsilon_k^2 + \epsilon_{k+1}^2) + \left(\alpha_k^2 \frac{L_{\nabla f}}{2} + \alpha_k (\lambda - \eta) \right) \|z_k\|^2.$$

Now, utilizing $\bar{\epsilon}_k = \max\{\epsilon_k, \epsilon_{k+1}\}$, and $w_k = \|\nabla g(\tilde{x}(\theta_k))\| + \|\nabla g(\tilde{x}(\theta_{k+1}))\|$, we derive the desired result. \blacksquare

The lemma below shows the validity of the line search condition in [Lemma 3.5](#) by establishing the existence of a non-empty interval of step sizes that satisfy it, given an interval of valid accuracies specified a priori.

Lemma 3.8. *Assume $\|z_k\| \neq 0$, $\lambda < \eta$, and $\|e_k\| \leq (1 - \eta)\|z_k\|$. Considering w_k and $\bar{\epsilon}$ as defined in [Lemma 3.7](#) and denoting*

$$(3.9) \quad s_k := \frac{1}{2L_{\nabla g}} \left(\sqrt{w_k^2 + \frac{L_{\nabla g}}{L_{\nabla f}} (\eta - \lambda)^2 \|z_k\|^2 - w_k} \right),$$

for all $\bar{\epsilon}_k \in [0, s_k)$, there exist $0 \leq \underline{\alpha}_k < \bar{\alpha}_k$ such that for all $\alpha_k \in [\underline{\alpha}_k, \bar{\alpha}_k]$ the sufficient decrease condition $\psi(\alpha_k) \leq 0$ holds. Moreover, denoting

$$(3.10) \quad \hat{s}_k := \sqrt{(\eta - \lambda)^2 - \frac{4L_{\nabla f}(w_k \bar{\epsilon}_k + L_{\nabla g} \bar{\epsilon}_k^2)}{\|z_k\|^2}} > 0,$$

the interval may be defined by

$$(3.11) \quad \underline{\alpha}_k = \frac{1}{L_{\nabla f}} (\eta - \lambda - \hat{s}_k), \quad \bar{\alpha}_k = \frac{1}{L_{\nabla f}} (\eta - \lambda + \hat{s}_k).$$

Note that $\lim_{\bar{\epsilon}_k \rightarrow 0} \underline{\alpha}_k = 0$, $\lim_{\bar{\epsilon}_k \rightarrow 0} \bar{\alpha}_k = \frac{2(\eta - \lambda)}{L_{\nabla f}} > 0$, and \hat{s}_k is monotonically decreasing in $\bar{\epsilon}_k$.

Proof. From [Lemma 3.7](#) it follows that $\phi(\alpha_k) \leq 0$ implies $\psi(\alpha_k) \leq 0$. Note that $\phi(\alpha_k)$ is quadratic in α_k with positive leading coefficient. The roots of ϕ are given by $\underline{\alpha}_k$ and $\bar{\alpha}_k$ [\(3.11\)](#). These roots are real if

$$(3.12) \quad (\eta - \lambda)^2 \|z_k\|^2 - 4L_{\nabla f}(w_k \bar{\epsilon}_k + L_{\nabla g} \bar{\epsilon}_k^2) \geq 0.$$

Following a similar argument, the inequality [\(3.12\)](#) holds if $\bar{\epsilon}_k \leq s_k$, given that the left-hand side of [\(3.12\)](#) is quadratic in $\bar{\epsilon}_k$ with a negative leading coefficient and roots s_k and $-s_k$. Moreover, $s_k > 0$ by definition, as $\|z_k\| \neq 0$. Since $\bar{\epsilon}_k \geq 0$, we have $\bar{\epsilon}_k \in [0, s_k)$. Referring to [\(3.12\)](#) and the definition of \hat{s}_k , we conclude that $0 < \hat{s}_k \leq \eta - \lambda$, implying $0 \leq \underline{\alpha}_k < \bar{\alpha}_k$. \blacksquare

As s_k [\(3.9\)](#) depends on ϵ_k and ϵ_{k+1} , and thus on $\bar{\epsilon}_k$, we need the following lemmas to show that for $\bar{\epsilon}_k$ sufficiently small, the inclusion $\bar{\epsilon}_k \in [0, s_k)$ holds a priori for any z_k , $\tilde{x}(\theta_k)$, and $\tilde{x}(\theta_{k+1})$ satisfying the conditions of computing the inexact gradient and the solution of the lower-level problem. It shows that if θ_k is non-stationary, there exists a nonempty positive range of accuracies for which we can find a nonempty positive range of step sizes satisfying the line search condition $\psi(\alpha_k) \leq 0$. For simplicity, we set $\delta_k = \bar{\epsilon}_k$ for the remainder of this section, however, the general case holds with a similar argument.

Lemma 3.9. *Let $\|\nabla f(\theta_k)\| > 0$, s_k be as defined in [\(3.9\)](#), and $\lambda \neq \eta$, then $\lim_{\bar{\epsilon}_k \rightarrow 0} s_k > 0$.*

Proof. As $\lim_{\bar{\epsilon}_k \rightarrow 0} \tilde{x}(\theta_k) = \hat{x}(\theta_k)$, $\lim_{\bar{\epsilon}_k \rightarrow 0} \tilde{x}(\theta_{k+1}) = \hat{x}(\theta_{k+1})$, and ∇g is continuous, we have $\lim_{\bar{\epsilon}_k \rightarrow 0} w_k = \bar{w}_k$, for some $\bar{w}_k \geq 0$. Moreover, since $\lim_{\bar{\epsilon}_k \rightarrow 0} z_k = \nabla f(\theta_k)$, $\|\nabla f(\theta_k)\| > 0$, $\lambda \neq \eta$, and $L_{\nabla f} > 0$,

$$\lim_{\bar{\epsilon}_k \rightarrow 0} s_k = \frac{1}{2L_{\nabla g}} \left(\sqrt{\bar{w}_k^2 + \frac{L_{\nabla g}}{L_{\nabla f}} (\eta - \lambda)^2 \|\nabla f(\theta_k)\|^2 - \bar{w}_k} \right) > 0.$$

as desired. \blacksquare

Now, we need the following lemma to show that for $\bar{\epsilon}_k$ small enough, the inclusion $\bar{\epsilon}_k \in [0, s_k)$ holds a priori.

Lemma 3.10. *Let $\{a_j\}_{j=0}^\infty \subseteq [0, \infty)$ and $\{b_j\}_{j=0}^\infty \subseteq [0, \infty)$ be sequences, with $a_j \rightarrow a > 0$ and $b_j \rightarrow 0$, respectively. Then, there exists $J \in \mathbb{N}$ such that for all $j \geq J$, $b_j \in [0, a_j]$.*

Proof. Since $a_j \rightarrow a > 0$, there exists $J_1 \in \mathbb{N}$, such that for all $j \geq J_1$, $a_j \geq \frac{a}{2}$. On the other hand, as $b_j \rightarrow 0$, there exists $J_2 \in \mathbb{N}$, such that for all $j \geq J_2$, $b_j \leq \frac{a}{2}$. Taking $J = \max\{J_1, J_2\}$, we have $\forall j \geq J$, $0 \leq b_j \leq \frac{a}{2} \leq a_j$, which gives us the desired conclusion. \blacksquare

Corollary 3.11. *Let $\|\nabla f(\theta_k)\| > 0$, $\bar{\epsilon}_k \rightarrow 0$, and $0 < \lambda < \eta < 1$. Then, there exists $\epsilon > 0$ such that if $\bar{\epsilon}_k \leq \epsilon$, then $\bar{\epsilon}_k < s_k$.*

That means, if the accuracy is small enough, there exist a range of step sizes satisfying $\psi(\alpha_k) \leq 0$. Note that ϵ is independent of z_k , $\tilde{x}(\theta_k)$, and $\hat{x}(\theta_k)$.

Proof. The assertion follows by setting $b_j = (\bar{\epsilon}_k)_j$ and $a_j = (s_k)_j$ using [Lemma 3.9](#) and [Lemma 3.10](#). \blacksquare

Lemma 3.12. *Let $\lambda < \eta$, $0 < \rho < 1$, and $\beta_k > 0$ as the initial step-size guess, with $\alpha_k = \beta_k \underline{\rho}^j$ for $j \in \mathbb{N}_0$. Let $\underline{\alpha}_k$ and $\bar{\alpha}_k$ be as defined in (3.11).*

- If $\bar{\epsilon}_k \geq 0$ is sufficiently small, then there exists $j_k \in \mathbb{N}_0$ such that $\underline{\alpha}_k \leq \beta_k \underline{\rho}^{j_k} \leq \bar{\alpha}_k$, i.e. $\psi(\alpha_k) \leq 0$ holds for $\alpha_k = \beta_k \underline{\rho}^{j_k}$.
- If $i_k > 0$ is the smallest integer for which $\psi(\alpha_k) \leq 0$ holds, then $\alpha_k = \beta_k \underline{\rho}^{i_k} > \underline{\rho} \bar{\alpha}_k$.

Proof. By [Lemma 3.8](#), for all $\bar{\epsilon}_k \in [0, s_k)$, $\bar{\alpha}_k > 0$. Let $\bar{\epsilon}_k \in [0, s_k)$; since $\lim_{j \rightarrow \infty} \beta_k \underline{\rho}^j = 0$, there exists an $j_k \in \mathbb{N}_0$ such that for all $j \geq j_k$, $\beta_k \underline{\rho}^j \leq \bar{\alpha}_k$. Referring again to [Lemma 3.8](#), we know that $\lim_{\bar{\epsilon}_k \rightarrow 0} \underline{\alpha}_k = 0$. Therefore, there exists $\epsilon \in [0, s_k)$ such that for all $\bar{\epsilon}_k \leq \epsilon$, $\underline{\alpha}_k \leq \beta_k \underline{\rho}^{j_k}$. Additionally, $\bar{\alpha}_k$ as a function of $\bar{\epsilon}_k$ is decreasing on the domain $[0, s_k)$. Hence, $\beta_k \underline{\rho}^{j_k} \leq \bar{\alpha}_k$ holds for $\bar{\epsilon}_k \leq \epsilon$.

As $i_k > 0$ be the smallest integer for which $\psi(\alpha_k) \leq 0$ holds, then $\alpha_k = \beta_k \underline{\rho}^{i_k} \leq \bar{\alpha}_k$. If $\alpha_k \leq \underline{\rho} \bar{\alpha}_k$, then $\beta_k \underline{\rho}^{i_k-1} \leq \bar{\alpha}_k$, which contradicts with the assumption that i_k is the smallest integer for which $\psi(\alpha_k) \leq 0$ holds. Hence, $\beta_k \underline{\rho}^{i_k} \leq \underline{\rho} \bar{\alpha}_k$. \blacksquare

Based on [Lemma 3.12](#), we can guarantee the existence of a suitable step size during the backtracking line search process if $\bar{\epsilon}_k$ is small enough.

3.2. Convergence Analysis. Prior to discussing the convergence theorem, we shall introduce the following auxiliary lemma.

Lemma 3.13. *Let $\epsilon > 0$ and $0 < \underline{c} < 1 < \bar{c}$. Further, let $\{a_k\}_{k=0}^\infty \subseteq [0, \infty)$ be the sequence*

$$a_{k+1} = \begin{cases} \bar{c} a_k & \text{if } a_k < \epsilon, \\ \underline{c}^{i_k} a_k & \text{if } a_k \geq \epsilon, \end{cases}$$

and i_k be the smallest non-negative integer such that $\underline{c}^{i_k} a_k < \epsilon$. Then, for all $k \geq 0$, it holds that $a_k \geq \min\{\underline{c} \epsilon, a_0\}$.

Proof. Assume, by induction, that for all $k = 0, \dots, t$, $a_t \geq \min\{\underline{c}\epsilon, a_0\}$. Now, consider a_{t+1} . If $a_t < \epsilon$, by definition, $a_{t+1} = \bar{c}a_t \geq \bar{c} \min\{\underline{c}\epsilon, a_0\} > \min\{\underline{c}\epsilon, a_0\}$. Otherwise, $a_t \geq \epsilon$. If $\underline{c}\epsilon \geq \underline{c}^{i_t}a_t$, then $\epsilon \geq \underline{c}^{i_t-1}a_t$, which contradicts the definition of i_t . Hence, $\underline{c}\epsilon < \underline{c}^{i_t}a_t$, which implies $a_{t+1} = \underline{c}^{i_t}a_t > \underline{c}\epsilon \geq \min\{\underline{c}\epsilon, a_0\}$. ■

Lemma 3.14. *Let Assumptions 2.1 and 2.2 hold, $\beta_0 > 0$, $\lambda < \eta$, and $0 < \underline{\rho} < 1 < \bar{\rho}$. Let i_k be the smallest non-negative integer such that $\alpha_k = \underline{\rho}^{i_k}\beta_k \in [\underline{\alpha}_k, \bar{\alpha}_k]$ and $\alpha_k > \underline{\rho}\bar{\alpha}_k$, if $i_k > 0$. Further, let $\beta_{k+1} = \bar{\rho}\alpha_k$, $\|z_k\| \neq 0$, and $\|e_k\| \leq (1 - \eta)\|z_k\|$. Denoting*

$$(3.13) \quad \tau := \min \left\{ \underline{\rho} \left(\frac{\eta - \lambda}{L_{\nabla f}} \right), \min \left\{ \underline{\rho} \left(\frac{2(\eta - \lambda)}{L_{\nabla f}} \right), \beta_0 \right\} \right\},$$

we have $\alpha_k \geq \tau$.

Proof. If $i_k > 0$, since $\alpha_k = \underline{\rho}^{i_k}\beta_k$, using Lemma 3.12 and (3.11) we have

$$(3.14) \quad \alpha_k > \underline{\rho}\bar{\alpha}_k \geq \underline{\rho} \left(\frac{\eta - \lambda}{L_{\nabla f}} \right).$$

Moreover, from Lemma 3.8, (3.11) yields

$$(3.15) \quad \alpha_k \leq \bar{\alpha}_k \leq \frac{2(\eta - \lambda)}{L_{\nabla f}}.$$

If $i_k = 0$, then $\alpha_k = \beta_k$. Considering the update $\beta_{k+1} = \bar{\rho}\alpha_k$, utilizing Lemma 3.13 by setting $a_k = \beta_k$, $\underline{c} = \underline{\rho}$, $\bar{c} = \bar{\rho}$, and $\epsilon = \frac{2(\eta - \lambda)}{L_{\nabla f}}$, based on (3.15), we derive $\beta_k \geq \min\{\underline{\rho} \left(\frac{2(\eta - \lambda)}{L_{\nabla f}} \right), \beta_0\}$. Taking τ as defined in (3.13) yields the desirable result. ■

Theorem 3.15. *Let the assumptions of Lemma 3.14 hold and $0 < \lambda < \eta < 1$. We have*

$$\lim_{k \rightarrow \infty} \|z_k\| = 0.$$

Proof. Lemma 3.5 and the backtracking line search rule $\psi(\alpha_k) \leq 0$ yield

$$(3.16) \quad f(\theta_{k+1}) - f(\theta_k) \leq -\lambda\alpha_k\|z_k\|^2.$$

Utilizing Lemma 3.14, since $0 < \lambda < \eta < 1$, we have $\tau > 0$, and recursively from the inequality (3.16) we can conclude

$$(3.17) \quad \sum_{k=1}^K \lambda\tau\|z_k\|^2 \leq f(\theta_1) - f(\theta_{K+1}).$$

As f is bounded below and $\lambda\tau > 0$ we conclude $\lim_{k \rightarrow \infty} \|z_k\| = 0$. ■

Corollary 3.16. *Under the assumptions of Theorem 3.15, we have*

$$\lim_{k \rightarrow \infty} \|\nabla f(\theta_k)\| = 0.$$

Proof. Since $\|e_k\| \leq (1 - \eta)\|z_k\|$, utilizing [Theorem 3.15](#) we have $\lim_{k \rightarrow \infty} \|z_k\| = 0$, which implies $\lim_{k \rightarrow \infty} \|e_k\| = 0$. Hence

$$\lim_{k \rightarrow \infty} \|\nabla f(\theta_k)\| \leq \lim_{k \rightarrow \infty} (\|e_k\| + \|z_k\|) = 0,$$

as required. ■

Note that, at each iteration $k \in \mathbb{N}_0$, based on [Proposition 3.1](#) and [Theorem 3.15](#), the approximate hypergradient should satisfy the inequality $\|e_k\| \leq (1 - \eta)\|z_k\|$. Otherwise, we can decrease the error in the hypergradient $\|e_k\|$ by means of multiplying ϵ_k and δ_k by a factor $0 < \underline{\rho} < 1$ as we do in the steps [10](#) and [11](#), then restarting the calculation of z_k using [Algorithm 3.2](#) until $\|e_k\| \leq (1 - \eta)\|z_k\|$ holds. Moreover, the steps of [Algorithm 3.1](#) are designed in a way to satisfy both the required accuracy for the line search $\psi(\alpha_k) \leq 0$. MAID also allows ϵ_k and δ_k to increase in steps [12](#) and [13](#) if they do not decrease in the subsequent iteration and remain suitable for the line search condition, which matches with $\bar{\epsilon}_k$ as mentioned in [Lemma 3.7](#) and [Lemma 3.8](#). This process in [Algorithm 3.1](#) and [Algorithm 3.2](#) provides us with an adaptive way of choosing the accuracy.

Proposition 3.17. *If $\|\nabla f(\theta_k)\| > 0$, the loop at line 3 in [Algorithm 3.2](#) terminates in finite time by finding a descent direction $-z_k$.*

Proof. From the bound [\(2.5\)](#) for the upper bound ω of the error $\|e_k\| = \|z_k - \nabla f(\theta_k)\|$, we have $\lim_{\epsilon_k, \delta_k \rightarrow 0} \omega = 0$. On the other hand, $\lim_{\epsilon_k, \delta_k \rightarrow 0} \|z_k\| = \|\nabla f(\theta_k)\| > 0$. Since $\epsilon_k \rightarrow 0$ and $\delta_k \rightarrow 0$ according to lines [11](#) and [10](#) of [Algorithm 3.2](#), respectively, and utilizing [Lemma 3.10](#), we conclude that the inequality $\omega \leq (1 - \eta)\|z_k\|$ holds for ϵ_k and δ_k sufficiently small. ■

Proposition 3.18. *The loop at line 3 in [Algorithm 3.1](#) terminates in finite time by successfully performing the update $\theta_{k+1} = \theta_k - \underline{\rho}^{i_k} \beta_k z_k$ for some $i_k \in \mathbb{N}_0$.*

Proof. If the inner loop at the line [5](#) fails to find a step size within the given maximum number of iterations j , two scenarios may occur. Firstly, if ϵ_k and δ_k are not sufficiently small, they are decreased in steps [9](#) and [10](#), followed by updating the descent direction z_k and a possible further decrease in accuracies ϵ_k and δ_k in step [4](#) as a part of the next iteration of the outer loop at line [3](#). This sequence of accuracy values will eventually become small enough to satisfy the assumptions in [Corollary 3.11](#) and [Lemma 3.12](#). Consequently, the required result holds.

Secondly, if $\underline{\rho}^{j-1} \beta_k > \bar{\alpha}_k$ and the backtracking process is unsuccessful, then the maximum number of backtracking iterations j goes to infinity during the iterations of the outer loop [3](#), and $\underline{\rho}^j \beta_k \rightarrow 0$. Since $\bar{\alpha}_k > 0$, this implies that there exists $j_k \in \mathbb{N}_0$ such that for all $i_k \geq j_k$, $\underline{\rho}^{i_k} \beta_k \leq \bar{\alpha}_k$. Therefore, the loop [Algorithm 3.1](#) terminates in a finite process. ■

Theorem 3.19. *Under Assumptions [2.1](#) and [2.2](#), the iterates θ_k of [Algorithm 3.1](#) satisfy*

$$\lim_{k \rightarrow \infty} \|\nabla f(\theta_k)\| = 0.$$

Proof. From [Proposition 3.18](#) we know that the loop at line [3](#) terminates after a finite number of iterations and finds a suitable step size α_k satisfying the line search condition $\psi(\alpha_k) \leq 0$. Moreover, the update of the initial step size in line [14](#) fulfills $\beta_{k+1} = \bar{\rho} \alpha_k$. Furthermore,

z_k computed in [Algorithm 3.2](#) satisfies $\|e_k\| \leq (1 - \eta)\|z_k\|$. Hence, all the assumptions of [Theorem 3.15](#) are satisfied by [Algorithm 3.1](#) and convergence follows from [Corollary 3.16](#). ■

Remark 3.20. The results presented in this section assume that the upper-level loss g does not depend on the parameters θ . However, if g depends on θ and is differentiable with respect to θ , the hypergradient takes the following form:

$$(3.18) \quad \nabla f(\theta) = \nabla(g \circ \hat{x})(\theta) = \nabla_{\theta} g(\hat{x}(\theta), \theta) + \partial \hat{x}(\theta)^T \nabla_x g(\hat{x}(\theta), \theta).$$

This modification affects the calculation of the inexact hypergradient, while the line search condition remains unchanged. This is because the bounds in [Lemma 3.3](#) were derived by expanding g with respect to its first argument (i.e., x). Moreover, assuming that $\nabla_{\theta} g(\hat{x}(\theta), \theta)$ is $L_{\nabla_{\theta} g}$ -Lipschitz in x , the right-hand side of bound (2.5) will include an additional term, $L_{\nabla_{\theta} g} \epsilon_k$.

4. Numerical Experiments. In this section, we present the numerical results of our proposed algorithm MAID ([Algorithm 3.1](#)) and compare it with HOAG [45] and the dynamic derivative-free algorithm DFO-LS [21]. For HOAG, we take the adaptive step size strategy used in the numerical experiments of [45] with accuracies either following a geometric sequence $\epsilon_k = 0.9^k \epsilon_0$, a quadratic sequence $\epsilon_k = k^{-2} \epsilon_0$, or a cubic sequence $\epsilon_k = k^{-3} \epsilon_0$ for $k = 1, 2, \dots$. The implementation of HOAG² and the Fast Iterative Shrinkage-Thresholding Algorithm (FISTA) variant of DFO-LS³ is based on the available code from the corresponding papers [45] and [21], respectively. We will make the codes of our implementations available on GitHub upon the acceptance of the paper. Moreover, we use warm-starting in the lower-level iterations, which means we save and use the lower-level solution as the initial point for the next call of the lower-level solver. This strategy is also utilized in DFO-LS and HOAG. For the constant $L_{H^{-1}}$ in the bound (2.5) used in [Algorithm 3.2](#), we approximate it as

$$\frac{\|\nabla_x^2 h(\tilde{x}(\theta), \theta)^{-1} \nabla g(\tilde{x}(\theta)) - \nabla_x^2 h(\tilde{x}(\theta), \theta)^{-1} u\|}{\|\nabla g(\tilde{x}(\theta)) - u\|},$$

where $u \sim \mathcal{N}(0, 1)$ is a random perturbation of $\nabla g(\tilde{x}(\theta))$. We then store the maximum value of this approximation across the upper-level iterations. In this approach, since we have already calculated $\nabla_x^2 h(\tilde{x}(\theta), \theta)^{-1} \nabla g(\tilde{x}(\theta))$ to compute the hypergradient, we possess an efficient approximation that improves by taking its maximum across upper-level iterations. We use a similar process for estimating L_J . Moreover, we estimate the $\|\nabla_{x\theta}^2 h(x(\theta), \theta)\|$ on the right-hand side of (2.5) using one iteration of the power method (to compute largest eigenvalue of $\nabla_{x\theta}^2 h(x(\theta), \theta)^T \nabla_{x\theta}^2 h(x(\theta), \theta)$). We include these calculations in our lower-level computational cost. Note that in all experiments where the upper-level loss is convex, the bound (3.6) from [Remark 3.6](#) will be used in step 6 of [Algorithm 3.1](#).

4.1. Validating the performance with known optimal parameters. Similar to [22], we first consider a simple linear least-squares problem taken from [36]:

$$(4.1a) \quad \min_{\theta \in \mathbb{R}^{10}} f(\theta) := \|A_1 \hat{x}(\theta) - b_1\|^2$$

$$(4.1b) \quad s.t. \quad \hat{x}(\theta) := \arg \min_{x \in \mathbb{R}^{10}} \|A_2 x + A_3 \theta - b_2\|^2,$$

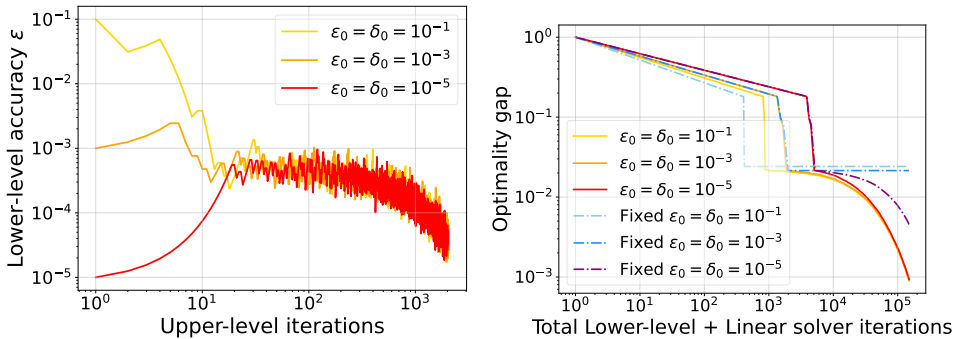
²<https://github.com/fabianp/hoag>

³https://github.com/lindonroberts/inexact_dfo_bilevel_learning

where $A_i \in \mathbb{R}^{1000 \times 10}$ have random entries uniformly distributed between 0 and 1, and $b_i \in \mathbb{R}^{1000}$ are given by $b_1 = A_1 \hat{x}_1 + 0.01y_1$ and $b_2 = A_2 \hat{x}_2 + A_3 \bar{\theta} + 0.01y_2$, where \hat{x}_1, \hat{x}_2 , and $\bar{\theta} \in \mathbb{R}^{10}$ have i.i.d. entries uniformly distributed between 0 and 1, and $y_1, y_2 \in \mathbb{R}^{1000}$ are independent standard Gaussian vectors with entries from $\mathcal{N}(0, 1)$. For our experiments, we pick θ_0 as the vector of all ones. In this problem, we have the capability to analytically compute $\hat{x}(\theta)$, $\nabla f(\theta)$, and all of the Lipschitz constants ($L_{\nabla f}$, $L_{\nabla g}$, $L_{H^{-1}}$, and L_J), along with determining the optimal θ^* . Therefore, this problem serves as a benchmark to evaluate the performance of our inexact algorithm.

For the experimental evaluation, we have chosen several configurations for the parameters in our algorithm. Specifically, these encompass $\epsilon_0 = \delta_0 = 10^{-1}$, $\epsilon_0 = \delta_0 = 10^{-3}$, $\epsilon_0 = \delta_0 = 10^{-5}$. Moreover, in MAID, we set the hyperparameters $\bar{\rho} = \frac{10}{9}$, $\underline{\rho} = 0.5$, $\bar{\nu} = 1.25$, and $\underline{\nu} = 0.5$. Across all instances of fixed accuracy, the line search mechanism of MAID ($\psi(\alpha_k) \leq 0$) is employed to determine the step size. Subsequently, we employ FISTA adapted to strongly convex problems to solve the lower-level problem [14].

Since the main computational cost involved in solving bilevel optimization problems lies in solving the lower-level problems, we have imposed a cap on the total number of lower-level iterations. Furthermore, recognizing the computational analogy between the cost of each CG iteration in the IFT approach, dominated by a single Hessian-vector product, and each iteration of the lower-level solver (e.g., FISTA), dominated by a single gradient evaluation, we treat CG iterations with comparable significance to that of an iteration of the lower-level solver. Consequently, the amalgamation of lower-level solver iterations and CG iterations is regarded as constituting the cumulative computational cost of the lower-level phase. For solving the problem (4.1), we set a termination budget of 1.5×10^5 lower-level iterations.



(a) Accuracy ϵ of the lower-level solver in each upper-level iteration (b) Upper-level loss vs total calls of linear solver and lower-level solver

Figure 1: The value of ϵ per upper-level iteration, alongside the performance and computational cost of MAID, is illustrated across various settings for solving the problem (4.1). All MAID configurations achieve lower loss values compared to the high fixed accuracy, and they do so at a lower computational cost. Furthermore, they consistently converge to the same range of the required accuracy $\epsilon = \delta$, regardless of the initial choices of ϵ_0 .

Figure 1a illustrates the determination of the necessary values for ϵ in solving the problem (4.1) through the MAID algorithm. As shown in Figure 1a, irrespective of the initial $\epsilon_0 = \delta_0$ as the accuracy of the lower-level solver and the linear solver, respectively, MAID adaptively reduces ϵ and δ , ultimately opting for accuracies around 2×10^{-5} .

Figure 1b illustrates the computational advantages of adaptivity in MAID compared to fixed low and high accuracies, with the main computational costs stemming from lower-level and linear solver iterations. Setting the low fixed accuracy to $\epsilon = 10^{-1}$ and $\epsilon = 10^{-3}$ results in loss stalling after a few upper-level iterations due to insufficient accuracy for line search progression. Conversely, using the higher fixed accuracy of $\epsilon = 10^{-5}$ yields a higher loss than dynamic configurations, due to exhausting the total lower-level iteration budget. This underscores the importance of the adaptive approach in MAID, as the suitable fixed accuracy is not known a priori; it may be too small, leading to failure (e.g., $\epsilon = 10^{-1}$), or require significantly higher computational costs for progress (e.g., $\epsilon = 10^{-5}$). Overall, dynamic configurations successfully solve problem (4.1) with lower computational costs compared to a fixed accuracy of 10^{-5} , regardless of the starting accuracy. While in Figure 1b we plot the optimality gap, it is also




Figure 2: The range of error in upper-level function evaluation given error bounds (3.3a) and (3.4), inexact loss given the approximated lower-level solution, and exact loss evaluation $f(\theta) = g(\hat{x}(\theta))$.

possible to compute the exact solution of the lower-level problem in the specific problem (4.1). This enables the calculation of the precise loss at each upper-level iteration. This ability allows us to verify both the sufficient decrease condition and the progress of MAID. To facilitate this verification, we present the exact loss alongside the loss obtained using MAID with $\epsilon_0 = \delta_0 = 1$ in Figure 2. Additionally, this figure visually represents the bounds for inexact loss evaluation, as defined by (3.3a) and (3.4), wherein the area between these bounds encapsulates both exact and inexact loss values. Furthermore, as observed in Figure 2, enhancing the accuracy of the lower-level solver leads to a narrower interval between the upper and lower bounds, thereby bringing the exact and inexact loss values closer.

4.2. Multinomial logistic loss. In order to compare MAID with HOAG [45] as a well-known gradient-based algorithm with inexact hypergradients, we replicate the multinomial logistic

regression bilevel problem used in [45, 38] on the MNIST⁴ dataset. Similar to [45], we take the entire $M = 60000$ training samples of MNIST and rescale them to 12×12 pixels. Moreover, we utilize all $N = 10000$ samples of the validation set for the upper-level problem. The bilevel problem corresponding to this model has the following form:

$$(4.2a) \quad \min_{\theta \in \mathbb{R}^d} f(\theta) := \sum_{i=1}^N \Psi(\bar{b}_i, \bar{a}_i^T \hat{x}(\theta))$$

$$(4.2b) \quad s.t. \quad \hat{x}(\theta) := \arg \min_{x \in \mathbb{R}^{p \times q}} \sum_{j=1}^M \Psi(b_j, a_j^T x) + \frac{1}{2} \sum_{k=1}^p \sum_{l=1}^q e^{\theta[(k-1)q+l]} x_{k,l}^2,$$

where $p = 12 \times 12 = 144$ represents the number of features, $q = 10$ represents the number of classes, and $d = p \times q = 1440$ denotes the number of parameters. Ψ represents the multinomial logistic loss [41, Section 10.3], while $a_j \in \mathbb{R}^p$ and $b_j \in \mathbb{R}^q$ correspond to the training images and labels for each $1 \leq j \leq M$. Similarly, $\bar{a}_i \in \mathbb{R}^p$ and $\bar{b}_i \in \mathbb{R}^q$ denote the validation images and labels in the upper-level problem for each $1 \leq i \leq N$. To solve this problem, we employ MAID, alongside HOAG using its heuristic adaptive step size with default parameters. We set the initial accuracy for both algorithms to $\epsilon_0 = \delta_0 = 10^{-1}$, as specified in [45]. Moreover, to ensure a fair comparison, similar to HOAG, we set the hyperparameters $\bar{\rho} = \frac{10}{9}$, $\underline{\rho} = 0.5$, $\bar{\nu} = 1.05$, and $\underline{\nu} = 0.5$. In HOAG, the initial step size $\beta_0 = \sqrt{d}/\|z_0\|$ was determined using the same strategy as outlined in [45]. We set the initial upper-level step size α_0 in MAID equal to the initial step size in HOAG. Additionally, we consider a fixed total lower-level budget of 6×10^5 iterations and a maximum of 300 upper-level iterations.

Considering the fixed computational budget, Figure 3 offers a comparison between HOAG and MAID in terms of upper-level step size, the accuracy of the lower-level solver, and loss. Throughout this experiment, all configurations of HOAG, employing the adaptive step size strategy from [45] and depending on their accuracy regime, demonstrated the ability to effectively adjust the step size and accuracy, leading to loss reduction in (4.2) within a few upper-level iterations and their corresponding lower-level computational cost. However, HOAG's adaptive step size eventually stops changing and becomes fixed at different levels, depending on the accuracy regime, while diminishing accuracy for the lower-level solver to negligible values. As a result, the lack of adaptivity in accuracy and reaching high accuracies lead to slower loss reduction using the quadratic and cubic sequences of accuracies.

In contrast, while using the same initial step size and accuracy as HOAG, MAID consistently identified an appropriate range for accuracy and upper-level step size. This led to driving the loss to lower values than those achieved by HOAG within the same budget. Moreover, upon fitting a polynomial regime to the accuracy sequence taken by MAID, the sequence $\epsilon_k = 0.35(k + 0.14)^{-1.05} \epsilon_0 - 0.0009$ for $k = 1, 2, \dots$ could potentially be a suitable option for HOAG when compared to the three other named sequences. Figure 4 supports this guess, showing that HOAG with the specified sequence performs comparably to the best of the named sequences, which in this case is geometric. Importantly, since this sequence is summable, the theoretical convergence of HOAG also holds. Nevertheless, knowing such a sequence a priori is not straightforward, and MAID finds it adaptively.

⁴<http://yann.lecun.com/exdb/mnist/>




Figure 3: The value of ϵ , and upper-level loss per total lower-level computational cost in the algorithms MAID and HOAG for solving the problem (4.2). It illustrates that MAID adaptively adjusts the upper-level step size and the required accuracy for the solving the lower-level problem, while HOAG with the adaptive step size does not adjust the step size after a few iterations and exhibits a decreasing sequence for the accuracy ϵ . In all HOAG and MAID settings, $\epsilon_0 = \delta_0$ is set.




Figure 4: Comparison of HOAG with geometric, quadratic, and cubic accuracy sequences against a custom sequence obtained by fitting an accuracy regime to the sequence chosen by MAID.

In summation, MAID demonstrates superior adaptability and performance compared to HOAG, achieving lower loss values within the same lower-level computational budget. MAID achieved this by adaptively adjusting accuracy and step size, and showcased an enhanced efficiency. Additionally, the step size selection in MAID comes with a theoretical convergence guarantee, whereas the strategy in HOAG is heuristic.

4.3. Total variation denoising. In this part, we consider the smoothed version of the Rudin-Osher-Fatemi (ROF) image-denoising model [48] as implemented in [21]. Here, a bilevel problem is formulated to learn the smoothing and regularization parameters of total variation (TV) regularization. In contrast to the first numerical experiment in this section, access to the optimal hyperparameters is not possible a priori.

Similar to [21], we apply this model to 2D images and the training set is $m = 25$ images of the Kodak dataset⁵ resized to 96×96 pixels, converted to grayscale, and Gaussian noise with $\mathcal{N}(0, \sigma^2)$ with $\sigma = 0.1$ is added. The bilevel problem we examine is as follows:

$$(4.3a) \quad \min_{\theta} \frac{1}{m} \sum_{t=1}^m \frac{1}{2} \|\hat{x}_t(\theta) - x_t^*\|^2$$

$$(4.3b) \quad s.t. \quad \hat{x}_t(\theta) = \arg \min_x \frac{1}{2} \|x - y_t\|^2 + e^{\theta[1]} \sum_{i,j} \sqrt{|\nabla x_i|^2 + |\nabla x_j|^2 + (e^{\theta[2]})^2},$$

Here, ∇x_i and ∇x_j represent the i -th and j -th elements of the forward differences of x in the horizontal and vertical directions, respectively. It is important to note that this problem satisfies the assumption 2.1. To initiate the process, we set $\theta_0 = (-5, -5)$ and utilize Algorithm 3.1 with hyperparameters $\bar{\rho} = \frac{10}{9}$, $\underline{\rho} = 0.5$, $\bar{\nu} = 1.25$, and $\underline{\nu} = 0.5$, along with the FISTA variant of DFO-LS with dynamic accuracy [21] to solve the problem (4.3). The DFO-LS with dynamic accuracy is a derivative-free bilevel algorithm that employs FISTA for solving the lower-level problem with an accuracy determined dynamically from the trust region radius in the derivative-free optimization algorithm for solving least-squares [13] in the upper-level problem. This algorithm serves as an efficient approach for solving bilevel problems and in particular for the problem (4.3). However, it should be noted that, as discussed in [17], in contrast with gradient-based algorithms, it may not scale well with the number of hyperparameters due to the nature of the derivative-free algorithm used for the upper-level problem. To maintain consistency with the first numerical example, we set a limit on the total number of lower-level iterations in this case.

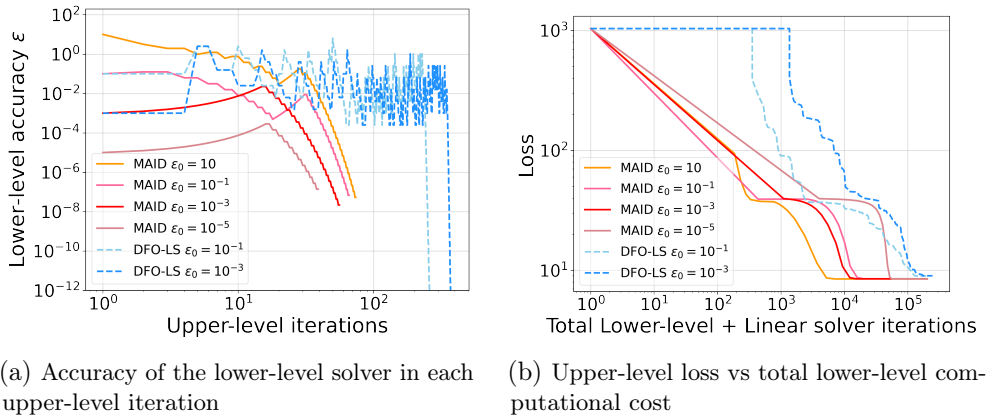


Figure 5: The values of ϵ , and upper-level loss per total lower-level computational cost in the algorithms MAID and DFO-LS for solving the problem (4.3). In various MAID settings, $\epsilon_0 = \delta_0$ is set.

Illustrated in Figure 5a, MAID demonstrates that, regardless of the starting accuracy of

⁵<https://r0k.us/graphics/kodak/>

the lower-level solver ϵ —chosen with different orders of magnitude—it adapts the accuracy and converges to the same range after a few upper-level iterations. This highlights the robustness of MAID in the selection of an initial accuracy, enabling it to find values as large as necessary and as small as needed to make progress in backtracking line search. On the other hand, the FISTA variant of DFO-LS initially selects an equal or higher value of ϵ compared to MAID when using the same starting accuracies. However, at a certain point, it begins to significantly decrease ϵ .

In terms of loss, as indicated in Figure 5b, with an equal total lower-level computational budget (assuming CG iterations of DFO-LS as zero), MAID surpasses the performance of the FISTA variant of DFO-LS, even in the case where the initial ϵ is chosen too small. All different setups of MAID converged to the loss value of 8.49, while DFO-LS with $\epsilon_0 = 10^{-1}$ and $\epsilon_0 = 10^{-3}$ reached 8.97 and 9.00, respectively. Furthermore, it is important to note that DFO-LS does not utilize any hypergradient information, which could place it at a disadvantage when compared to gradient-based methods.

To further compare MAID with the existing inexact first-order method HOAG, we selected the best initial setting from Figure 5b (i.e., $\epsilon_0 = \delta_0 = 10$) and tuned HOAG with geometric, quadratic, and cubic accuracy sequences. Figure 6 presents a performance comparison between these variants of HOAG and MAID. As shown in Figure 6a, MAID adaptively selected the required accuracy, while the HOAG variants followed three different a priori accuracy sequences, each giving different results in Figure 6b. After 10^4 lower-level computations (lower-level solver plus linear solver iterations), both MAID and the quadratic variant HOAG converged. Extending the computational budget allowed HOAG with cubic accuracy to converge as well, though at a significantly higher computational cost. In contrast, HOAG with a geometric sequence converged to a higher loss. These outcomes, illustrated in Figure 6, highlight MAID’s adaptability and robustness in selecting suitable accuracy without tuning or requiring prior knowledge. Moreover, MAID’s faster convergence with lower computational costs underscores its better performance compared to HOAG for solving (4.3).

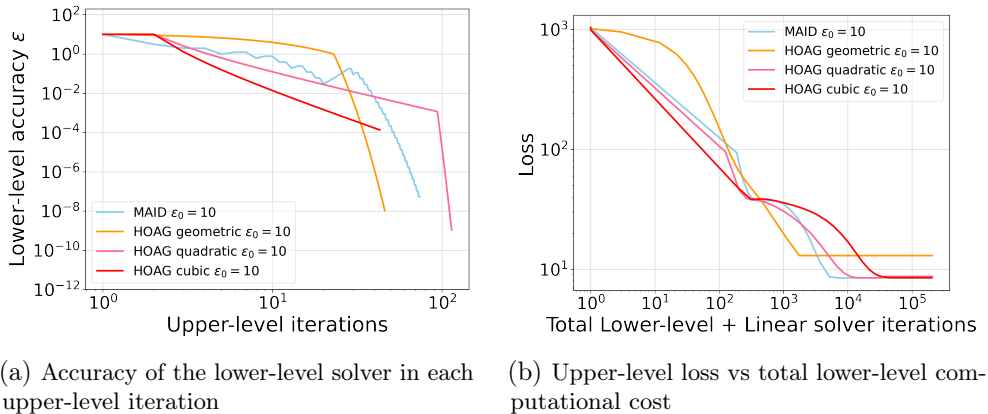


Figure 6: The values of ϵ , and upper-level loss per total lower-level computational cost in the algorithms MAID and HOAG for solving the problem (4.3). In all settings of HOAG and MAID $\epsilon_0 = \delta_0$ is set.

Now, to evaluate the effectiveness of the learned parameters obtained through solving (4.3) using MAID and dynamic DFO-LS, We solved the training problem (4.3) for 20 images of the Kodak dataset, each of size 256×256 with Gaussian noise $\sigma = 0.1$, using MAID with $\epsilon_0 = \delta_0 = 10^{-1}$ and the FISTA variant of dynamic DFO-LS with $\epsilon_0 = \Delta_0 = 10^{-1}$, where Δ_0 represents the trust region radius in the first upper-level iteration. Subsequently, we applied these learned parameters to the lower-level total variation denoising problem for a test noisy image of size 256×256 and standard deviation $\sigma = 0.1$.

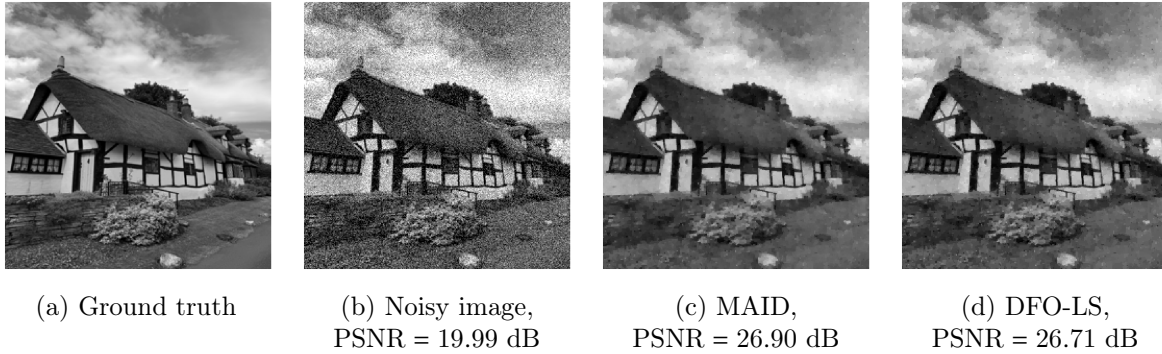


Figure 7: Total variation denoising with learned parameters by MAID ($\epsilon_0 = \delta_0 = 10^{-1}$), dynamic DFO-LS ($\epsilon_0 = 10^{-1}$).

Figure 7 shows the denoising results for both algorithms. It displays the original and noisy image as well as the denoised images obtained using MAID and DFO-LS, respectively. Each denoised image is accompanied by its corresponding peak signal-to-noise ratio (PSNR) value. The learned hyperparameters using MAID result in slightly higher PSNR values than DFO-LS learned hyperparameters on test images, which is consistent with the training results. This indicates that both MAID and DFO-LS effectively solve (4.3) with adaptability and robustness to initial accuracy. However, MAID leverages first-order information, resulting in superior performance and scalability.

Using the same setup as in the comparison with DFO-LS illustrated in Figure 7, we solve (4.3) using MAID and HOAG variants. Given a computational budget of 10^4 , the resulting denoised test image based on parameters learned by MAID and HOAG is shown in Figure 8a. As expected from Figure 6b, only the quadratic variant of HOAG achieves comparable results to MAID, while the cubic variant requires more iterations, and the geometric variant remains stalled in a suboptimal solution, potentially requiring further tuning. Thus, Figure 8 underscores the performance and robustness of MAID over HOAG in this example.

We now demonstrate MAID's performance on a problem involving a non-convex upper-level loss. Consider the total variation denoising lower-level problem (4.3b) within the bilevel setting (4.3). We replace the upper-level loss in (4.3a) with the following smooth, non-convex loss from [3, 11]:

$$(4.4) \quad g(x) = \frac{\|x - x^*\|^2}{1 + \|x - x^*\|^2}.$$




Figure 8: Total variation denoising with learned parameters by MAID and HOAG ($\epsilon_0 = \delta_0 = 10$) with 10^4 computational budget.

This results in the modified upper-level problem below:

$$(4.5) \quad \min_{\theta} \frac{1}{m} \sum_{t=1}^m \frac{\|x_t - x_t^*\|^2}{1 + \|x_t - x_t^*\|^2}.$$

Using the same setting as in (4.3) and applying (3.5) in [Algorithm 3.1](#), we solve this problem. [Figure 9](#) illustrates MAID's convergence behavior and robustness to the initial accuracy when solving the bilevel problem with the non-convex upper-level loss (4.4).

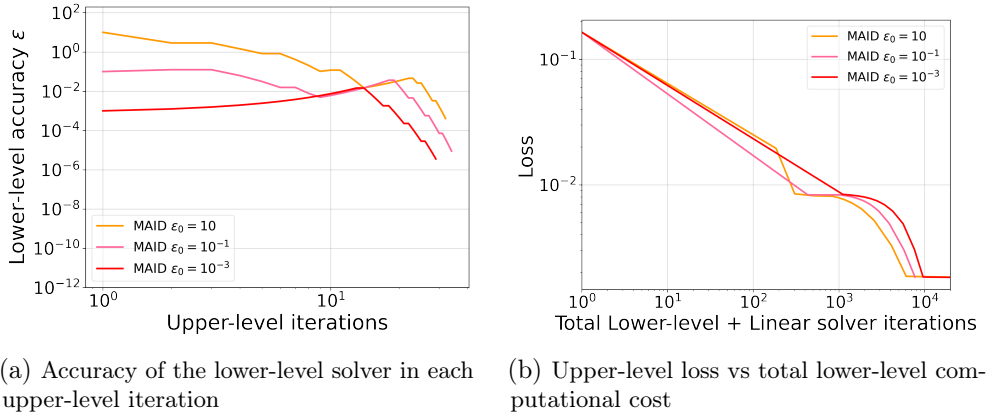


Figure 9: Values of ϵ , and upper-level loss per total lower-level computational cost in MAID for learning parameters in the ROF model (4.3b) with the non-convex upper-level loss (4.4) and corresponding upper-level problem (4.5). In various MAID settings, $\epsilon_0 = \delta_0$ is set.

4.4. Field of experts denoising. In this part, we explore image denoising using a more data-adaptive regularizer known as the Field of Experts (FoE) approach. This approach was

previously employed in a bilevel setting in [16] to learn the parameters of the regularizer in the variational setting. However, their choice of regularizer was non-convex and does not align with our assumptions. Therefore, we adopt the convex case, as described in [17], resulting in the following bilevel problem:

$$(4.6a) \quad \min_{\theta} \frac{1}{m} \sum_{i=1}^m \frac{1}{2} \|\hat{x}_i(\theta) - x_i^*\|^2$$

$$(4.6b) \quad s.t. \quad \hat{x}_i(\theta) = \arg \min_x \left\{ \frac{1}{2} \|x - y_i\|^2 + e^{\theta[0]} \sum_{k=1}^K e^{\theta[k]} \|c_k * x\|_{\theta[K+k]} \right\}, \quad i = 1, \dots, m,$$

where each c_k is a $2D$ convolution kernel. We set $K = 30$ and take each c_k as a 7×7 kernel with elements $[\theta[2K + 49(k - 1) + 1], \dots, \theta[2K + 49k]]$. Also, $\|x\|_{\nu} = \sum_{j=1}^n \sqrt{x_j^2 + \nu^2}$ represents the smoothed 1-norm with smoothing parameter ν for any $x \in \mathbb{R}^n$. So, this problem has 1531 hyperparameters that we represent by θ . We initialize the kernels by drawing randomly from a Gaussian distribution with mean 0 and standard deviation 1, and we set the initial weights to e^{-3} . During training, similar to TV denoising, we use 25 images from the Kodak dataset and rescale them to 96×96 pixels, converting them to grayscale. Additionally, we add Gaussian noise with $\mathcal{N}(0, \sigma^2)$, where $\sigma = 0.1$, to the ground truth images. Additionally, as in the previous experiment, we set the MAID hyperparameter $\bar{\rho} = \frac{10}{9}$, $\underline{\rho} = 0.5$, $\bar{\nu} = 1.25$, and $\underline{\nu} = 0.5$.

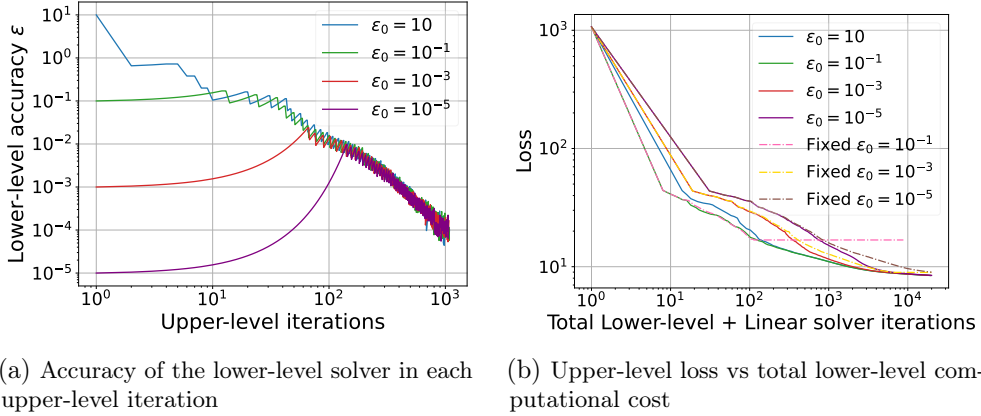


Figure 10: Investigating the impact of MAID configurations on the accuracy of the lower-level solver, and loss at each lower-level computation unit with varied $\epsilon_0 = \delta_0$ for solving the problem (4.6).

As seen in Figure 10a, by dynamically adjusting the accuracy within the framework of MAID, After a few upper-level iterations in various settings, the algorithm ultimately selects the same range of accuracies. Akin to the patterns observed in previous numerical experiments, choosing a too-large initial accuracy $\epsilon_0 = 10$ in this case prompts the algorithm to undergo reductions and recalculations to find a descent direction, resulting in higher computational costs

initially. However, after sufficiently reducing ϵ , it manages to make progress with a reasonable computational cost.

The advantage of adaptive accuracy in MAID becomes evident when comparing it to fixed accuracies. Using a low fixed accuracy $\epsilon_0 = 10^{-1}$ along with the line search mechanism of MAID, one can see a rapid progress initially, as indicated by the pink dashed curve, but it reaches a plateau after a few upper-level iterations. Furthermore, it does not use the full lower-level budget because the upper-level progress stalls at some point, and due to warm-starting, the lower-level solution is accurate enough and does not require many more iterations. A medium fixed accuracy, such as 10^{-3} , performs comparably to adaptive cases; however, it cannot be known a priori and must be chosen experimentally. A high fixed accuracy, like $\epsilon_0 = 10^{-5}$, appears sufficiently small but exhausts the lower-level solver and runs out of budget while achieving higher loss. All adaptive instances demonstrate enhanced overall efficiency compared to their fixed counterparts, as showcased in Figure 10b.

To compare MAID with HOAG, a scalable inexact first-order algorithm for solving (4.6), we ran HOAG with geometric, quadratic, and cubic accuracy sequences under the same settings, with $\epsilon_0 = \delta_0 = 10^{-1}$. As illustrated in Figure 11a, MAID adaptively selected larger ϵ values while maintaining steady progress in Figure 11b. For any budget exceeding 10^3 , MAID outperforms HOAG. Different HOAG variants vary in performance and likely require further tuning, while MAID remains robust to accuracy choices. The best-performing HOAG variant, quadratic, fluctuates significantly after 2×10^4 computations, whereas MAID exhibits a steady, monotonic behavior. Overall, Figure 11 highlights MAID's superior performance and stability in this experiment.

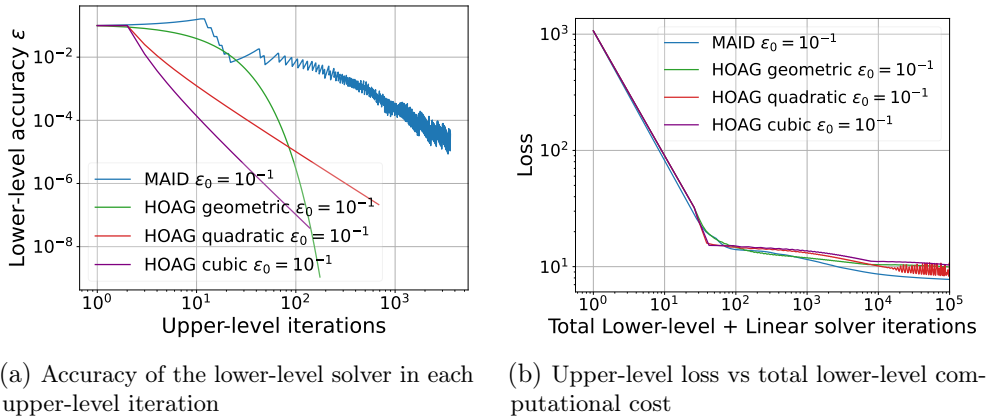


Figure 11: Comparison of MAID and HOAG on the lower-level solver accuracy and loss per lower-level computation unit, with initial settings $\epsilon_0 = \delta_0 = 10^{-1}$ for solving problem (4.6).

Noting the importance of the problem (4.6) as a large-scale and data-adaptive problem, to evaluate the learned parameters, we utilized the same set of 20 images used in the total variation problem, each of size 256×256 . Additionally, we employed 48 filters of size 7×7 , initialized with zero-mean random filters. After training using MAID and variation of HOAG

with $\epsilon_0 = \delta_0 = 10^{-1}$, we evaluated the learned model on a noisy test image of size 256×256 pixels with Gaussian noise of $\sigma = 0.1$, as depicted in Figure 12a and Figure 12b respectively.

The denoised image obtained using parameters learned by MAID is shown in Figure 12c, achieving a PSNR of 29.73 dB. In comparison, the best-performing variant of HOAG in this experiment—the quadratic variant—produces a denoised image with a PSNR of 28.79 dB, as shown in Figure 12d. The quality drops to 28.03 dB and 27.96 dB when using the geometric and cubic variants of HOAG, respectively, as shown in Figure 12e and Figure 12f. In contrast, MAID maintains robust performance without requiring fine-tuning or prior knowledge of accuracy.

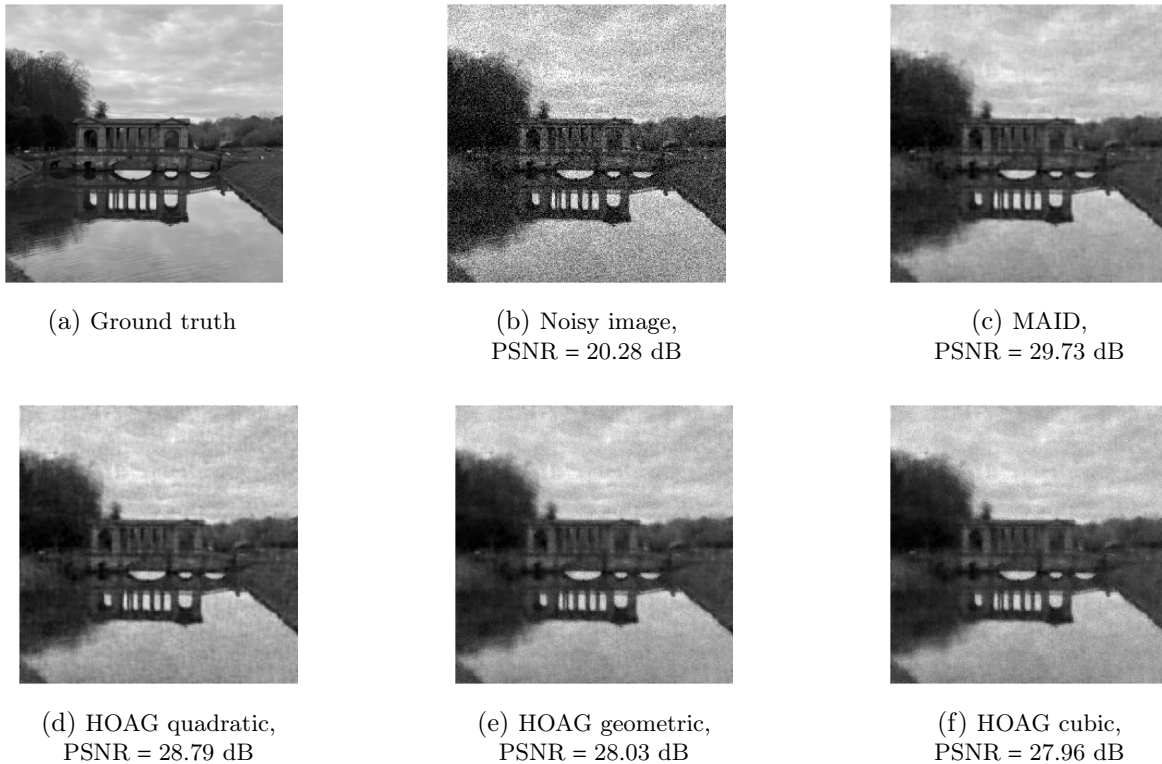


Figure 12: Field of experts denoising with learned parameters utilizing MAID in comparison with HOAG.

4.5. Runtime comparison. In our numerical results, the lower-level computational cost includes iterations of the lower-level solver as well as the matrix-vector products in the linear solver and power method involved in computing the bound (2.5). These were considered the primary computational costs, as CPU time can vary across problems and depends heavily on implementation. To support this choice, we present a comparison of CPU times for the lower-level solver, the matrix-vector product (linear solver plus power method), the computation of the bound (2.5) for checking a descent direction, and the line search condition (3.5) in

Lemma 3.5 across all experiments in [Table 1](#). The results show that, as expected, the majority of CPU time per upper-level iteration is spent on lower-level iterations and matrix-vector product.

In addition to these results, note that, in terms of computational complexity, the extra calculations of $\nabla_{x\theta}^2 h(x(\theta), \theta)$ on the right-hand side of [\(2.5\)](#) (for example, when approximating $\|\nabla_{x\theta}^2 h(x(\theta), \theta)\|$ using one step of the power method) are approximately twice as costly as computing a single gradient $\nabla_x h(x, \theta)$, assuming the dimension d of the parameter θ is comparable to the dimension n of the input x . This is because the calculations are performed in a matrix-vector product manner, resulting in roughly two additional lower-level gradient computations per upper-level iteration. Moreover, we found the choice of using only one step of the power method to be numerically effective when computing the operator norm in [\(2.5\)](#). Increasing the number of power method iterations did not change the behavior of the algorithm but led to higher computational costs than necessary.

Experiment	FISTA	Matrix-vector	Bound (2.5)	Line search (3.5)
Quadratic Section 4.1	7.8×10^{-2}	7.1×10^{-2}	6.9×10^{-5}	5.8×10^{-5}
Regression Section 4.2	1.2×10^1	2.9×10^0	2.3×10^{-2}	2.4×10^{-3}
TV denoising Section 4.3	6.2×10^0	1.5×10^{-1}	2.6×10^{-4}	1.4×10^{-4}
FoE denoising Section 4.4	1.8×10^1	4.0×10^0	2.8×10^{-4}	2.4×10^{-4}

Table 1: Average CPU time (seconds) for total lower-level computational cost (lower-level solver (FISTA) iterations + matrix-vector products), descent direction check (evaluating bound [\(2.5\)](#)), and backtracking line search per upper-level iteration of MAID. All experiments use the settings from [section 4](#) with $\epsilon_0 = \delta_0 = 10^{-1}$.

5. Conclusions and future work. In this paper, we proposed the MAID algorithm, along with its convergence analysis and numerical evaluation. For future work, it would be valuable to explore the convergence rate of MAID in addition to the asymptotic convergence results presented here. One promising direction is to study non-monotone line search techniques [\[54, 25\]](#) as an extension to the current analysis, which is based on an Armijo-type monotone line search. Additionally, given the success of stochastic line search methods [\[52, 25, 5\]](#), investigating stochastic variants of MAID could enable scalability to large datasets and potentially provide speed-ups.

A future direction could involve extending the analysis of our method to accommodate potentially non-convex lower-level functions utilizing the *value function* approach [\[18\]](#). Additionally, comparing fully first-order methods [\[37, 34, 35\]](#) with inexact IFT+CG-based approaches, as well as exploring ways to combine the adaptivity and robustness of MAID with fully first-order methods to leverage their avoidance of second-order computations, presents a promising avenue for further study.

MAID is the first algorithm which actively exploits the fundamental trade-off in the accuracy and computational cost in estimating function values and gradients in bilevel learning problems. Rather than a-priori selected, MAID selects accuracies for both function values and gradients following the overarching paradigm to be as low accurate as possible to save computational

cost but as accurate as needed to guarantee global convergence to critical points. A key novelty is that MAID performs backtracking using inexact computations to assure monotonic descent. Our numerical results underscore the superiority of MAID over state-of-the-art derivative-free and first-order methods for bilevel learning across a range of problems in data science. Importantly, it demonstrates remarkable robustness in the face of hyperparameters such as the initial accuracy and starting step size choices, alleviating the need for tuning these algorithm parameters specifically for each application.

Acknowledgments. This research made use of Hex, the GPU Cloud in the Department of Computer Science at the University of Bath.

References.

- [1] B. AMOS, L. XU, AND J. Z. KOLTER, *Input convex neural networks*, in Proceedings of the 34th International Conference on Machine Learning, D. Precup and Y. W. Teh, eds., vol. 70 of Proceedings of Machine Learning Research, PMLR, 06–11 Aug 2017, pp. 146–155.
- [2] C. ARNDT, *Regularization theory of the analytic deep prior approach*, Inverse Problems, 38 (2022), p. 115005.
- [3] A. E. BEATON AND J. W. TUKEY, *The fitting of power series, meaning polynomials, illustrated on band-spectroscopic data*, Technometrics, 16 (1974), pp. 147–185.
- [4] Y. BENGIO, *Gradient-based optimization of hyperparameters*, Neural Computation, 12 (2000), pp. 1889–1900.
- [5] A. S. BERAHAS, L. CAO, AND K. SCHEINBERG, *Global convergence rate analysis of a generic line search algorithm with noise*, SIAM Journal on Optimization, 31 (2021), pp. 1489–1518.
- [6] J. BERGSTRA AND Y. BENGIO, *Random search for hyper-parameter optimization*, J. Mach. Learn. Res., 13 (2012), p. 281–305.
- [7] L. BOGENSPERGER, A. CHAMBOLLE, AND T. POCK, *Convergence of a piggyback-style method for the differentiation of solutions of standard saddle-point problems*, SIAM Journal on Mathematics of Data Science, 4 (2022), pp. 1003–1030.
- [8] J. BOLTE, E. PAUWELS, AND S. VAITER, *Automatic differentiation of nonsmooth iterative algorithms*, in Advances in Neural Information Processing Systems, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, eds., vol. 35, Curran Associates, Inc., 2022, pp. 26404–26417.
- [9] K. BREDIES, M. CARIONI, AND M. HOLLER, *Regularization graphs—a unified framework for variational regularization of inverse problems*, Inverse Problems, 38 (2022), p. 105006.
- [10] T. A. BUBBA, L. CALATRONI, A. CATOZZI, S. CRISCI, T. POCK, M. PRAGLIOLA, S. RAUTIO, D. RICCIO, AND A. SEBASTIANI, *Bilevel learning of regularization models and their discretization for image deblurring and super-resolution*, in Advanced Techniques in Optimization for Machine Learning and Imaging, A. Benfenati, F. Porta, T. A. Bubba, and M. Viola, eds., Singapore, 2024, Springer Nature Singapore, pp. 55–81.
- [11] Y. CARMON, J. C. DUCHI, O. HINDER, AND A. SIDFORD, *"convex until proven guilty": Dimension-free acceleration of gradient descent on non-convex functions*, in International Conference on Machine Learning, 2017.
- [12] C. CARTIS, J. FIALA, B. MARTEAU, AND L. ROBERTS, *Improving the flexibility and robustness of model-based derivative-free optimization solvers*, ACM Trans. Math. Softw.,

- 45 (2019).
- [13] C. CARTIS AND L. ROBERTS, *A derivative-free Gauss–Newton method*, Mathematical Programming Computation, 11 (2017), pp. 631 – 674.
 - [14] A. CHAMBOLLE AND T. POCK, *An introduction to continuous optimization for imaging*, Acta Numerica, 25 (2016), pp. 161–319.
 - [15] X. CHEN, T. XIAO, AND K. BALASUBRAMANIAN, *Optimal algorithms for stochastic bilevel optimization under relaxed smoothness conditions*, Journal of Machine Learning Research, 25 (2024), pp. 1–51.
 - [16] Y. CHEN, R. RANFTL, AND T. POCK, *Insights into analysis operator learning: From patch-based sparse models to higher order MRFs*, IEEE Transactions on Image Processing, 23 (2014), pp. 1060–1072.
 - [17] C. CROCKETT AND J. A. FESSLER, *Bilevel methods for image reconstruction*, Foundations and Trends® in Signal Processing, 15 (2022), pp. 121–289.
 - [18] S. DEMPE AND A. ZEMKOHO, eds., *Bilevel Optimization: Advances and Next Challenges*, vol. 161 of Springer Optimization and Its Applications, Springer International Publishing, Cham, 2020.
 - [19] S. DITTMER, T. KLUTH, P. MAASS, AND D. O. BAGUER, *Regularization by architecture: A deep prior approach for inverse problems*, Journal of Mathematical Imaging and Vision, 62 (2019), pp. 456–470.
 - [20] S. DOWNING, S. GAZZOLA, I. G. GRAHAM, AND E. A. SPENCE, *Optimising seismic imaging design parameters via bilevel learning*, Inverse Problems, 40 (2024), p. 115008.
 - [21] M. J. EHRHARDT AND L. ROBERTS, *Inexact derivative-free optimization for bilevel learning*, J. Math. Imaging Vis., 63 (2021), pp. 580–600.
 - [22] M. J. EHRHARDT AND L. ROBERTS, *Analyzing inexact hypergradients for bilevel learning*, IMA Journal of Applied Mathematics, 89 (2023), pp. 254–278.
 - [23] C. FAN, G. CHONÉ-DUCASSE, M. SCHMIDT, AND C. THRAMPOULIDIS, *BiSLS/SPS: Auto-tune step sizes for stable bi-level optimization*, in Thirty-seventh Conference on Neural Information Processing Systems, 2023.
 - [24] P. I. FRAZIER, *A tutorial on bayesian optimization*, 2018, <https://arxiv.org/abs/1807.02811>.
 - [25] L. GALLI, H. RAUHUT, AND M. SCHMIDT, *Don't be so monotone: Relaxing stochastic line search in over-parameterized models*, in Thirty-seventh Conference on Neural Information Processing Systems, 2023.
 - [26] S. GHADIMI AND M. WANG, *Approximation methods for bilevel programming*, 2018, <https://arxiv.org/abs/1802.02246>.
 - [27] A. GOUJON, S. NEUMAYER, P. BOHRA, S. DUCOTTERD, AND M. A. UNSER, *A neural-network-based convex regularizer for inverse problems*, IEEE Transactions on Computational Imaging, 9 (2022), pp. 781–795.
 - [28] R. GRAZZI, L. FRANCESCHI, M. PONTIL, AND S. SALZO, *On the iteration complexity of hypergradient computation*, in Proceedings of the 37th International Conference on Machine Learning, H. D. III and A. Singh, eds., vol. 119 of Proceedings of Machine Learning Research, PMLR, 13–18 Jul 2020, pp. 3748–3758.
 - [29] A. GRIEWANK AND C. FAURE, *Piggyback differentiation and optimization*, in Large-Scale PDE-Constrained Optimization, L. T. Biegler, M. Heinkenschloss, O. Ghattas, and B. van

- Bloemen Waanders, eds., Berlin, Heidelberg, 2003, Springer Berlin Heidelberg, pp. 148–164.
- [30] M. HONG, H.-T. WAI, Z. WANG, AND Z. YANG, *A two-timescale stochastic algorithm framework for bilevel optimization: Complexity analysis and application to actor-critic*, SIAM Journal on Optimization, 33 (2023), pp. 147–180.
- [31] F. HUANG, J. LI, AND S. GAO, *Biadam: Fast adaptive bilevel optimization methods*, 2023, <https://arxiv.org/abs/2106.11396>.
- [32] K. JI, J. YANG, AND Y. LIANG, *Bilevel optimization: Convergence analysis and enhanced design*, in Proceedings of the 38th International Conference on Machine Learning, M. Meila and T. Zhang, eds., vol. 139 of Proceedings of Machine Learning Research, PMLR, 18–24 Jul 2021, pp. 4882–4892.
- [33] K. KUNISCH AND T. POCK, *A bilevel optimization approach for parameter learning in variational models*, SIAM J. Imaging Sci., 6 (2013), pp. 938–983.
- [34] J. KWON, D. KWON, S. WRIGHT, AND R. D. NOWAK, *A fully first-order method for stochastic bilevel optimization*, in Proceedings of the 40th International Conference on Machine Learning, A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett, eds., vol. 202 of Proceedings of Machine Learning Research, PMLR, 23–29 Jul 2023, pp. 18083–18113.
- [35] J. KWON, D. KWON, S. WRIGHT, AND R. D. NOWAK, *On penalty methods for nonconvex bilevel optimization and first-order stochastic approximation*, in The Twelfth International Conference on Learning Representations, 2024.
- [36] J. LI, B. GU, AND H. HUANG, *A fully single loop algorithm for bilevel optimization without hessian inverse*, Proceedings of the AAAI Conference on Artificial Intelligence, 36 (2022), pp. 7426–7434.
- [37] B. LIU, M. YE, S. WRIGHT, P. STONE, AND QIANG LIU, *BOME! bilevel optimization made easy: A simple first-order approach*, in Advances in Neural Information Processing Systems, A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, eds., 2022.
- [38] D. MACLAURIN, D. DUVENAUD, AND R. ADAMS, *Gradient-based hyperparameter optimization through reversible learning*, in Proceedings of the 32nd International Conference on Machine Learning, F. Bach and D. Blei, eds., vol. 37 of Proceedings of Machine Learning Research, Lille, France, 07–09 Jul 2015, PMLR, pp. 2113–2122.
- [39] S. MEHMOOD AND P. OCHS, *Automatic differentiation of some first-order methods in parametric optimization*, in Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics, S. Chiappa and R. Calandra, eds., vol. 108 of Proceedings of Machine Learning Research, PMLR, 26–28 Aug 2020, pp. 1584–1594.
- [40] S. MUKHERJEE, S. DITTMER, Z. SHUMAYLOV, S. LUNZ, O. ÖKTEM, AND C.-B. SCHÖNLIEB, *Data-driven convex regularizers for inverse problems*, in ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2024, pp. 13386–13390.
- [41] K. P. MURPHY, *Machine Learning: A Probabilistic Perspective*, The MIT Press, 2012.
- [42] Y. NESTEROV, *Lectures on Convex Optimization*, Springer Publishing Company, Incorporated, 2nd ed., 2018.
- [43] J. NOCEDAL AND S. J. WRIGHT, *Numerical Optimization*, Springer Series in Operations Research and Financial Engineering, Springer, New York, 2 ed., 2006.
- [44] P. OCHS, R. RANFTL, T. BROX, AND T. POCK, *Techniques for gradient-based bilevel*

- optimization with non-smooth lower level problems*, Journal of Mathematical Imaging and Vision, 56 (2016), pp. 175–194.
- [45] F. PEDREGOSA, *Hyperparameter optimization with approximate gradient*, in Proceedings of The 33rd International Conference on Machine Learning, M. F. Balcan and K. Q. Weinberger, eds., vol. 48 of Proceedings of Machine Learning Research, New York, New York, USA, 20–22 Jun 2016, PMLR, pp. 737–746.
- [46] Z. RAMZI, F. MANNEL, S. BAI, J.-L. STARCK, P. CIUCIU, AND T. MOREAU, *SHINE: SHaring the INverse estimate from the forward pass for bi-level optimization and implicit models*, in International Conference on Learning Representations, 2022.
- [47] J. C. D. L. REYES AND D. VILLACÍS, *Bilevel Optimization Methods in Imaging*, Springer International Publishing, Cham, 2023, pp. 909–941.
- [48] L. I. RUDIN, S. OSHER, AND E. FATEMI, *Nonlinear total variation based noise removal algorithms*, Physica D: Nonlinear Phenomena, 60 (1992), pp. 259–268.
- [49] F. SHERRY, M. BENNING, J. C. DE LOS REYES, M. J. GRAVES, G. MAIERHOFER, G. WILLIAMS, C.-B. SCHÖNLIEB, AND M. J. EHRHARDT, *Learning the sampling pattern for MRI*, IEEE Transactions on Medical Imaging, 39 (2020), pp. 4310–4321.
- [50] H.-J. M. SHI, Y. XIE, R. BYRD, AND J. NOCEDAL, *A noise-tolerant Quasi-Newton algorithm for unconstrained optimization*, SIAM Journal on Optimization, 32 (2022), pp. 29–55.
- [51] E. SUONPERÄ AND T. VALKONEN, *Linearly convergent bilevel optimization with single-step inner methods*, 2023, <https://arxiv.org/abs/2205.04862>.
- [52] S. VASWANI, A. MISHKIN, I. H. LARADJI, M. SCHMIDT, G. GIDEL, AND S. LACOSTE-JULIEN, *Painless stochastic gradient: Interpolation, line-search, and convergence rates*, in Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada, H. M. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. A. Fox, and R. Garnett, eds., 2019, pp. 3727–3740.
- [53] H. YANG, L. LUO, C. J. LI, M. JORDAN, AND M. FAZEL, *Accelerating inexact hypergradient descent for bilevel optimization*, in OPT 2023: Optimization for Machine Learning, 2023.
- [54] H. ZHANG AND W. W. HAGER, *A nonmonotone line search technique and its application to unconstrained optimization*, SIAM J. Optim., 14 (2004), pp. 1043–1056.
- [55] N. ZUCCHET AND J. SACRAMENTO, *Beyond backpropagation: Bilevel optimization through implicit differentiation and equilibrium propagation*, Neural Computation, 34 (2022), pp. 2309–2346.