



HAL
open science

Second-order step-size tuning of SGD for non-convex optimization

Camille Castera, Cédric Févotte, Jérôme Bolte, Edouard Pauwels

► **To cite this version:**

Camille Castera, Cédric Févotte, Jérôme Bolte, Edouard Pauwels. Second-order step-size tuning of SGD for non-convex optimization. *Neural Processing Letters*, 2022, pp.1–26. 10.1007/s11063-021-10705-5 . hal-03161775v2

HAL Id: hal-03161775

<https://hal.science/hal-03161775v2>

Submitted on 23 Nov 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Second-order step-size tuning of SGD for non-convex optimization

Camille Castera*

CNRS - IRIT,
Université de Toulouse,
Toulouse, France

Jérôme Bolte†

Toulouse School of Economics
Université de Toulouse
Toulouse, France

Cédric Févotte†

CNRS - IRIT,
Université de Toulouse,
Toulouse, France

Edouard Pauwels†

CNRS - IRIT,
Université de Toulouse,
DEEL, IRT Saint Exupery
Toulouse, France

Abstract

In view of a direct and simple improvement of vanilla SGD, this paper presents a fine-tuning of its step-sizes in the mini-batch case. For doing so, one estimates curvature, based on a local quadratic model and using only noisy gradient approximations. One obtains a new stochastic first-order method (Step-Tuned SGD), enhanced by second-order information, which can be seen as a stochastic version of the classical Barzilai-Borwein method. Our theoretical results ensure almost sure convergence to the critical set and we provide convergence rates. Experiments on deep residual network training illustrate the favorable properties of our approach. For such networks we observe, during training, both a sudden drop of the loss and an improvement of test accuracy at medium stages, yielding better results than SGD, RMSprop, or ADAM.

1 Introduction

In the recent years, machine learning has generated a growing need for methods to solve non-convex optimization problems. In particular, the training of deep neural networks (DNNs) has received tremendous attention. Designing methods for this purpose is particularly difficult as one deals with both expensive function evaluations and limited storage capacities. This explains why stochastic gradient descent (SGD) remains the central algorithm in deep learning (DL). It consists in the iterative scheme,

$$\theta_{k+1} = \theta_k - \gamma_k \nabla \mathcal{J}_{\mathcal{B}_k}(\theta_k), \quad (1)$$

where \mathcal{J} is the function to minimize (usually the empirical loss) parameterized by $\theta \in \mathbb{R}^P$ (the weights of the DNN), $\nabla \mathcal{J}_{\mathcal{B}_k}(\theta_k)$ is a stochastic estimation of the gradient of \mathcal{J} (randomness

*Corresponding author: camille.castera@protonmail.com

†Last three authors are listed in alphabetical order.

being related to the sub-sampled mini-batch B_k), and $\gamma_k > 0$ is a step-size whose choice is critical in terms of empirical performance.

In order to improve the basic SGD method, a common practice is to use adaptive methods [16, 24, 43]. They act as preconditioners, reducing the importance of the choice of the sequence of step-sizes $(\gamma_k)_{k \in \mathbb{N}}$. This paper focuses instead exclusively on the step-size issue: how can we take advantage of curvature information of non-convex landscapes in a stochastic context in order to design step-sizes adapted to each iteration?

Our starting point to answer this question is an infinitesimal second-order variational model along the gradient direction. The infinitesimal feature is particularly relevant in DL since small steps constitute standard practice in training due to sub-sampling noise. Second-order information is approximated with first-order quantities using finite differences. In deterministic (full-batch) setting, our method corresponds to a non-convex version of the Barzilai-Borwein (BB) method [6, 8, 14, 46] and is somehow a discrete non-convex adaption of the continuous gradient system in Alvarez and Cabot [3]. It is also close to earlier work [35], with the major difference that our algorithm is supported by a variational model. This is essential to generalize the method to accommodate noisy gradients, providing a convexity test similar to those in Babaie-Kafaki and Fatemi [4], Curtis and Guo [12].

Our main contribution is a fine step-size tuning method which accelerates stochastic gradient algorithms, it is built on a strong geometrical principle: step-sizes are deduced from a carefully derived discrete approximation of a curvature term of the expected loss. We provide general convergence guarantees to critical points and rates of convergence. Extensive computations on DL problems show that our method is particularly successful for residual networks [18]. In that case, we observe a surprising phenomenon: in the early training stage the method shows similar performances to standard DL algorithms (SGD, ADAM, RMSprop), then at medium stage, we observe simultaneously a sudden drop of the training loss and a notable increase in test accuracy.

To summarize, our contributions are as follows:

- Exploit the vanishing step-size nature of DL training to use infinitesimal second-order optimization for fine-tuning the step-size at each iteration.
- Use our geometrical perspective to discretize and adapt the method to noisy gradients despite strong non-linearities.
- Prove the convergence of the proposed algorithm and provide rates of convergence for twice-differentiable non-convex functions either under boundedness assumption or under Lipschitz-continuity assumptions (see Theorem 1 and Corollary 2).
- Show that our method has remarkable practical properties, in particular when training residual networks in DL, for which one observes an advantageous “drop down” of the loss during the training phase.

Structure of the paper. A preliminary deterministic (i.e., full-batch) algorithm is derived in Section 3.1. We then build a stochastic mini-batch variant in Section 3.2, which is our core contribution. Theoretical results are stated in Section 4 and DL experiments are conducted in Section 5.

2 Related work

Methods using second-order information for non-convex optimization have been actively studied in the last years, both for deterministic and stochastic applications, see, e.g., Allen-Zhu [2],

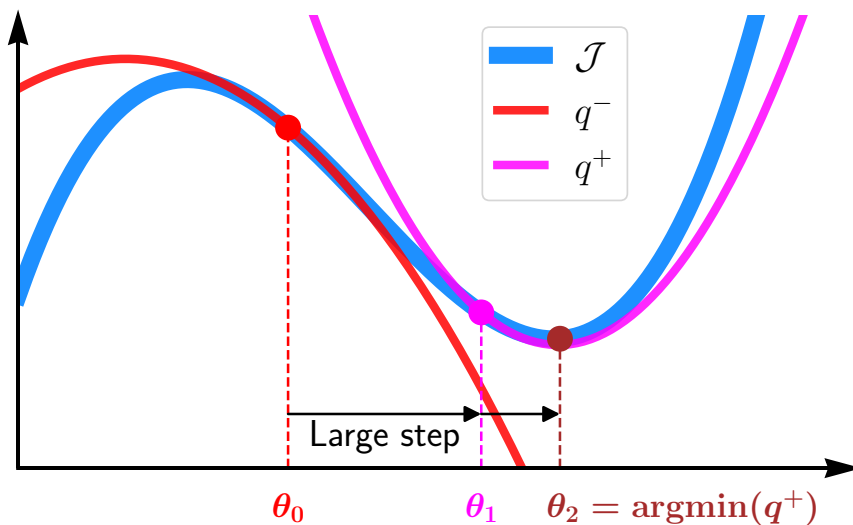


Figure 1: Illustration of negative and positive curvature steps. The function q^- represents the variational model at θ_0 , with negative curvature. Concavity suggests to take a large step to reach θ_1 . Then, at θ_1 , the variational model q^+ has positive curvature and can be minimized to obtain θ_2 .

Carmon et al. [10], Curtis and Robinson [13], Krishnan et al. [25], Liu and Yang [32], Martens and Grosse [33], Royer and Wright [40].

BB-like methods are very sensitive to noisy gradient estimates. Most existing stochastic BB algorithms [30, 38, 42] overcome this issue with stabilization methods in the style of SVRG [23], which allows to prescribe a new step-size at every epoch only (i.e., after a full pass over the data). Doing so, one cannot capture variations of curvature within a full epoch, and one is limited to using absolute values to prevent negative step-sizes caused by non-convexity. On the contrary, our stochastic approximation method can adapt to local curvature every two iterations. Regarding the utilization of flatness and concavity of DL loss functions, the AdaBelief algorithm of Zhuang et al. [47] is worth mentioning. The latter uses the difference between the current stochastic gradient estimate and an average of past gradients, this difference being used to prescribe a vector step-size in the style of ADAM. In comparison, our method uses scalar step-sizes and aims to capture subtle variations as it computes a stabilized difference between consecutive gradient estimates *before* averaging.

There are few techniques to analyze stochastic methods in non-convex settings. An important category is the ODE machinery used for SGD [9, 15], ADAM [5] and INNA [11]. In this paper, we use instead direct and more traditional arguments, such as in Li and Orabona [29] in the context of DL.

3 Design of the algorithm

We first build a preliminary algorithm based upon a simple second-order variational model. We then adapt this algorithm to address mini-batch stochastic approximations.

3.1 Deterministic full-batch algorithm

We start with the following variational considerations.

3.1.1 Second-order infinitesimal step-size tuning.

Given a positive integer P , assume that \mathcal{J} is a twice-differentiable function, denote $\nabla\mathcal{J}$ and $\nabla^2\mathcal{J}$ the gradient and the Hessian matrix of \mathcal{J} respectively. Let $\theta \in \mathbb{R}^P$. Given an update direction $d \in \mathbb{R}^P$, a natural strategy is to choose $\gamma \in \mathbb{R}$ that minimizes $\mathcal{J}(\theta + \gamma d)$. Let us approximate $\gamma \mapsto \mathcal{J}(\theta + \gamma d)$ around 0 with a Taylor expansion,

$$q_d(\gamma) \stackrel{\text{def}}{=} \mathcal{J}(\theta) + \gamma \langle \nabla\mathcal{J}(\theta), d \rangle + \frac{\gamma^2}{2} \langle \nabla^2\mathcal{J}(\theta)d, d \rangle. \quad (2)$$

If the curvature term $\langle \nabla^2\mathcal{J}(\theta)d, d \rangle$ is positive, then q_d has a unique minimizer at,

$$\gamma^* = -\frac{\langle \nabla\mathcal{J}(\theta), d \rangle}{\langle \nabla^2\mathcal{J}(\theta)d, d \rangle}. \quad (3)$$

On the contrary when $\langle \nabla^2\mathcal{J}(\theta)d, d \rangle \leq 0$, the infinitesimal model q_d is concave (or equivalently \mathcal{J} is locally concave in the direction d) which suggests to take a large step-size. These considerations are illustrated on Figure 1.

3.1.2 Tuning gradient descent.

In order to tune gradient descent we choose the direction $d = -\nabla\mathcal{J}(\theta)$ which gives,

$$\gamma(\theta) \stackrel{\text{def}}{=} \frac{\|\nabla\mathcal{J}(\theta)\|^2}{\langle \nabla^2\mathcal{J}(\theta)\nabla\mathcal{J}(\theta), \nabla\mathcal{J}(\theta) \rangle}. \quad (4)$$

According to our previous considerations, an ideal iterative process $\theta_{k+1} = \theta_k - \gamma_k \nabla\mathcal{J}(\theta_k)$ would use $\gamma_k = \gamma(\theta_k)$ when $\gamma(\theta_k) > 0$. But for computational reasons and discretization purposes, we shall rather seek a step-size γ_k such that, $\gamma_k \simeq \gamma(\theta_{k-1})$ (again when $\gamma(\theta_{k-1}) > 0$). Let us assume that, for $k \geq 1$, $\theta_{k-1}, \gamma_{k-1}$ are known and let us approximate the quantity,

$$\gamma(\theta_{k-1}) = \frac{\|\nabla\mathcal{J}(\theta_{k-1})\|^2}{\langle \nabla^2\mathcal{J}(\theta_{k-1})\nabla\mathcal{J}(\theta_{k-1}), \nabla\mathcal{J}(\theta_{k-1}) \rangle}, \quad (5)$$

using only first-order objects. We rely on two identities,

$$\Delta\theta_k \stackrel{\text{def}}{=} \theta_k - \theta_{k-1} = -\gamma_{k-1} \nabla\mathcal{J}(\theta_{k-1}), \quad (6)$$

$$\Delta g_k \stackrel{\text{def}}{=} \nabla\mathcal{J}(\theta_k) - \nabla\mathcal{J}(\theta_{k-1}) \simeq -\gamma_{k-1} \mathcal{C}_{\mathcal{J}}(\theta_{k-1}), \quad (7)$$

where $\mathcal{C}_{\mathcal{J}}(\theta) \stackrel{\text{def}}{=} \nabla^2\mathcal{J}(\theta)\nabla\mathcal{J}(\theta)$ and (7) is obtained by Taylor's formula. Combining the above, we are led to consider the following step-size,

$$\gamma_k = \begin{cases} \frac{\|\Delta\theta_k\|^2}{\langle \Delta\theta_k, \Delta g_k \rangle} & \text{if } \langle \Delta\theta_k, \Delta g_k \rangle > 0 \\ \nu & \text{otherwise} \end{cases}, \quad (8)$$

where $\nu > 0$ is an hyper-parameter of the algorithm representing the large step-sizes to use in locally concave regions.

Algorithm 1 Full-batch preliminary algorithm

```
1: Input:  $\alpha > 0, \nu > 0$ 
2: Initialize  $\theta_0 \in \mathbb{R}^P$ 
3:  $\theta_1 = \theta_0 - \alpha \nabla \mathcal{J}(\theta_0)$ 
4: for  $k = 1, \dots$  do
5:    $\Delta g_k = \nabla \mathcal{J}(\theta_k) - \nabla \mathcal{J}(\theta_{k-1})$ 
6:    $\Delta \theta_k = \theta_k - \theta_{k-1}$ 
7:   if  $\langle \Delta g_k, \Delta \theta_k \rangle > 0$  then
8:      $\gamma_k = \frac{\|\Delta \theta_k\|^2}{\langle \Delta g_k, \Delta \theta_k \rangle}$ 
9:   else
10:     $\gamma_k = \nu$ 
11:   end if
12:    $\theta_{k+1} = \theta_k - \alpha \gamma_k \nabla \mathcal{J}(\theta_k)$ 
13: end for
```

The resulting full-batch non-convex optimization method is Algorithm 1, in which α is a so-called *learning rate* or *scaling factor*. This algorithm is present in the literature under subtle variants [8, 14, 35, 46]. It may be seen as a non-convex version of the BB method (originally designed for strongly convex functions) as it coincides with the BB step-size when $\langle \Delta \theta_k, \Delta g_k \rangle$ is positive. In the classical optimization literature the BB step-size and its variants are often combined with line-search procedures which is impossible in our large-scale DL context. This is why we replace the line-search by a scaling factor α , present in most DL optimizers and which generally requires tuning. Our purpose is not however to get rid of hyper-parameters pre-tuning but rather to combine this α with an automatic fine tuning able to capture the local variations in \mathcal{J} . Although Algorithm 1 is close to existing methods, the interest of our variational viewpoint is the characterization of the underlying geometrical mechanism supporting the algorithm, which is key in designing an efficient stochastic version of Algorithm 1 in Section 3.2.

3.1.3 Illustrative experiment.

Before presenting the stochastic version, we illustrate the interest of exploiting negative curvature through the *large-step* parameter ν with a synthetic experiment inspired from Carmon et al. [10]. We apply Algorithm 1 to a non-convex regression problem of the form $\min_{\theta \in \mathbb{R}^P} \phi(A\theta - b)$ where ϕ is a non-convex real-valued function (see Section C of the Supplementary). We compare Algorithm 1 with the methods à la BB where absolute values are used when the step-size is negative¹ (see, e.g., Liang et al. [30], Tan et al. [42] in stochastic settings) and with Armijo’s line-search gradient method. As shown on Figure 2, Algorithm 1 efficiently exploits local curvature and converges much faster than other methods.

3.2 Stochastic mini-batch algorithm

We wish to adapt Algorithm 1 in cases where gradients can only be approximated through mini-batch sub-sampling. This is necessary in particular for DL applications.

¹For a fair comparison we implement this method with the scaling-factor α of Algorithm 1.

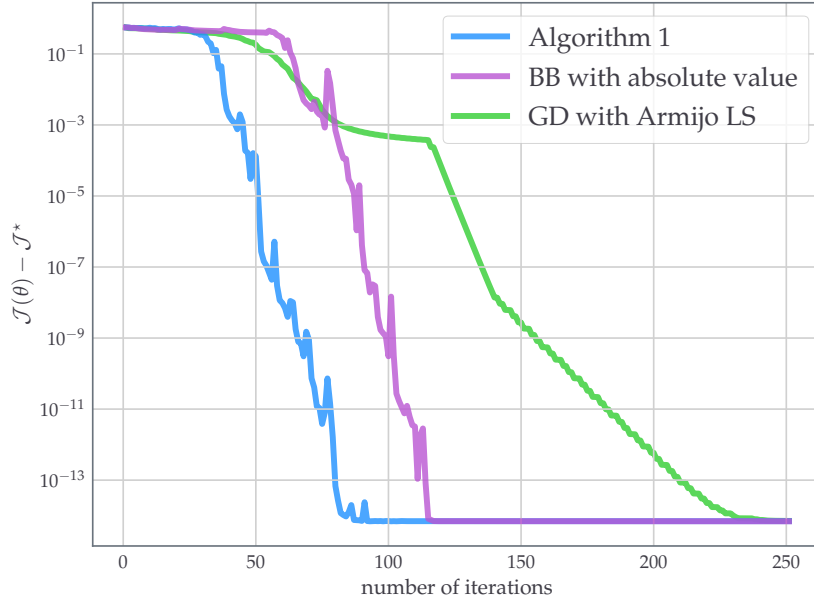


Figure 2: Values of the loss function $\mathcal{J}(\theta)$ against iterations (each corresponding to a gradient step) for the synthetic non-convex regression problem detailed in Section C of the Supplementary. The optimal value \mathcal{J}^* is unknown and is estimated by taking the best value obtained among all algorithms after 10^5 iterations.

3.2.1 Mini-batch sub-sampling.

We assume the following sum-structure of the loss function, for $N \in \mathbb{N}_{>0}$,

$$\mathcal{J}(\theta) = \frac{1}{N} \sum_{n=1}^N \mathcal{J}_n(\theta), \quad (9)$$

where each \mathcal{J}_n is a twice continuously-differentiable function. Given any fixed subset $\mathbf{B} \subset \{1, \dots, N\}$, we define the following quantities for any $\theta \in \mathbb{R}^P$,

$$\mathcal{J}_{\mathbf{B}}(\theta) \stackrel{\text{def}}{=} \frac{1}{|\mathbf{B}|} \sum_{n \in \mathbf{B}} \mathcal{J}_n(\theta), \quad (10)$$

$$\nabla \mathcal{J}_{\mathbf{B}}(\theta) \stackrel{\text{def}}{=} \frac{1}{|\mathbf{B}|} \sum_{n \in \mathbf{B}} \nabla \mathcal{J}_n(\theta), \quad (11)$$

where $|\mathbf{B}|$ denotes the number of elements of the set \mathbf{B} . Throughout this paper we will consider independent copies of a random subset $\mathbf{S} \subset \{1, \dots, N\}$ referred to as mini-batch. The distribution of this subset is fixed throughout the paper and taken such that the expectation over the realization of \mathbf{S} in (10) corresponds to the empirical expectation in (9). This is valid for example if \mathbf{S} is taken uniformly at random over all possible subsets of fixed size. As a consequence, we have the identity $\mathcal{J} = \mathbb{E}[\mathcal{J}_{\mathbf{S}}]$, where \mathbb{E} denotes the expectation operator, here taken over the random draw of \mathbf{S} . This allows to interpret mini-batch sub-sampling as a stochastic approximation process since we also have $\nabla \mathcal{J} = \mathbb{E}[\nabla \mathcal{J}_{\mathbf{S}}]$.

Our algorithm is of stochastic gradient type where stochasticity is related to the randomness of mini-batches. We start with an initialization $\theta_0 \in \mathbb{R}^P$, and a sequence of i.i.d. random mini-batches $(\mathbf{B}_k)_{k \in \mathbb{N}}$, whose common distribution is the same as \mathbf{S} . The algorithm produces a random sequence of iterates $(\theta_k)_{k \in \mathbb{N}}$. For $k \in \mathbb{N}$, \mathbf{B}_k is used to estimate an update direction $-\nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k)$ which is used in place of $-\nabla \mathcal{J}(\theta_k)$ in the same way as gradient descent algorithm.

Algorithm 2 Step-Tuned SGD

```

1: Input:  $\alpha > 0, \nu > 0$ 
2: Input:  $\beta \in [0, 1], \tilde{m} > 0, \tilde{M} > 0, \delta \in (0, 1/2)$ 
3: Initialize  $\theta_0 \in \mathbb{R}^P, G_{-1} = \mathbf{0}_P, \gamma_0 = 1$ 
4: Draw independent random mini-batches  $(\mathbf{B}_k)_{k \in \mathbb{N}}$ .
5: for  $k = 0, 1, \dots$  do
6:    $\theta_{k+\frac{1}{2}} = \theta_k - \frac{\alpha}{(k+1)^{1/2+\delta}} \gamma_k \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k)$ 
7:    $\theta_{k+1} = \theta_{k+\frac{1}{2}} - \frac{\alpha}{(k+1)^{1/2+\delta}} \gamma_k \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_{k+\frac{1}{2}})$ 
8:    $\Delta \theta_{\mathbf{B}_k} = \theta_{k+\frac{1}{2}} - \theta_k$ 
9:    $\Delta g_{\mathbf{B}_k} = \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_{k+\frac{1}{2}}) - \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k)$ 
10:   $G_k = \beta G_{k-1} + (1 - \beta) \Delta g_{\mathbf{B}_k}$ 
11:   $\hat{G}_k = G_k / (1 - \beta^{k+1})$ 
12:  if  $\langle \hat{G}_k, \Delta \theta_{\mathbf{B}_k} \rangle > 0$  then
13:     $\gamma_{k+1} = \frac{\|\Delta \theta_{\mathbf{B}_k}\|^2}{\langle \hat{G}_k, \Delta \theta_{\mathbf{B}_k} \rangle}$ 
14:  else
15:     $\gamma_{k+1} = \nu$ 
16:  end if
17:   $\gamma_{k+1} = \min(\max(\gamma_{k+1}, \tilde{m}), \tilde{M})$ 
18: end for

```

3.2.2 Second-order tuning of mini-batch SGD: Step-Tuned SGD.

Our goal is to devise a step-size strategy, based on the variational ideas developed earlier and on the quantity $\mathcal{C}_{\mathcal{J}}$, in the context of mini-batch sub-sampling. First observe that for $\theta \in \mathbb{R}^P$,

$$\mathcal{C}_{\mathcal{J}}(\theta) = \nabla^2 \mathcal{J}(\theta) \nabla \mathcal{J}(\theta) = \nabla \left(\frac{1}{2} \|\nabla \mathcal{J}(\theta)\|^2 \right). \quad (12)$$

So rewriting \mathcal{J} as an expectation,

$$\mathcal{C}_{\mathcal{J}}(\theta) = \nabla \left(\frac{1}{2} \|\mathbb{E}[\nabla \mathcal{J}_{\mathbf{S}}(\theta)]\|^2 \right), \quad (13)$$

where \mathbf{S} denotes like in the previous paragraph a random subset of $\{1, \dots, N\}$, or mini-batch. This suggests the following estimator,

$$\mathcal{C}_{\mathcal{J}_{\mathbf{B}}}(\theta) \stackrel{\text{def}}{=} \nabla \left(\frac{1}{2} \|\nabla \mathcal{J}_{\mathbf{B}}(\theta)\|^2 \right) = \nabla^2 \mathcal{J}_{\mathbf{B}}(\theta) \nabla \mathcal{J}_{\mathbf{B}}(\theta), \quad (14)$$

to build an infinitesimal model as in (4), where $\mathbf{B} \subset \{1, \dots, N\}$ and $\theta \in \mathbb{R}^P$.

Like in the deterministic case we approximate the new target (14) with a Taylor expansion of $\mathcal{J}_{\mathbf{B}}$ between two iterations. We obtain for any $\mathbf{B} \subset \{1, \dots, N\}$, $\theta \in \mathbb{R}^P$, and small $\gamma > 0$

$$-\gamma \mathcal{C}_{\mathcal{J}_{\mathbf{B}}}(\theta) \simeq \nabla \mathcal{J}_{\mathbf{B}}(\underbrace{\theta - \gamma \nabla \mathcal{J}_{\mathbf{B}}(\theta)}_{\text{next iterate}}) - \nabla \mathcal{J}_{\mathbf{B}}(\theta). \quad (15)$$

This suggests to use each mini-batch twice and compute a difference of gradients every two iterations.² We adopt the following convention, at iteration $k \in \mathbb{N}$, the random mini-batch \mathbf{B}_k

²There is also the possibility of computing additional estimates as Schraudolph et al. [41] previously did for a stochastic BFGS algorithm, but this would double the computational cost.

is used to compute a stochastic gradient, $\nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k)$ and at iteration $k + \frac{1}{2}$ the same mini-batch is used to compute another stochastic gradient $\nabla \mathcal{J}_{\mathbf{B}_k}(\theta_{k+\frac{1}{2}})$, for a given $\theta_{k+\frac{1}{2}}$. Let us define,

$$\Delta g_{\mathbf{B}_k} \stackrel{\text{def}}{=} \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_{k+\frac{1}{2}}) - \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k), \quad (16)$$

thereby $\Delta g_{\mathbf{B}_k}$ forms an approximation of $-\gamma_k \mathcal{C}_{\mathcal{J}_{\mathbf{B}_k}}(\theta_k)$ that we can use to compute the next step-size γ_{k+1} . We define the difference between two iterates accordingly,

$$\Delta \theta_{\mathbf{B}_k} \stackrel{\text{def}}{=} \theta_{k+\frac{1}{2}} - \theta_k. \quad (17)$$

Finally, we stabilize the approximation of the target quantity in (14) by using an exponential moving average of the previously computed $(\Delta g_{\mathbf{B}_j})_{j \leq k}$. More precisely, we recursively compute G_k defined by,

$$G_k = \beta G_{k-1} + (1 - \beta) \Delta g_{\mathbf{B}_k}. \quad (18)$$

We finally introduce $\hat{G}_k = G_k / (1 - \beta^{k+1})$ to debias the estimator G_k such that the sum of the weights in the average equals 1. This mostly impacts the first iterations as β^{k+1} vanishes quickly; a similar process is used in ADAM [24].

Altogether we obtain our main method: Algorithm 2, which we name *Step-Tuned SGD*, as it aims to tune the step-size every two iterations and not at every epoch like most stochastic BB methods. Note that the main idea behind Step-Tuned SGD remains the same than in the deterministic setting: we exploit the curvature properties of $\mathcal{J}_{\mathbf{B}_k}$ through the quantities $\langle \hat{G}_k, \Delta \theta_{\mathbf{B}_k} \rangle$ to devise our method. Compared to Algorithm 1, the iteration index is shifted by 1 so that the estimated step-size γ_{k+1} only depends on mini-batches \mathbf{B}_0 up to \mathbf{B}_k and is therefore conditionally independent of \mathbf{B}_{k+1} . This conditional dependency structure is crucial to obtain the convergence guarantees given in Section 4.

Remark 1. *It is worth precising that like most methods, Algorithm 2 does not alleviate the need of tuning the scaling factor α . The choice of this parameter remains indeed important in most practical applications. The purpose of Step-Tuned SGD is rather to speed up SGD by fine-tuning the step-size at each iteration (through the introduction of γ_k). This is analogous to the BB methods for deterministic applications which often accelerate algorithms but must be stabilized with line-search strategies. To the best of our knowledge, in comparison to the epoch-wise BB methods [30, 42], Algorithm 2 is the first method that manages to mimic the iteration-wise behavior of deterministic BB methods for mini-batch applications.*

3.3 Heuristic construction of Step-Tuned SGD

In this section we present the main elements which led us to the step-tuned method of Algorithm 2 and discuss its hyper-parameters. Throughout this paragraph, the term *gradient variation* (GV) denotes the local variations of the gradient; it is simply the difference of consecutive gradients along a sequence. Our heuristic discussion blends discretization arguments and experimental considerations. We use the non-convex regression experiment of Section 3.1 as a test for our intuition and algorithms. A complete description of the methods below is given in Section D, we only sketch the main ideas.

3.3.1 First heuristic experiment with exact GVs.

Assume that along any ordered collection $\theta_1, \dots, \theta_k \in \mathbb{R}^P$, one is able to evaluate the GVs of \mathcal{J} , that is, terms of the form $\nabla \mathcal{J}(\theta_i) - \nabla \mathcal{J}(\theta_{i-1})$. Recall that we denote $\Delta \theta_i = \theta_i - \theta_{i-1}$, the

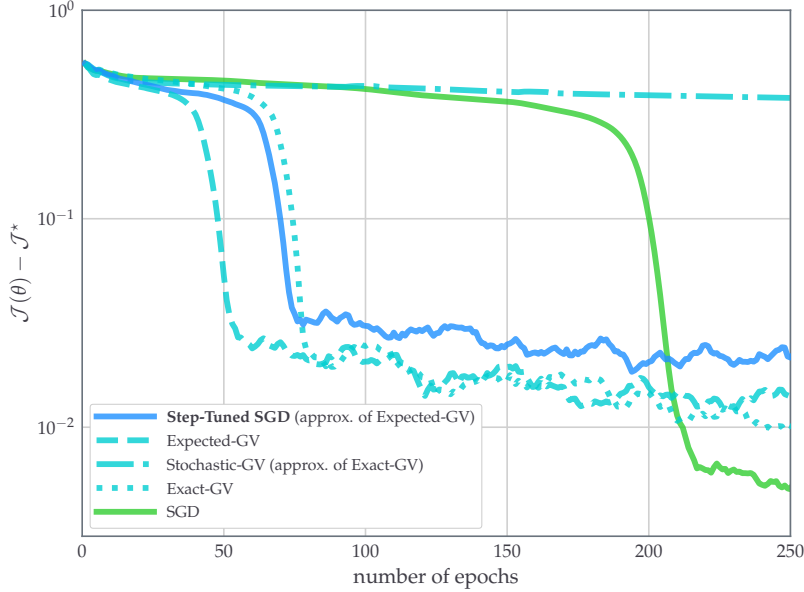


Figure 3: Values of the loss function against epochs for non-convex regression: heuristic methods (dashed lines) of Section 3.3 are compared with Step-Tuned SGD (plain blue). SGD serves as a reference to evidence the fast drop down effect of other methods. The additional computational cost of Expected-GV and Exact-GV is ignored as these methods are here only for illustration purposes (see Section 3.3).

difference between two consecutive iterates, for all $i \geq 1$. In the deterministic (i.e., noiseless) setting, Algorithm 1 is based on these GVs, indeed,

$$\theta_{k+1} = \theta_k - \frac{\alpha \|\Delta\theta_k\|^2}{\langle \nabla \mathcal{J}(\theta_k) - \nabla \mathcal{J}(\theta_{k-1}), \Delta\theta_k \rangle} \nabla \mathcal{J}(\theta_k), \quad (19)$$

whenever the denominator is positive. Given our sequence of independent random mini-batches $(\mathbf{B}_k)_{k \in \mathbb{N}}$, a heuristic stochastic approximation version of this recursion could be as follows,

$$\theta_{k+1} = \theta_k - \frac{\alpha_k \|\Delta\theta_k\|^2}{\langle \nabla \mathcal{J}(\theta_k) - \nabla \mathcal{J}(\theta_{k-1}), \Delta\theta_k \rangle} \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k), \quad (\text{Exact-GV})$$

where the difference between (19) and Exact-GV lies in the randomness of the search direction and the dependency of the scaling factor α_k which aims to moderate the effect of noise (generally $\alpha_k \rightarrow 0$). As shown in Figure 3 the recursion Exact-GV is much faster than SGD especially for the first ~ 150 epochs which is often the main concern for DL applications. Indeed, although SGD achieves a smaller value of \mathcal{J} after a larger number of iterations (due to other methods using larger step-sizes), this happens when the value of \mathcal{J} is already low.³ Overall the quantity Exact-GV seems very promising.

Yet, for large sums, the gradient-variation in Exact-GV is too computationally expensive. One should therefore adapt (Exact-GV) to the mini-batch context. A direct adaption would simply consist in the algorithm,

$$\theta_{k+1} = \theta_k - \frac{\alpha_k \|\Delta\theta_k\|^2}{\langle \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k) - \nabla \mathcal{J}_{\mathbf{B}_{k-1}}(\theta_{k-1}), \Delta\theta_k \rangle} \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k), \quad (\text{Stochastic-GV})$$

³Step-Tuned SGD achieves the same small level of error as SGD when doing additional epochs thanks to the decay schedule present in Algorithm 2.

where mini-batches are used both to obtain a search direction and to approximate the GV. For this naive approach, we observe a dramatic loss of performance, as illustrated in Figure 3. This reveals the necessity to use accurate stochastic approximation of GVs.

3.3.2 Second heuristic experiment using expected gradient variations.

Towards a more stable approximation of the GVs, we consider the following recursion,

$$\theta_{k+1} = \theta_k - \frac{\alpha_k \|\Delta\theta_k\| \|\nabla \mathcal{J}_{\mathbf{B}_{k-1}}(\theta_{k-1})\|}{\langle -\mathbb{E}[\mathcal{C}_{\mathcal{J}_{\mathbf{S}}}(\theta_{k-1})], \Delta\theta_k \rangle} \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k), \quad (\text{Expected-GV})$$

where $\mathcal{C}_{\mathcal{J}_{\mathbf{B}}}$ is defined in (14) for any $\mathbf{B} \subset \{1, \dots, N\}$ and the expectation taken is over the independent draw of $\mathbf{S} \subset \{1, \dots, N\}$, conditioned on the other random variables. The main difference with Exact-GV is the use of expected GVs instead of exact GVs, the minus sign ensures a coherent interpretation in term of GVs. The numerator in Expected-GV is also modified to ensure homogeneity of the steps with the other variations of the algorithm. Indeed $\mathcal{C}_{\mathcal{J}_{\mathbf{B}}}(\theta_k)$ approximates a difference of gradients modulo a step-size, see (15). As illustrated in Figure 3, the recursion Expected-GV provides performances comparable (and even superior) to Exact-GV, and in particular for both algorithms, we also recover the loss drop which was observed in the deterministic setting.

Algorithm 2 is nothing less than an approximate version of Expected-GV which combines a double use of mini-batches with a moving average. Indeed, from (15), considering the expectation over the random draw of \mathbf{S} , for any $\theta \in \mathbb{R}^P$ and small $\gamma > 0$, we have,

$$-\gamma \mathbb{E}[\mathcal{C}_{\mathcal{J}_{\mathbf{S}}}(\theta)] \simeq \mathbb{E}[\nabla \mathcal{J}_{\mathbf{S}}(\theta - \gamma \nabla \mathcal{J}_{\mathbf{S}}(\theta)) - \nabla \mathcal{J}_{\mathbf{S}}(\theta)]. \quad (20)$$

The purpose of the term \hat{G}_k in Algorithm 2 is precisely to mimic this last quantity. The experimental results of Algorithm 2 are very similar to those of Expected-GV, see Figure 3.

Let us conclude by saying that the above considerations on gradient variations (GVs) led us to propose Algorithm 2 as a possible stochastic version of Algorithm 1. The similarity between the performances of the two methods and the underlying geometric aspects (see Section 3.2) were also major motivations.

3.3.3 Parameters of the algorithm.

Algorithm 2 contains more hyper-parameters than in the deterministic case, but we recommend keeping the default values for most of them.⁴ Like in most optimizers (SGD, ADAM, RMSprop, etc.), only the parameter $\alpha > 0$ has to be carefully tuned to get the most of Algorithm 2. The value $\beta = 0.9$ is a common default choice for exponential moving averages (see e.g., Kingma and Ba [24]). Note that we enforce $\gamma_k \in [\tilde{m}, \tilde{M}]$. The bounds stabilize the algorithm and also play an important role for the convergence as we will show in Section 4. While a fine tuning of these bounds may improve the performances, we chose rather tight default values for the sake of numerical stability so that practitioners need not tuning them. The same choice was made for the parameter ν . Note that we also enforce the step-size to decrease using a decay of the form $1/k^{1/2+\delta}$ where the value of $\delta > 0$ is of little importance as long as it is taken close to 0. This standard procedure goes back to Robbins and Monro [36] and is again necessary to obtain the convergence results presented next.

⁴Default values: $(\nu, \beta, \tilde{m}, \tilde{M}, \delta) = (2, 0.9, 0.5, 2, 0.001)$

4 Theoretical results

We study the convergence of Step-Tuned SGD for smooth non-convex stochastic optimization which encompasses in particular smooth DL problems.

4.1 Main result.

We recall that \mathcal{J} is a finite sum of twice continuously-differentiable functions $(\mathcal{J}_n)_{n=1,\dots,N}$. Hence, the gradient of \mathcal{J} and the gradients of each \mathcal{J}_n are locally Lipschitz continuous. A function g is locally Lipschitz continuous on \mathbb{R}^P if for any $\theta \in \mathbb{R}^P$, there exists a neighborhood V of θ and a constant $L \in \mathbb{R}_+$ such that for all $\psi_1, \psi_2 \in V$,

$$\|g(\psi_1) - g(\psi_2)\| \leq L\|\psi_1 - \psi_2\|. \quad (21)$$

We assume that \mathcal{J} is lower-bounded on \mathbb{R}^P , which holds for most DL loss functions by construction (they are usually non-negative). We denote by $\frac{1}{2}\mathbb{N} = \{0, \frac{1}{2}, 1, \frac{3}{2}, 2, \dots\}$ the set of half integers so that the iterations of Step-Tuned SGD are indexed by $k \in \frac{1}{2}\mathbb{N}$. The main theoretical result of this paper follows.

Theorem 1. *Let $\theta_0 \in \mathbb{R}^P$, and let $(\theta_k)_{k \in \frac{1}{2}\mathbb{N}}$ be a sequence generated by Step-Tuned SGD initialized at θ_0 . Assume that there exists a constant $C_1 > 0$ such that almost surely $\sup_{k \in \frac{1}{2}\mathbb{N}} \|\theta_k\| < C_1$. Then the sequence of values $(\mathcal{J}(\theta_k))_{k \in \mathbb{N}}$ converges almost surely and $(\|\nabla \mathcal{J}(\theta_k)\|^2)_{k \in \mathbb{N}}$ converges to 0 almost surely. In addition, for $k \in \mathbb{N}_{>0}$,*

$$\min_{j \in \{0, \dots, k-1\}} \mathbb{E} [\|\nabla \mathcal{J}(\theta_j)\|^2] = O\left(\frac{1}{k^{1/2-\delta}}\right).$$

The results above state in particular that a realization of the algorithm reaches a point where the gradient is arbitrarily small with probability one. Note that the rate depends on the parameter $\delta \in (0, 1/2)$ which can be chosen by the user and corresponds to the decay schedule $1/(k+1)^{1/2+\delta}$. In most cases, one will want to slowly decay the step-size so $\delta \simeq 0$ and the rate is close to $1/\sqrt{k+1}$.

4.2 An alternative to the boundedness assumption.

In Theorem 1 the assumption that almost surely the iterates $(\theta_k)_{k \in \frac{1}{2}\mathbb{N}}$ are uniformly bounded is made. While this is usual for non-convex problems tackled with stochastic algorithms [11, 15, 17] it may be hard to check in practice. This assumption is however consistent with numerical experiments since practitioners usually choose the hyper-parameters (in particular the step-size) so that the weights of the DNN remain “not too large” for the sake of numerical stability.

One can alternatively replace the boundedness assumption by leveraging additional regularity assumptions on the loss function as Li and Orabona [29] did for example for the scalar variant of ADAGRAD. This is more restrictive than the locally-Lipschitz-continuous property of the gradient that we used but for completeness we provide below an alternative version of Theorem 1 where we assume that for each $n \in \{1, \dots, N\}$, the function \mathcal{J}_n and its gradient $\nabla \mathcal{J}_n$ are Lipschitz continuous.

Corollary 2. *Let $\theta_0 \in \mathbb{R}^P$, and let $(\theta_k)_{k \in \frac{1}{2}\mathbb{N}}$ be a sequence generated by Step-Tuned SGD initialized at θ_0 . Assume that each \mathcal{J}_n and $\nabla \mathcal{J}_n$ are Lipschitz continuous on \mathbb{R}^P and that each \mathcal{J}_n is bounded below, for all $n \in \{1, \dots, N\}$. Then the same conclusions as in Theorem 1 apply.*

Note that the assumption that each \mathcal{J}_n is Lipschitz continuous is similar to assumptions (H4) and (H4') from Li and Orabona [29, Section 3], yet a little less general. This is due to the fact that using each mini-batch twice brings additional difficulties when studying the convergence (see the proof of Theorem 1).

4.3 Proof sketch of Theorem 1.

The proof of our main theorem is fully detailed in Section B of the Supplementary. Here we present the key elements of this proof.

- The proof relies on the descent lemma: for any compact subset $C \subset \mathbb{R}^P$ there exists $L > 0$ such that for any $\theta \in C$ and $d \in \mathbb{R}^P$ such that $\theta + d \in C$,

$$\mathcal{J}(\theta + d) \leq \mathcal{J}(\theta) + \langle \nabla \mathcal{J}(\theta), d \rangle + \frac{L}{2} \|d\|^2. \quad (22)$$

- Let $(\theta_k)_{k \in \frac{1}{2}\mathbb{N}}$ be a realization of the algorithm. Using the boundedness assumption, almost surely the iterates belong to a compact subset C on which $\nabla \mathcal{J}$ and the gradients estimates $\nabla \mathcal{J}_{B_k}$ are uniformly bounded. So at any iteration $k \in \mathbb{N}$, we may use the descent lemma (22) on the update direction $d = -\gamma_k \nabla \mathcal{J}_{B_k}(\theta_k)$ to bound the difference $\mathcal{J}(\theta_{k+1}) - \mathcal{J}(\theta_k)$.
- As stated in Section 3.2, conditioning on B_0, \dots, B_{k-1} the step-size γ_k is constructed to be independent of the current mini-batch B_k . Using this and the descent lemma, we show that there exist $M_1, M_2 > 0$ such that, for all $k \in \mathbb{N}_{>0}$,

$$\mathbb{E}[\mathcal{J}(\theta_{k+1}) \mid B_0, \dots, B_{k-1}] \leq \mathcal{J}(\theta_k) - \frac{M_1}{(k+1)^{1/2+\delta}} \|\nabla \mathcal{J}(\theta_k)\|^2 + \frac{M_2}{(k+1)^{1+2\delta}}, \quad (23)$$

where $\mathbb{E}[\mathcal{J}(\theta_{k+1}) \mid B_0, \dots, B_{k-1}]$ denotes the conditional expectation of the random value $\mathcal{J}(\theta_{k+1})$ conditionally on the mini-batches B_0, \dots, B_{k-1} .

- Applying Robbins-Siegmund convergence theorem [37] for martingales to (23), using the fact that $\sum_{k=0}^{+\infty} \frac{1}{(k+1)^{1+2\delta}} < \infty$, we obtain almost surely that the sequence $(\mathcal{J}(\theta_k))_{k \in \mathbb{N}}$ converges and

$$\sum_{k=0}^{+\infty} \frac{1}{(k+1)^{1/2+\delta}} \|\nabla \mathcal{J}(\theta_k)\|^2 < +\infty, \quad (24)$$

Since $\sum_{k=0}^{+\infty} \frac{1}{(k+1)^{1/2+\delta}} = +\infty$, we deduce that $\nabla \mathcal{J}(\theta_k)$ converges to zero almost surely, using the local Lipschitz continuity of the gradient (from twice differentiability) and an argument of Alber et al. [1]. The rate follows from considering expectations on both sides of (23).

5 Application to Deep Learning

We finally evaluate the performance of Step-Tuned SGD by training DNNs. We consider six different problems presented next and fully-specified in Section A of the Supplementary. The results with Problems (a) to (d) are first presented here while the results with Problems (e) and (f) are discussed at the end of this section. We compare Step-Tuned SGD with two of the most popular non-momentum methods, SGD and RMSprop [43], and we also consider the momentum method ADAM [24] which is a very popular DL optimizer. Our method is detailed below.

5.1 Setting of the experiments

- We consider image classification problems with color images CIFAR-10 and CIFAR-100 [26] and the training of an auto-encoder on MNIST [28].
- The networks are slightly modified versions of Lenet [27], ResNet-20 [18], Network-in-Network (NiN) [31] and the auto-encoder of Hinton and Salakhutdinov [19].
- As specified in Table 1 of the Supplementary, we used either smooth (ELU, SiLU) or nonsmooth (ReLU) activations. Though our theoretical analysis only applies to smooth activations, we did not in practice observe a significant qualitative difference between ReLU or its smooth versions.
- For image classification tasks, the dissimilarity measure is the cross-entropy, and for the auto-encoder, it is the mean-squared error. In each problem we also add a ℓ^2 -regularization parameter (a.k.a. weight decay) of the form $\frac{\lambda}{2}\|\theta\|_2^2$.
- For each algorithm, we selected the learning rate parameter α from the set $\{10^{-4}, \dots, 10^0\}$. The value is selected as the one yielding minimum training loss after 10% of the total number of epochs. For example, if we intend to train the network during 100 epochs, the grid-search is carried on the first 10 epochs. For Step-Tuned SGD, the parameter ν was selected with the same criterion from the set $\{1, 2, 5\}$ and for ADAM the momentum parameter was chosen in $\{0.1, 0.5, 0.9, 0.99\}$. All other parameters of the algorithms are left to their default values.
- Decay-schedule: To meet the conditions of Theorem 1 the step-size decay schedule of SGD and Step-Tuned SGD takes the form $1/q^{1/2+\delta}$ where q is the current epoch index and $\delta = 0.001$. It slightly differs from what is given in Algorithm 2 as we apply the decay at each epoch instead of each iteration. This slower schedule still satisfies the conditions for the convergence of Theorem 1.⁵ RMSprop and ADAM rely on their own adaptive procedure and are usually used without step-size decay schedule.
- The experiments were run on a Nvidia GTX 1080 TI GPU, with an Intel Xeon 2640 V4 CPU. The code was written in *Python* 3.6.9 and *PyTorch* 1.4 [34].

Second experiment: mini-batch sub-sampling. Step-Tuned SGD departs from the standard process of drawing a new mini-batch after each gradient update. Indeed, we use each mini-batch twice in order to properly approximate the curvature of the sliding loss, but also to maintain a computing time similar to standard algorithms. We performed additional experiments to understand the consequences of using the same mini-batch twice, and in particular make sure that this is not the source of the observed advantage of Step-Tuned SGD. In these experiments all methods are used with the mini-batch drawing procedure of Step-Tuned SGD detailed in Algorithm 2 (each mini-batch being used to perform two consecutive gradient steps).

5.2 Results

We describe the results for the two types of experiments, the comparative one to assess the quality of Step-Tuned SGD against concurrent optimization algorithms, and the other one to

⁵An alternative common practice consists in manually decaying the step-size at pre-defined epochs. This technique although efficient in practice to achieve state-of-the-art results makes the comparison of algorithms harder, hence we stick to a usual Robbins-Monro type of decay.

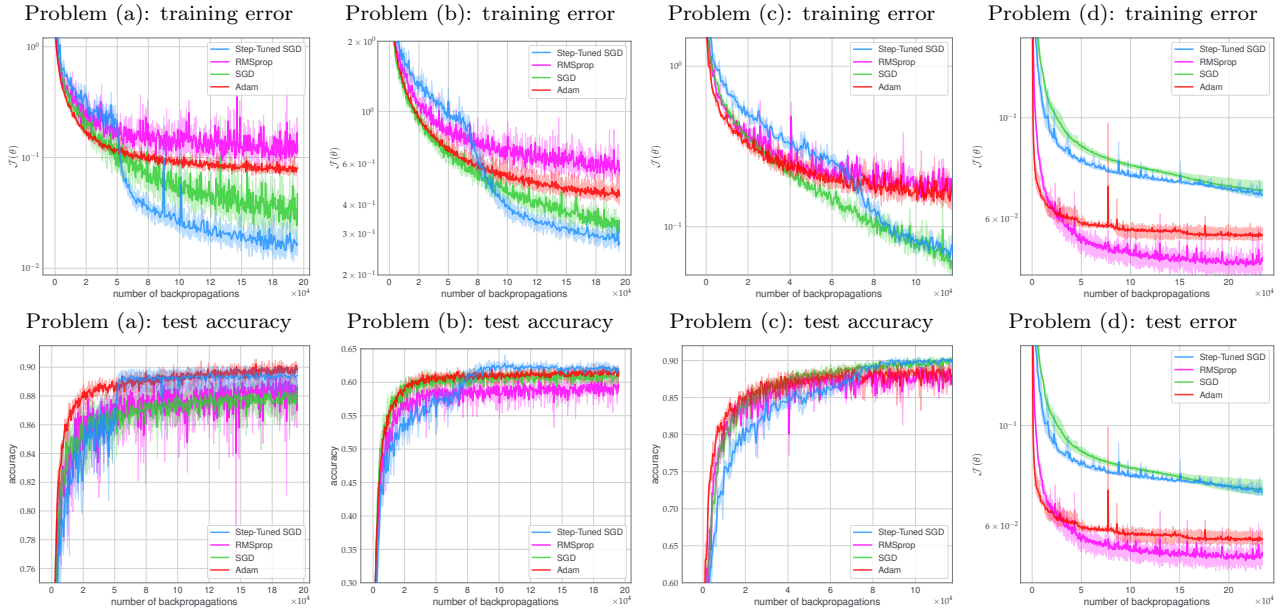


Figure 4: Classification of CIFAR-10 and CIFAR-100 with ResNet-20 (left and middle-left respectively), CIFAR-10 with NiN (middle-right) and training of an auto-encoder on MNIST (right). This corresponds to Problems (a) to (d) specified in Table 1. Continuous lines: average values from 3 random initializations. Limits of shadow area: best and worst runs (in training loss). For fair comparison values are plotted against the number of gradient estimates computed (using back-propagation).

study the effect of changing the way mini-batches are used.

5.2.1 Comparison with standard methods.

The results for problems (a) to (d) are displayed on Figure 4. For each problem we display the evolution of the values of the loss function and of the test accuracy during the training phase. We observe a recurrent behavior: during early training Step-Tuned SGD behaves similarly than other methods, then there is a sudden drop of the loss (combined with an improvement in terms of test accuracy which we discuss below). As a result, Step-Tuned SGD achieves the best training performance among all algorithms on problems (a) and (b) and at least outperforms SGD in five of the six problems considered (result for Problems (e) and (f) are on Figure 6). The sudden drop observed is in accordance with our preliminary observations in Figure 3. We note that a similar drop and improved results are reported for SGD and ADAM when used with a manually enforced reduction of the learning rate see, e.g. He et al. [18]. Our experiments show however that Step-Tuned SGD behaves similarly but in an automatic way: the *drop down* is caused by the automatic fine-tuning the algorithm is designed to achieve and not by user-defined reduction of the step-size.

We remark that in problem (d) ADAM and RMSprop are notably better than SGD and Step-Tuned SGD. This may be explained by their vector step-sizes (a scalar step-size for each coordinate of θ) as auto-encoders are often ill-conditioned, making methods with scalar step-sizes less efficient. To conclude on these comparative experiments, in most cases Step-Tuned SGD represents a significant improvement compared to SGD. It also seems to be a good alternative to adaptive methods like RMSprop or ADAM especially on residual networks. Note also that while stochastic second-order methods usually perform well mostly when combined with

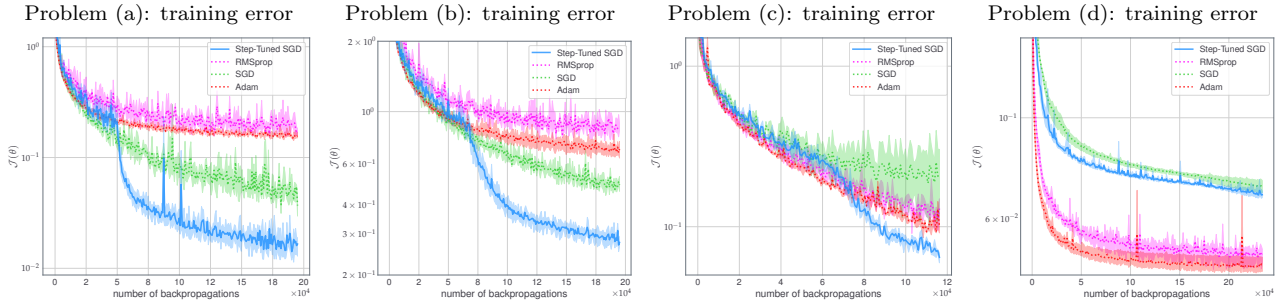


Figure 5: Experiment where each algorithms receives the same mini-batch for two consecutive iterations as in Algorithm 2. This allows to compare algorithms with respect to the number of data processed. The problems and the framework are the same as in Figure 4.

large mini-batches (hence with less-noisy gradients), we obtain satisfactory performances with mini-batches of standard sizes.

In addition to efficient training performances (in terms of loss function values), Step-Tuned SGD generalizes well (as measured by test accuracy). For example Figure 4 shows a correlation between test accuracy and training loss. Conditions or explanations for when this happens are not fully understood to this day. Yet, SGD is often said to behave well with respect to this matter [45] and hence it is satisfactory to observe that Step-Tuned SGD seems to inherit this property.

5.2.2 Effect of the mini-batch sub-sampling of Step-Tuned SGD.

The results are presented on Figure 5. We observe that using each mini-batch twice usually reduces the performance of SGD, ADAM and RMSprop, except on problem (c) where it benefits the latter two in term of training error. Thus, on these problems, changing the way of using mini-batches is not the reason for the success of our method. On the contrary, it seems that our goal which was to obtain a fine-tuned step-size specifically for each iteration is clearly achieved, but processing data more slowly, like Step-Tuned SGD does, can sometimes impact the performances of the algorithm.

Arguably these results show that the need for using each mini-batch twice appears to be the main downside of Step-Tuned SGD. Thus in problems where mini-batches may be very different we should expect other methods to be more efficient as they process data twice faster. We actually remark that our method achieves its best results on networks where batch normalization (BatchNorm) is used, a technique that aims to normalize the inputs of neural networks [22]. Figure 6 corroborates these observations: BatchNorm has a positive effect on Step-Tuned SGD.

6 Conclusion

We presented a new method to tune SGD’s step-sizes for stochastic non-convex optimization within a first-order computational framework. In addition to the new algorithm, we also presented a generic strategy (Section 3.1 and 3.3) on how to use empirical and geometric considerations to address the major difficulty of preserving favorable behaviors of deterministic algorithms while dealing with mini-batches. In particular, we tackled the problem of adapting the step-sizes to the local landscape of non-convex loss functions with noisy estimations. For a

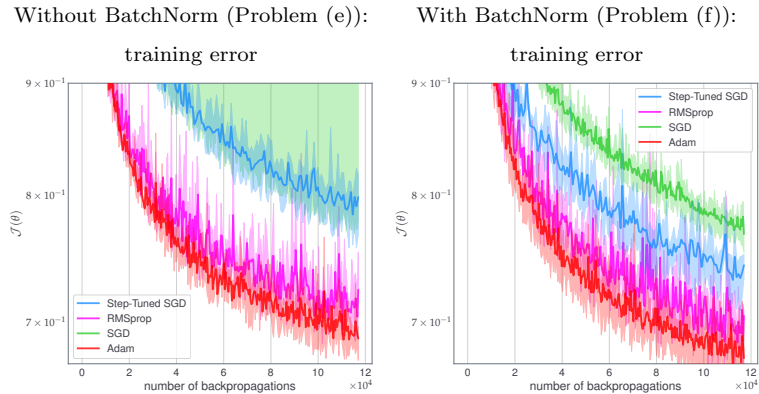


Figure 6: Classification of CIFAR-10 with LeNet with and without batch normalization, corresponding to Problems (e) and (f) specified in Section A of the Supplementary. These experiments illustrate how batch normalization has a positive effect on Step-Tuned SGD.

computational cost similar to SGD, our method uses a step-size schedule changing every two iterations unlike other stochastic methods à la Barzilai-Borwein. Our algorithm comes with asymptotic convergence results and convergence rates.

While our method does not alleviate hyper-parameter pre-tuning, it shows how an efficient *automatic* fine-tuning of a *simple* scalar step-size can improve the training of DNNs. Step-Tuned SGD processes data more slowly than other DL methods but by doing so manages to fine-tune step-sizes, leading to faster training in some DL problems with a typical sudden drop of the error rate at medium stages, especially on ResNets.

Acknowledgements

The authors acknowledge the support of the European Research Council (ERC FACTORY-CoG-6681839), the Agence Nationale de la Recherche (ANR 3IA-ANITI, ANR-17-EURE-0010 CHESS, ANR-19-CE23-0017 MASDOL) and the Air Force Office of Scientific Research (FA9550-18-1-0226).

Part of the numerical experiments were done using the OSIRIM platform of IRIT, supported by the CNRS, the FEDER, Région Occitanie and the French government (<http://osirim.irit.fr/site/en>). We thank the development teams of the following libraries that were used in the experiments: Python [39], Numpy [44], Matplotlib [20], PyTorch [34], and the PyTorch implementation of ResNets from Idelbayev [21].

We thank Emmanuel Soubies and Sixin Zhang for useful discussions and Sébastien Gadat for pointing out flaws in the original proof.

References

- [1] Alber, Y. I., Iusem, A. N., and Solodov, M. V. (1998). On the projected subgradient method for nonsmooth convex optimization in a hilbert space. *Mathematical Programming*, 81(1):23–35.
- [2] Allen-Zhu, Z. (2018). Natasha 2: Faster non-convex optimization than SGD. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2675–2686.

- [3] Alvarez, F. and Cabot, A. (2004). Steepest descent with curvature dynamical system. *Journal of Optimization Theory and Applications*, 120(2):247–273.
- [4] Babaie-Kafaki, S. and Fatemi, M. (2013). A modified two-point stepsize gradient algorithm for unconstrained minimization. *Optimization Methods and Software*, 28(5):1040–1050.
- [5] Barakat, A. and Bianchi, P. (2018). Convergence of the ADAM algorithm from a dynamical system viewpoint. *arXiv preprint:1810.02263*.
- [6] Barzilai, J. and Borwein, J. M. (1988). Two-point step size gradient methods. *IMA journal of Numerical Analysis*, 8(1):141–148.
- [7] Bertsekas, D. P., Hager, W., and Mangasarian, O. (1998). *Nonlinear programming*. Athena Scientific Belmont, MA.
- [8] Biglari, F. and Solimanpur, M. (2013). Scaling on the spectral gradient method. *Journal of Optimization Theory and Applications*, 158:626–635.
- [9] Bolte, J. and Pauwels, E. (2020). A mathematical model for automatic differentiation in machine learning. In *Advances in Neural Information Processing Systems (NIPS)*.
- [10] Carmon, Y., Duchi, J. C., Hinder, O., and Sidford, A. (2017). Convex until proven guilty: Dimension-free acceleration of gradient descent on non-convex functions. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 654–663.
- [11] Castera, C., Bolte, J., Févotte, C., and Pauwels, E. (2021). An inertial Newton algorithm for deep learning. *Journal of Machine Learning Research*, 22(134):1–31.
- [12] Curtis, F. E. and Guo, W. (2016). Handling nonpositive curvature in a limited memory steepest descent method. *IMA Journal of Numerical Analysis*, 36(2):717–742.
- [13] Curtis, F. E. and Robinson, D. P. (2019). Exploiting negative curvature in deterministic and stochastic optimization. *Mathematical Programming*, 176(1-2):69–94.
- [14] Dai, Y., Yuan, J., and Yuan, Y.-X. (2002). Modified two-point stepsize gradient methods for unconstrained optimization. *Computational Optimization and Applications*, 22(1).
- [15] Davis, D., Drusvyatskiy, D., Kakade, S., and Lee, J. D. (2020). Stochastic subgradient method converges on tame functions. *Foundations of Computational mathematics*, 20(1):119–154.
- [16] Duchi, J., Hazan, E., and Singer, Y. (2011). Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7).
- [17] Duchi, J. C. and Ruan, F. (2018). Stochastic methods for composite and weakly convex optimization problems. *SIAM Journal on Optimization*, 28(4):3229–3259.
- [18] He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778.
- [19] Hinton, G. E. and Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507.

- [20] Hunter, J. D. (2007). Matplotlib: A 2d graphics environment. *Computing in science & engineering*, 9(3):90–95.
- [21] Idelbayev, Y. (2018). Proper ResNet implementation for CIFAR10/CIFAR100 in PyTorch. https://github.com/akamaster/pytorch_resnet_cifar10.
- [22] Ioffe, S. and Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 448–456.
- [23] Johnson, R. and Zhang, T. (2013). Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in Neural Information Processing Systems (NIPS)*, pages 315–323.
- [24] Kingma, D. P. and Ba, J. (2015). Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- [25] Krishnan, S., Xiao, Y., and Saurous, R. A. (2018). Neumann optimizer: A practical optimization algorithm for deep neural networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- [26] Krizhevsky, A. (2009). Learning multiple layers of features from tiny images. Technical report, Canadian Institute for Advanced Research.
- [27] LeCun, Y., Bottou, L., Bengio, Y., Haffner, P., et al. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324.
- [28] LeCun, Y., Cortes, C., and Burges, C. (2010). MNIST handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>.
- [29] Li, X. and Orabona, F. (2019). On the convergence of stochastic gradient descent with adaptive stepsizes. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 983–992.
- [30] Liang, J., Xu, Y., Bao, C., Quan, Y., and Ji, H. (2019). Barzilai–Borwein-based adaptive learning rate for deep learning. *Pattern Recognition Letters*, 128:197 – 203.
- [31] Lin, M., Chen, Q., and Yan, S. (2013). Network in network. *arXiv preprint:1312.4400*.
- [32] Liu, M. and Yang, T. (2017). On noisy negative curvature descent: Competing with gradient descent for faster non-convex optimization. *arXiv preprint:1709.08571*.
- [33] Martens, J. and Grosse, R. (2015). Optimizing neural networks with kronecker-factored approximate curvature. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 2408–2417.
- [34] Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. (2019). Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems (NIPS)*, pages 8026–8037.
- [35] Raydan, M. (1997). The Barzilai and Borwein gradient method for the large scale unconstrained minimization problem. *SIAM Journal on Optimization*, 7(1):26–33.

- [36] Robbins, H. and Monro, S. (1951). A stochastic approximation method. *The Annals of Mathematical Statistics*, 22(1):400–407.
- [37] Robbins, H. and Siegmund, D. (1971). A convergence theorem for non negative almost supermartingales and some applications. In *Optimizing methods in statistics*, pages 233–257. Elsevier.
- [38] Robles-Kelly, A. and Nazari, A. (2019). Incorporating the Barzilai-Borwein adaptive step size into subgradient methods for deep network training. In *2019 Digital Image Computing: Techniques and Applications (DICTA)*, pages 1–6.
- [39] Rossum, G. (1995). *Python reference manual*. CWI (Centre for Mathematics and Computer Science).
- [40] Royer, C. W. and Wright, S. J. (2018). Complexity analysis of second-order line-search algorithms for smooth nonconvex optimization. *SIAM Journal on Optimization*, 28(2):1448–1477.
- [41] Schraudolph, N. N., Yu, J., and Günter, S. (2007). A stochastic Quasi-Newton method for online convex optimization. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*.
- [42] Tan, C., Ma, S., Dai, Y.-H., and Qian, Y. (2016). Barzilai-Borwein step size for stochastic gradient descent. In *Advances in Neural Information Processing Systems (NIPS)*, pages 685–693.
- [43] Tieleman, T. and Hinton, G. (2012). Lecture 6.5-RMSprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4(2):26–31.
- [44] Walt, S. v. d., Colbert, S. C., and Varoquaux, G. (2011). The numpy array: a structure for efficient numerical computation. *Computing in science & engineering*, 13(2):22–30.
- [45] Wilson, A. C., Roelofs, R., Stern, M., Srebro, N., and Recht, B. (2017). The marginal value of adaptive gradient methods in machine learning. In *Advances in Neural Information Processing Systems (NIPS)*, pages 4148–4158.
- [46] Xiao, Y., Wang, Q., and Wang, D. (2010). Notes on the Dai–Yuan–Yuan modified spectral gradient method. *Journal of Computational and Applied Mathematics*, 234(10):2986 – 2992.
- [47] Zhuang, J., Tang, T., Ding, Y., Tatikonda, S. C., Dvornik, N., Papademetris, X., and Duncan, J. (2020). Adabelief optimizer: Adapting stepsizes by the belief in observed gradients. *Advances in Neural Information Processing Systems (NIPS)*, 33.

A Details about deep learning experiments

In addition to the method described in Section 5.1, we provide in Table 1 a summary of each problem considered.

In the DL experiments of Section 5, we display the training error and the test accuracy of each algorithm as a function of the number of stochastic gradient estimates computed. Due to

their adaptive procedures, ADAM, RMSprop and Step-Tuned SGD have additional sub-routines in comparison to SGD. Thus, in Table 2 we additionally provide the wall-clock time per epoch of these methods relatively to SGD. Unlike the number of back-propagations performed, wall-clock time depends on many factors: the network and datasets considered, the computer used, and most importantly, the implementation. Regarding implementation, we would like to emphasize the fact that we used the versions of SGD, ADAM and RMSprop provided in *PyTorch*, which are fully optimized (and in particular parallelized). Table 2 indicates that Step-Tuned SGD is slower than other adaptive methods for large networks but this is due to our non-parallel implementation. Actually on small networks (where the benefits of parallel computing is small), we observe that running Step-Tuned SGD for one epoch is actually faster than for SGD. As a conclusion, the number of back-propagations is a more suitable metric for comparing the algorithms, and all methods considered require a single back-propagation per iteration.

Table 1: Setting of the four different deep learning experiments.

	Problem (a)	Problem (b)	Problem (c)
Type	Classification	Classification	Classification
Dataset	CIFAR-10	CIFAR-100	CIFAR-10
Network	ResNet-20 (Residual)	ResNet-20 (Residual)	Network-in-Network (Nested)
BatchNorm	Yes	Yes	Yes
Batch-size	128	128	128
Activation functions	ReLU	ReLU	ELU
Dissimilarity measure	Cross-entropy	Cross-entropy	Cross-entropy
Regularization	$\lambda = 10^{-4}$	$\lambda = 10^{-4}$	$\lambda = 10^{-4}$
Grid-search	50 epochs	50 epochs	30 epochs
Stop-criterion	500 epochs	500 epochs	300 epochs
	Problem (d)	Problem (e)	Problem (f)
Type	Auto-encoder	Classification	Classification
Dataset	MNIST	CIFAR-10	CIFAR-10
Network	Auto-Encoder (Dense)	LeNet (Convolutional)	LeNet (Convolutional)
BatchNorm	No	No	Yes
Batch-size	128	128	128
Activation functions	SiLU	ELU	ELU
Dissimilarity measure	Mean square	Cross-entropy	Cross-entropy
Regularization	$\lambda = 10^{-4}$	$\lambda = 10^{-4}$	$\lambda = 10^{-4}$
Grid-search	50 epochs	30 epochs	30 epochs
Stop-criterion	500 epochs	300 epochs	300 epochs

Table 2: Relative wall-clock time per epoch compared to SGD.

	Prob.(a)	Prob.(b)	Prob.(c)	Prob.(d)	Prob.(e)	Prob.(f)
ADAM	1.13	1.13	1.03	1.18	1.04	1.00
RMSprop	1.06	1.08	1.02	1.13	1.00	1.01
Step-Tuned SGD	1.67	1.71	1.20	1.47	0.71	0.88

B Proof of the theoretical results

We state a lemma that we will use to prove Theorem 1.

B.1 Preliminary lemma

The result is the following.

Lemma 3 (Alber et al. [1, Proposition 2]). *Let $(u_k)_{k \in \mathbb{N}}$ and $(v_k)_{k \in \mathbb{N}}$ two non-negative real sequences. Assume that $\sum_{k=0}^{+\infty} u_k v_k < +\infty$, and $\sum_{k=0}^{+\infty} v_k = +\infty$. If there exists a constant $C > 0$ such that $\forall k \in \mathbb{N}, |u_{k+1} - u_k| \leq C v_k$, then $u_k \xrightarrow[k \rightarrow +\infty]{} 0$.*

B.2 Proof of the main theorem

We can now prove Theorem 1.

Proof of Theorem 1. We first clarify the random process induced by the draw of the mini-batches. Algorithm 2 takes a sequence of mini-batches as input. This sequence is represented by the random variables $(\mathbf{B}_k)_{k \in \mathbb{N}}$ as described in Section 3.2. Each of these random variables is independent of the others. In particular, for $k \in \mathbb{N}_{>0}$, \mathbf{B}_k is independent of the previous mini-batches $\mathbf{B}_0, \dots, \mathbf{B}_{k-1}$. For convenience, we will denote $\underline{\mathbf{B}}_k = \{\mathbf{B}_0, \dots, \mathbf{B}_k\}$, the mini-batches up to iteration k . Due to the randomness of the mini-batches, the algorithm is a random process as well. As such, θ_k is a random variable with a deterministic dependence on $\underline{\mathbf{B}}_{k-1}$ and is independent of \mathbf{B}_k . However, $\theta_{k+\frac{1}{2}}$ and \mathbf{B}_k are not independent. Similarly, we constructed γ_k such that it is a random variable with a deterministic dependence on $\underline{\mathbf{B}}_{k-1}$, which is independent of \mathbf{B}_k . This dependency structure will be crucial to derive and bound conditional expectations. Finally, we highlight the following important identity, for any $k \in \mathbb{N}_{>0}$,

$$\mathbb{E} [\nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k) | \underline{\mathbf{B}}_{k-1}] = \nabla \mathcal{J}(\theta_k). \quad (25)$$

Indeed, the iterate θ_k is a deterministic function of $\underline{\mathbf{B}}_{k-1}$, so taking the expectation over \mathbf{B}_k , which is independent of $\underline{\mathbf{B}}_{k-1}$, we recover the full gradient of \mathcal{J} as the distribution of \mathbf{B}_k is the same as that of \mathbf{S} in Section 3.2. Notice in addition that a similar identity does not hold for $\theta_{k+\frac{1}{2}}$ (as it depends on \mathbf{B}_k).

We now provide estimates that will be used extensively in the rest of the proof. The gradient of the loss function $\nabla \mathcal{J}$ is locally Lipschitz continuous as \mathcal{J} is twice continuously differentiable. By assumption, there exists a compact convex set $\mathbf{C} \subset \mathbb{R}^P$, such that with probability 1, the sequence of iterates $(\theta_k)_{k \in \frac{1}{2}\mathbb{N}}$ belongs to \mathbf{C} . Therefore, by local Lipschitz continuity, the restriction of $\nabla \mathcal{J}$ to \mathbf{C} is Lipschitz continuous on \mathbf{C} . Similarly, each $\nabla \mathcal{J}_n$ is also Lipschitz continuous on \mathbf{C} . We denote by $L > 0$ a Lipschitz constant common to each $\nabla \mathcal{J}_n$, $n = 1, \dots, N$. Notice that the Lipschitz continuity is preserved by averaging, in other words,

$$\forall \mathbf{B} \subseteq \{1, \dots, N\}, \forall \psi_1, \psi_2 \in \mathbf{C}, \quad \|\nabla \mathcal{J}_{\mathbf{B}}(\psi_1) - \nabla \mathcal{J}_{\mathbf{B}}(\psi_2)\| \leq L \|\psi_1 - \psi_2\|. \quad (26)$$

In addition, using the continuity of the $\nabla \mathcal{J}_n$'s, there exists a constant $C_2 > 0$, such that,

$$\forall \mathbf{B} \subseteq \{1, \dots, N\}, \forall \psi \in \mathbf{C}, \quad \|\nabla \mathcal{J}_{\mathbf{B}}(\psi)\| \leq C_2. \quad (27)$$

Finally, for a function $g : \mathbb{R}^P \rightarrow \mathbb{R}$ with L -Lipschitz continuous gradient, we recall the following inequality called descent lemma (see for example Bertsekas et al. [7, Proposition A.24]). For any $\theta \in \mathbb{R}^P$ and any $d \in \mathbb{R}^P$,

$$g(\theta + d) \leq g(\theta) + \langle \nabla g(\theta), d \rangle + \frac{L}{2} \|d\|^2. \quad (28)$$

In our case since we only have the L -Lipschitz continuity of $\nabla \mathcal{J}$ on \mathbf{C} which is convex, we have a similar bound for $\nabla \mathcal{J}$ on \mathbf{C} : for any $\theta \in \mathbf{C}$ and any $d \in \mathbb{R}^P$ such that $\theta + d \in \mathbf{C}$,

$$\mathcal{J}(\theta + d) \leq \mathcal{J}(\theta) + \langle \nabla \mathcal{J}(\theta), d \rangle + \frac{L}{2} \|d\|^2. \quad (29)$$

Let $\theta_0 \in \mathbb{R}^P$ and let $(\theta_k)_{k \in \frac{1}{2}\mathbb{N}}$ a sequence generated by Algorithm 2 initialized at θ_0 . By assumption this sequence belongs to \mathbf{C} almost surely. To simplify, for $k \in \mathbb{N}$, we denote $\eta_k = \alpha \gamma_k (k+1)^{-(1/2+\delta)}$. Fix an iteration $k \in \mathbb{N}$, we can use (29) with $\theta = \theta_k$ and $d = -\eta_k \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k)$, almost surely (with respect to the boundedness assumption),

$$\mathcal{J}(\theta_{k+\frac{1}{2}}) \leq \mathcal{J}(\theta_k) - \eta_k \langle \nabla \mathcal{J}(\theta_k), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k) \rangle + \frac{\eta_k^2}{2} L \|\nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k)\|^2. \quad (30)$$

Similarly with $\theta = \theta_{k+\frac{1}{2}}$ and $d = -\eta_k \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_{k+\frac{1}{2}})$, almost surely,

$$\mathcal{J}(\theta_{k+1}) \leq \mathcal{J}(\theta_{k+\frac{1}{2}}) - \eta_k \langle \nabla \mathcal{J}(\theta_{k+\frac{1}{2}}), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_{k+\frac{1}{2}}) \rangle + \frac{\eta_k^2}{2} L \|\nabla \mathcal{J}_{\mathbf{B}_k}(\theta_{k+\frac{1}{2}})\|^2. \quad (31)$$

We combine (30) and (31), almost surely,

$$\begin{aligned} \mathcal{J}(\theta_{k+1}) \leq & \mathcal{J}(\theta_k) - \eta_k \left(\langle \nabla \mathcal{J}(\theta_k), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k) \rangle + \langle \nabla \mathcal{J}(\theta_{k+\frac{1}{2}}), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_{k+\frac{1}{2}}) \rangle \right) \\ & + \frac{\eta_k^2}{2} L \left(\|\nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k)\|^2 + \|\nabla \mathcal{J}_{\mathbf{B}_k}(\theta_{k+\frac{1}{2}})\|^2 \right). \end{aligned} \quad (32)$$

Using the boundedness assumption and (27), almost surely,

$$\|\nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k)\|^2 \leq C_2 \quad \text{and} \quad \|\nabla \mathcal{J}_{\mathbf{B}_k}(\theta_{k+\frac{1}{2}})\|^2 \leq C_2. \quad (33)$$

So almost surely,

$$\mathcal{J}(\theta_{k+1}) \leq \mathcal{J}(\theta_k) - \eta_k \left(\langle \nabla \mathcal{J}(\theta_k), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k) \rangle + \langle \nabla \mathcal{J}(\theta_{k+\frac{1}{2}}), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_{k+\frac{1}{2}}) \rangle \right) + \eta_k^2 L C_2. \quad (34)$$

Then, we take the conditional expectation of (34) over \mathbf{B}_k conditionally on $\underline{\mathbf{B}}_{k-1}$ (the mini-batches used up to iteration $k-1$), we have,

$$\begin{aligned} \mathbb{E} [\mathcal{J}(\theta_{k+1}) | \underline{\mathbf{B}}_{k-1}] \leq & \mathbb{E} [\mathcal{J}(\theta_k) | \underline{\mathbf{B}}_{k-1}] + \mathbb{E} [\eta_k^2 L C_2 | \underline{\mathbf{B}}_{k-1}] \\ & - \mathbb{E} \left[\eta_k \left(\langle \nabla \mathcal{J}(\theta_k), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k) \rangle + \langle \nabla \mathcal{J}(\theta_{k+\frac{1}{2}}), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_{k+\frac{1}{2}}) \rangle \right) \middle| \underline{\mathbf{B}}_{k-1} \right]. \end{aligned} \quad (35)$$

As explained at the beginning of the proof, θ_k is a deterministic function of $\underline{\mathbf{B}}_{k-1}$, thus, $\mathbb{E} [\mathcal{J}(\theta_k) | \underline{\mathbf{B}}_{k-1}] = \mathcal{J}(\theta_k)$. Similarly, by construction η_k is independent of the current mini-batch \mathbf{B}_k , it is a deterministic function of $\underline{\mathbf{B}}_{k-1}$. Hence, (35) reads,

$$\begin{aligned} \mathbb{E} [\mathcal{J}(\theta_{k+1}) | \underline{\mathbf{B}}_{k-1}] \leq & \mathcal{J}(\theta_k) + \eta_k^2 L C_2 - \eta_k \langle \nabla \mathcal{J}(\theta_k), \mathbb{E} [\nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k) | \underline{\mathbf{B}}_{k-1}] \rangle \\ & - \eta_k \mathbb{E} \left[\langle \nabla \mathcal{J}(\theta_{k+\frac{1}{2}}), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_{k+\frac{1}{2}}) \rangle \middle| \underline{\mathbf{B}}_{k-1} \right]. \end{aligned} \quad (36)$$

Then, we use the fact that $\mathbb{E} [\nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k) | \underline{\mathbf{B}}_{k-1}] = \nabla \mathcal{J}(\theta_k)$. Overall, we obtain,

$$\begin{aligned} \mathbb{E} [\mathcal{J}(\theta_{k+1}) | \underline{\mathbf{B}}_{k-1}] \leq & \mathcal{J}(\theta_k) + \eta_k^2 L C_2 - \eta_k \|\nabla \mathcal{J}(\theta_k)\|^2 \\ & - \eta_k \mathbb{E} \left[\langle \nabla \mathcal{J}(\theta_{k+\frac{1}{2}}), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_{k+\frac{1}{2}}) \rangle \middle| \underline{\mathbf{B}}_{k-1} \right]. \end{aligned} \quad (37)$$

We will now bound the last term of (37). First we write,

$$\begin{aligned} & - \langle \nabla \mathcal{J}(\theta_{k+\frac{1}{2}}), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_{k+\frac{1}{2}}) \rangle \\ & = - \langle \nabla \mathcal{J}(\theta_{k+\frac{1}{2}}), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_{k+\frac{1}{2}}) - \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k) \rangle - \langle \nabla \mathcal{J}(\theta_{k+\frac{1}{2}}), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k) \rangle. \end{aligned} \quad (38)$$

Using the Cauchy-Schwarz inequality, as well as (26) and (27), almost surely,

$$\begin{aligned} |\langle \nabla \mathcal{J}(\theta_{k+\frac{1}{2}}), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_{k+\frac{1}{2}}) - \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k) \rangle| & \leq \|\nabla \mathcal{J}(\theta_{k+\frac{1}{2}})\| \|\nabla \mathcal{J}_{\mathbf{B}_k}(\theta_{k+\frac{1}{2}}) - \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k)\| \\ & \leq \|\nabla \mathcal{J}(\theta_{k+\frac{1}{2}})\| L \|\theta_{k+\frac{1}{2}} - \theta_k\| \\ & \leq \|\nabla \mathcal{J}(\theta_{k+\frac{1}{2}})\| L \|\eta_k \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k)\| \\ & \leq LC_2^2 \eta_k. \end{aligned} \quad (39)$$

Hence,

$$- \langle \nabla \mathcal{J}(\theta_{k+\frac{1}{2}}), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_{k+\frac{1}{2}}) \rangle \leq LC_2^2 \eta_k - \langle \nabla \mathcal{J}(\theta_{k+\frac{1}{2}}), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k) \rangle. \quad (40)$$

We perform similar computations on the last term of (40), almost surely

$$\begin{aligned} & - \langle \nabla \mathcal{J}(\theta_{k+\frac{1}{2}}), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k) \rangle \\ & = - \langle \nabla \mathcal{J}(\theta_{k+\frac{1}{2}}) - \nabla \mathcal{J}(\theta_k), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k) \rangle - \langle \nabla \mathcal{J}(\theta_k), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k) \rangle \\ & \leq \|\nabla \mathcal{J}(\theta_{k+\frac{1}{2}}) - \nabla \mathcal{J}(\theta_k)\| \|\nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k)\| - \langle \nabla \mathcal{J}(\theta_k), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k) \rangle \\ & \leq LC_2 \|\theta_{k+\frac{1}{2}} - \theta_k\| - \langle \nabla \mathcal{J}(\theta_k), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k) \rangle \\ & \leq LC_2^2 \eta_k - \langle \nabla \mathcal{J}(\theta_k), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k) \rangle. \end{aligned} \quad (41)$$

Finally we obtain by combining (38), (40) and (41), almost surely,

$$- \langle \nabla \mathcal{J}(\theta_{k+\frac{1}{2}}), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_{k+\frac{1}{2}}) \rangle \leq 2LC_2^2 \eta_k - \langle \nabla \mathcal{J}(\theta_k), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k) \rangle. \quad (42)$$

Going back to the last term of (37), we have, taking the conditional expectation of (42), almost surely

$$\begin{aligned} & - \eta_k \mathbb{E} \left[\langle \nabla \mathcal{J}(\theta_{k+\frac{1}{2}}), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_{k+\frac{1}{2}}) \rangle \middle| \underline{\mathbf{B}}_{k-1} \right] \\ & \leq 2LC_2^2 \eta_k^2 - \eta_k \mathbb{E} \left[\langle \nabla \mathcal{J}(\theta_k), \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k) \rangle \middle| \underline{\mathbf{B}}_{k-1} \right] \\ & \leq 2LC_2^2 \eta_k^2 - \eta_k \langle \nabla \mathcal{J}(\theta_k), \mathbb{E} [\nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k) \middle| \underline{\mathbf{B}}_{k-1}] \rangle \\ & = 2LC_2^2 \eta_k^2 - \eta_k \|\nabla \mathcal{J}(\theta_k)\|^2. \end{aligned} \quad (43)$$

In the end we obtain, for an arbitrary iteration $k \in \mathbb{N}$, almost surely

$$\mathbb{E} [\mathcal{J}(\theta_{k+1}) \middle| \underline{\mathbf{B}}_{k-1}] \leq \mathcal{J}(\theta_k) - 2\eta_k \|\nabla \mathcal{J}(\theta_k)\|^2 + \eta_k^2 L(C_2 + 2C_2^2). \quad (44)$$

To simplify we assume that $\tilde{M} \geq \nu$ (otherwise set $\tilde{M} = \max(\tilde{M}, \nu)$). We use the fact that, $\eta_k \in [\frac{\alpha \tilde{m}}{(k+1)^{1/2+\delta}}, \frac{\alpha \tilde{M}}{(k+1)^{1/2+\delta}}]$, to obtain almost surely,

$$\mathbb{E} [\mathcal{J}(\theta_{k+1}) \middle| \underline{\mathbf{B}}_{k-1}] \leq \mathcal{J}(\theta_k) - 2 \frac{\alpha \tilde{m}}{(k+1)^{1/2+\delta}} \|\nabla \mathcal{J}(\theta_k)\|^2 + \frac{\alpha^2 \tilde{M}^2}{(k+1)^{1+2\delta}} L(C_2 + 2C_2^2). \quad (45)$$

Since by assumption, the last term is summable, we can now invoke Robbins-Siegmund convergence theorem [37] to obtain that, almost surely, $(\mathcal{J}(\theta_k))_{k \in \mathbb{N}}$ converges and,

$$\sum_{k=0}^{+\infty} \frac{1}{(k+1)^{1/2+\delta}} \|\nabla \mathcal{J}(\theta_k)\|^2 < +\infty. \quad (46)$$

Since $\sum_{k=0}^{+\infty} \frac{1}{(k+1)^{1/2+\delta}} = +\infty$, this implies at least that almost surely,

$$\liminf_{k \rightarrow \infty} \|\nabla \mathcal{J}(\theta_k)\|^2 = 0. \quad (47)$$

To prove that in addition $\lim_{k \rightarrow \infty} \|\nabla \mathcal{J}(\theta_k)\|^2 = 0$, we will use Lemma 3 with $u_k = \|\nabla \mathcal{J}(\theta_k)\|^2$ and $v_k = \frac{1}{(k+1)^{1/2+\delta}}$, for all $k \in \mathbb{N}$. So we need to prove that there exists $C_3 > 0$ such that $|u_{k+1} - u_k| \leq C_3 v_k$. To do so, we use the L -Lipschitz continuity of the gradients on \mathbb{C} , triangle inequalities and (27). It holds, almost surely, for all $k \in \mathbb{N}$

$$\begin{aligned} & \left| \|\nabla \mathcal{J}(\theta_{k+1})\|^2 - \|\nabla \mathcal{J}(\theta_k)\|^2 \right| \\ &= \left(\|\nabla \mathcal{J}(\theta_{k+1})\| + \|\nabla \mathcal{J}(\theta_k)\| \right) \times \left| \|\nabla \mathcal{J}(\theta_{k+1})\| - \|\nabla \mathcal{J}(\theta_k)\| \right| \\ &\leq 2C_2 \left| \|\nabla \mathcal{J}(\theta_{k+1})\| - \|\nabla \mathcal{J}(\theta_k)\| \right| \\ &\leq 2C_2 \|\nabla \mathcal{J}(\theta_{k+1}) - \nabla \mathcal{J}(\theta_k)\| \\ &\leq 2C_2 L \|\theta_{k+1} - \theta_k\| \\ &\leq 2C_2 L \left\| -\eta_k \nabla \mathcal{J}_{\mathbb{B}_k}(\theta_k) - \eta_k \nabla \mathcal{J}_{\mathbb{B}_k}(\theta_{k+\frac{1}{2}}) \right\| \\ &\leq 2C_2 L \frac{\alpha \tilde{M}}{(k+1)^{1/2+\delta}} \|\nabla \mathcal{J}_{\mathbb{B}_k}(\theta_k) + \nabla \mathcal{J}_{\mathbb{B}_k}(\theta_{k+\frac{1}{2}})\| \\ &\leq 4C_2^2 L \frac{\alpha \tilde{M}}{(k+1)^{1/2+\delta}}. \end{aligned} \quad (48)$$

So taking $C_3 = 4C_2^2 L \alpha \tilde{M}$, by Lemma 3, almost surely, $\lim_{k \rightarrow +\infty} \|\nabla \mathcal{J}(\theta_k)\|^2 = 0$. This concludes the almost sure convergence proof.

As for the rate, consider the expectation of (45) (with respect to the random variables $(\mathbb{B}_k)_{k \in \mathbb{N}}$). The tower property of the conditional expectation gives,

$$\mathbb{E}[\mathbb{E}[\mathcal{J}(\theta_{k+1}) | \mathbb{B}_{k-1}]] = \mathbb{E}[\mathcal{J}(\theta_{k+1})],$$

so we obtain, for all $k \in \mathbb{N}$,

$$2 \frac{\alpha \tilde{m}}{(k+1)^{1/2+\delta}} \mathbb{E}[\|\nabla \mathcal{J}(\theta_k)\|^2] \leq \mathbb{E}[\mathcal{J}(\theta_k)] - \mathbb{E}[\mathcal{J}(\theta_{k+1})] + \frac{\alpha^2 \tilde{M}^2}{(k+1)^{1+2\delta}} L(C_2 + 2C_2^2). \quad (49)$$

Then for $K \geq 1$, we sum from 0 to $K-1$,

$$\begin{aligned} & \sum_{k=0}^{K-1} 2 \frac{\alpha \tilde{m}}{(k+1)^{1/2+\delta}} \mathbb{E}[\|\nabla \mathcal{J}(\theta_k)\|^2] \\ &\leq \sum_{k=0}^{K-1} \mathbb{E}[\mathcal{J}(\theta_k)] - \sum_{k=0}^{K-1} \mathbb{E}[\mathcal{J}(\theta_{k+1})] + \sum_{k=0}^{K-1} \frac{\alpha^2 \tilde{M}^2}{(k+1)^{1+2\delta}} L(C_2 + 2C_2^2) \\ &= \mathcal{J}(\theta_0) - \mathbb{E}[\mathcal{J}(\theta_K)] + \sum_{k=0}^{K-1} \frac{\alpha^2 \tilde{M}^2}{(k+1)^{1+2\delta}} L(C_2 + 2C_2^2) \\ &\leq \mathcal{J}(\theta_0) - \inf_{\psi \in \mathbb{R}^P} \mathcal{J}(\psi) + \sum_{k=0}^{K-1} \frac{\alpha^2 \tilde{M}^2}{(k+1)^{1+2\delta}} L(C_2 + 2C_2^2), \end{aligned} \quad (50)$$

The right-hand side is finite, so there is a constant $C_4 > 0$ such that for any $K \in \mathbb{N}$, it holds,

$$C_4 \geq \sum_{k=0}^K \frac{1}{(k+1)^{1/2+\delta}} \mathbb{E} [\|\nabla \mathcal{J}(\theta_k)\|^2] \geq \min_{k \in \{1, \dots, K\}} \mathbb{E} [\|\nabla \mathcal{J}(\theta_k)\|^2] \sum_{k=0}^K \frac{1}{(k+1)^{1/2+\delta}} \\ \geq (K+1)^{1/2-\delta} \min_{k \in \{1, \dots, K\}} \mathbb{E} [\|\nabla \mathcal{J}(\theta_k)\|^2], \quad (51)$$

and we obtain the rate. \square

B.3 Proof of the corollary

Before proving the corollary we recall the following result.

Lemma 4. *Let $g : \mathbb{R}^P \rightarrow \mathbb{R}$ a L -Lipschitz continuous and differentiable function. Then ∇g is uniformly bounded on \mathbb{R}^P .*

We can now prove the corollary.

Proof of Corollary 2. The proof is very similar to the one of Theorem 1. Denote L the Lipschitz constant of $\nabla \mathcal{J}$. Then, the descent lemma (30) holds surely. Furthermore, since for all $n \in \{1, \dots, N\}$, each \mathcal{J}_n is Lipschitz, so is \mathcal{J} , and globally Lipschitz functions have uniformly bounded gradients so $\nabla \mathcal{J}$ has bounded gradient. This is enough to obtain (45). Similarly, at iteration $k \in \mathbb{N}$, $\mathbb{E} [\|\nabla \mathcal{J}_{\mathcal{B}_k}(\theta_k)\|]$ is also uniformly bounded. Overall these arguments allows to follow the lines of the proof of Theorem 1 and the same conclusions follow by repeating the same arguments. \square

C Details on the synthetic experiments

We detail the non-convex regression problem that we presented in Figure 2 and 3. Given a matrix $A \in \mathbb{R}^{N \times P}$ and a vector $b \in \mathbb{R}^N$, denote A_n the n -th line of A . The problem consists in minimizing a loss function of the form,

$$\theta \in \mathbb{R}^P \mapsto \mathcal{J}(\theta) = \frac{1}{N} \sum_n \phi(A_n^T \theta - b_n), \quad (52)$$

where the non-convexity comes from the function $t \in \mathbb{R} \mapsto \phi(t) = t^2/(1+t^2)$. For more details on the initialization of A and b we refer to Carmon et al. [10] where this problem is initially proposed. In the experiments of Figure 3, the mini-batch approximation was made by selecting a subset of the lines of A , which amounts to computing only a few terms of the full sum in (52). We used $N = 500$, $P = 30$ and mini-batches of size 50.

In the deterministic setting we ran each algorithm during 250 iterations and selected the hyper-parameters of each algorithm such that they achieved $|\mathcal{J}(\theta) - \mathcal{J}^*| < 10^{-1}$ as fast as possible. In the mini-batch experiments we ran each algorithm during 250 epochs and selected the hyper-parameters that yielded the smallest value of $\mathcal{J}(\theta)$ after 50 epochs.

D Description of auxiliary algorithms

We precise the heuristic algorithms used in Figure 3 and discussed in Section 3.3. Note that the step-size in Algorithm 5 is equivalent to Expected-GV but is written differently to avoid storing an additional gradient estimate.

Algorithm 3 Stochastic-GV SGD

```

1: Input:  $\alpha > 0, \nu > 0$ 
2: Input:  $\tilde{m} > 0, \tilde{M} > 0, \delta \in (0, 1/2)$ 
3: Initialize  $\theta_0 \in \mathbb{R}^P, \gamma_0 = 1$ 
4: Draw mini-batches  $(\mathbf{B}_k)_{k \in \mathbb{N}}$  independently and uniformly at random with replacement.
5:  $\theta_1 = \theta_0 - \alpha \gamma_0 \nabla \mathcal{J}_{\mathbf{B}_0}(\theta_0)$ 
6: for  $k = 1, \dots$  do
7:    $\Delta \theta_k = \theta_k - \theta_{k-1}$ 
8:    $\Delta g_k^{\text{naive}} = \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k) - \nabla \mathcal{J}_{\mathbf{B}_{k-1}}(\theta_{k-1})$ 
9:   if  $\langle \Delta g_k^{\text{naive}}, \Delta \theta_{\mathbf{B}_k} \rangle > 0$  then
10:     $\gamma_k = \frac{\|\Delta \theta_k\|^2}{\langle \Delta g_k^{\text{naive}}, \Delta \theta_k \rangle}$ 
11:   else
12:     $\gamma_k = \nu$ 
13:   end if
14:    $\gamma_k = \min(\max(\gamma_k, \tilde{m}), \tilde{M})$ 
15:    $\theta_{k+1} = \theta_k - \frac{\alpha}{(k+1)^{1/2+\delta}} \gamma_k \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k)$ 
16: end for

```

Algorithm 4 Exact-GV SGD

```

1: Input:  $\alpha > 0, \nu > 0$ 
2: Input:  $\tilde{m} > 0, \tilde{M} > 0, \delta \in (0, 1/2)$ 
3: Initialize  $\theta_0 \in \mathbb{R}^P, \gamma_0 = 1$ 
4: Draw mini-batches  $(\mathbf{B}_k)_{k \in \mathbb{N}}$  independently and uniformly at random with replacement.
5:  $\theta_1 = \theta_0 - \alpha \gamma_0 \nabla \mathcal{J}_{\mathbf{B}_0}(\theta_0)$ 
6: for  $k = 1, \dots$  do
7:    $\Delta \theta_k = \theta_k - \theta_{k-1}$ 
8:    $G_k = \nabla \mathcal{J}(\theta_k) - \nabla \mathcal{J}(\theta_{k-1})$ 
9:   if  $\langle G_k, \Delta \theta_{\mathbf{B}_k} \rangle > 0$  then
10:     $\gamma_k = \frac{\|\Delta \theta_k\|^2}{\langle G_k, \Delta \theta_k \rangle}$ 
11:   else
12:     $\gamma_k = \nu$ 
13:   end if
14:    $\gamma_k = \min(\max(\gamma_k, \tilde{m}), \tilde{M})$ 
15:    $\theta_{k+1} = \theta_k - \frac{\alpha}{(k+1)^{1/2+\delta}} \gamma_k \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k)$ 
16: end for

```

Algorithm 5 Expected-GV SGD

1: **Input:** $\alpha > 0, \nu > 0$
2: **Input:** $\tilde{m} > 0, \tilde{M} > 0, \delta \in (0, 1/2)$
3: **Initialize** $\theta_0 \in \mathbb{R}^P, \gamma_0 = 1$
4: **Draw** mini-batches $(\mathbf{B}_k)_{k \in \mathbb{N}}$ independently and uniformly at random with replacement.
5: $\theta_1 = \theta_0 - \alpha \gamma_0 \nabla \mathcal{J}_{\mathbf{B}_0}(\theta_0)$
6: **for** $k = 1, \dots$ **do**
7: $\Delta\theta_k = \theta_k - \theta_{k-1}$
8: $G_k = -\frac{\alpha}{(k-1)^{1/2+\delta}} \gamma_{k-1} \mathbb{E} \left[\mathcal{C}_{\mathcal{J}_{\mathbf{B}_{k-1}}}(\theta_{k-1}) \right]$
9: **if** $\langle G_k, \Delta\theta_{\mathbf{B}_k} \rangle > 0$ **then**
10: $\gamma_k = \frac{\|\Delta\theta_k\|^2}{\langle G_k, \Delta\theta_k \rangle}$
11: **else**
12: $\gamma_k = \nu$
13: **end if**
14: $\gamma_k = \min(\max(\gamma_k, \tilde{m}), \tilde{M})$
15: $\theta_{k+1} = \theta_k - \frac{\alpha}{(k+1)^{1/2+\delta}} \gamma_k \nabla \mathcal{J}_{\mathbf{B}_k}(\theta_k)$
16: **end for**
