

Stochastic adaptive regularization with dynamic sampling for machine learning

S. Lotfi, D. Orban, A. Lodi

G–2020–51

October 2020

Revised: June 2021

La collection *Les Cahiers du GERAD* est constituée des travaux de recherche menés par nos membres. La plupart de ces documents de travail a été soumis à des revues avec comité de révision. Lorsqu'un document est accepté et publié, le pdf original est retiré si c'est nécessaire et un lien vers l'article publié est ajouté.

Citation suggérée : S. Lotfi, D. Orban, A. Lodi (Octobre 2020). Stochastic adaptive regularization with dynamic sampling for machine learning, Rapport technique, Les Cahiers du GERAD G–2020–51, GERAD, HEC Montréal, Canada.

Avant de citer ce rapport technique, veuillez visiter notre site Web (<https://www.gerad.ca/fr/papers/G-2020-51>) afin de mettre à jour vos données de référence, s'il a été publié dans une revue scientifique.

The series *Les Cahiers du GERAD* consists of working papers carried out by our members. Most of these pre-prints have been submitted to peer-reviewed journals. When accepted and published, if necessary, the original pdf is removed and a link to the published article is added.

Suggested citation: S. Lotfi, D. Orban, A. Lodi (October 2020). Stochastic adaptive regularization with dynamic sampling for machine learning, Technical report, Les Cahiers du GERAD G–2020–51, GERAD, HEC Montréal, Canada.

Before citing this technical report, please visit our website (<https://www.gerad.ca/en/papers/G-2020-51>) to update your reference data, if it has been published in a scientific journal.

La publication de ces rapports de recherche est rendue possible grâce au soutien de HEC Montréal, Polytechnique Montréal, Université McGill, Université du Québec à Montréal, ainsi que du Fonds de recherche du Québec – Nature et technologies.

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2020
– Bibliothèque et Archives Canada, 2020

The publication of these research reports is made possible thanks to the support of HEC Montréal, Polytechnique Montréal, McGill University, Université du Québec à Montréal, as well as the Fonds de recherche du Québec – Nature et technologies.

Legal deposit – Bibliothèque et Archives nationales du Québec, 2020
– Library and Archives Canada, 2020

Stochastic adaptive regularization with dynamic sampling for machine learning

Sanae Lotfi

Tiphaine Bonniot de Ruisselet

Dominique Orban

Andrea Lodi

GERAD & Department of Mathematical and Industrial Engineering, Polytechnique Montreal, Montreal, Montréal (Qc), Canada, H3T 1J4

sanae.lotfi@polymtl.ca

tiphaine.bonniot@gmail.com

dominique.orban@polymtl.ca

andrea.lodi@polymtl.ca

October 2020

Revised: June 2021

Les Cahiers du GERAD

G–2020–51

Copyright © 2020 GERAD, Lotfi, Orban, Lodi

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs. Les auteurs conservent leur droit d'auteur et leurs droits moraux sur leurs publications et les utilisateurs s'engagent à reconnaître et respecter les exigences légales associées à ces droits. Ainsi, les utilisateurs:

- Peuvent télécharger et imprimer une copie de toute publication du portail public aux fins d'étude ou de recherche privée;
- Ne peuvent pas distribuer le matériel ou l'utiliser pour une activité à but lucratif ou pour un gain commercial;
- Peuvent distribuer gratuitement l'URL identifiant la publication.

Si vous pensez que ce document enfreint le droit d'auteur, contactez-nous en fournissant des détails. Nous supprimerons immédiatement l'accès au travail et enquêterons sur votre demande.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*. Copyright and moral rights for the publications are retained by the authors and the users must commit themselves to recognize and abide the legal requirements associated with these rights. Thus, users:

- May download and print one copy of any publication from the public portal for the purpose of private study or research;
- May not further distribute the material or use it for any profit-making activity or commercial gain;
- May freely distribute the URL identifying the publication.

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Abstract : In this paper, we consider both first- and second-order techniques to address continuous optimization problems arising in machine learning. In the first-order case, we propose a framework of transition from deterministic or semi-deterministic to stochastic quadratic regularization methods. We leverage the two-phase nature of stochastic optimization to propose a novel first-order algorithm with adaptive sampling and adaptive step size. In the second-order case, we propose a novel stochastic damped L-BFGS method that improves on previous algorithms in the highly nonconvex context of deep learning. Both algorithms are evaluated on well-known deep learning datasets and exhibit promising performance.

Keywords: Machine learning, deep learning, stochastic optimization, adaptive sampling, adaptive regularization, adaptive step size, large-scale optimization

1 Introduction and related work

In this paper, we explore promising research directions to improve upon existing first- and second-order optimization algorithms in both deterministic and stochastic settings, with special emphasis on solving problems that arise in the machine learning context.

We consider general unconstrained stochastic optimization problems of the form

$$\min_{x \in \mathbb{R}^n} \mathbb{E}_\xi[F(x, \xi)], \quad (1)$$

where $\xi \in \mathbb{R}^d$ denotes a random variable and $F : \mathbb{R}^n \times \mathbb{R}^d \rightarrow \mathbb{R}$ is continuously differentiable and may be nonconvex. We introduce additional assumptions about F as needed. We refer to (1) as an *online* optimization problem, which indicates that we typically sample data points during the optimization process so as to use *new* samples at each iteration, instead of using a fixed dataset that is known up-front. We refer to $\mathbb{E}_\xi[F(x, \xi)]$ as the *expected risk* or the *expected loss*.

In machine learning, a special case of (1) is *empirical risk minimization*, which consists in solving

$$\min_{x \in \mathbb{R}^n} f(x), \quad f(x) := \frac{1}{N} \sum_{i=1}^N f_i(x), \quad (2)$$

where f_i is the loss function that corresponds to the i -th element of our dataset and N is the number of samples, i.e., the size of the dataset. We refer to (2) as a *finite-sum* problem and to f as the *empirical risk* or the *empirical loss*.

Stochastic Gradient Descent (SGD) [9, 52] and its variants [23, 32, 43, 49, 54], including variance-reduced algorithms [24, 31, 44, 56], are among the most popular methods for (1) and (2) in machine learning. They are the typical benchmark with respect to developments in both first- and second-order methods.

In the context of first-order methods, we focus on the case where it is not realistic to evaluate $\nabla f(x)$ exactly; rather, we have access to an approximation $g(x_k, \xi_k)$. If we consider (2), a popular choice for $g(x_k, \xi_k)$ is the *mini-batch* gradient

$$g(x_k, \xi_k) = \frac{1}{m_k} \sum_{i=1}^{m_k} \nabla f_{\xi_{k,i}}(x_k), \quad (3)$$

where ξ_k is the subset of samples considered at iteration k , from a given set of realizations of ξ , $\xi_{k,i}$ is the i -th sample of ξ_k and m_k is the batch size used at iteration k , i.e., the number of samples used to evaluate the gradient approximation. It follows that $f_{\xi_{k,i}}$ represents the loss function that corresponds to the sample $\xi_{k,i}$.

As discussed previously, stochastic optimization methods are the default choice for solving large-scale optimization problems. However, the search for the most problem-adapted hyperparameters involves significant computational costs. Asi and Duchi [2] highlighted the alarming computational and engineering energy used to find the best set of hyperparameters in training neural networks by citing 3 recent works [17, 51, 60], where the amount of time needed to tune the optimization algorithm and find the best neural structure can be equivalent to up to 750,000 central processing unit days [17]. Therefore, we need to design optimization algorithms with adaptive parameters, that automatically adjust to the nature of the problem without the need for a hyperparameter search. Blanchet et al. [4], Chen et al. [16], Larson and Billups [34], and Curtis et al. [19] proposed stochastic methods with adaptive step size using the trust-region framework. The use of an adaptive batch size throughout the optimization process is another promising direction towards adaptive algorithms. Exiting works include Byrd et al. [11], Friedlander and Schmidt [27], Hashemi et al. [29], and Bollapragada et al. [6]. We propose a novel algorithm with both adaptive step size and batch size in [Section 2](#).

Another potential area of improvement is in the context of solving machine learning problems that are highly nonconvex and ill-conditioned [10], which are more effectively treated using (approximate) second-order information. Second-order algorithms are well studied in the deterministic case [1, 20–22] but there are many areas to explore in the stochastic context that go beyond existing works [8, 12, 28, 41, 53]. Among these areas, the use of damping in L-BFGS is an interesting research direction to be leveraged in the stochastic case. Specifically, Wang et al. [55] proposed a stochastic damped L-BFGS (SdLBFGS) algorithm and proved almost sure convergence to a stationary point. However, damping does not prevent the inverse Hessian approximation H_k from being ill-conditioned [15]. The convergence of SdLBFGS may be heavily affected if the Hessian approximation becomes nearly singular during the iterations. In order to remedy this issue, Chen et al. [15] proposed to combine SdLBFGS with regularized BFGS [40]. The approach we propose in Section 3 differs.

Notation

We use $\|\cdot\| := \|\cdot\|_2$ to denote the Euclidean norm. Other types of norms will be introduced by the notation $\|\cdot\|_p$, where the value of p will be specified as needed. Capital Latin letters such as A , B , and H represent matrices, lowercase Latin letters such as s , x , and y represent vectors in \mathbb{R}^n , and lowercase Greek letters such as α , β and γ represent scalars. The operators $\mathbb{E}_\xi[\cdot]$ and $\mathbb{V}_\xi[\cdot]$ represent the expectation and variance over random variable ξ , respectively. For a symmetric matrix A , we use $A \succ 0$ and $A \succeq 0$ to indicate that A is, respectively, positive definite or positive semidefinite, and $\lambda_{\min}(A)$ and $\lambda_{\max}(A)$ to denote its smallest and largest eigenvalue, respectively. The identity matrix of appropriate size is denoted by I . We use $\lceil \cdot \rceil$ to denote the ceiling function that maps a real number y to the smallest integer greater than or equal to y .

1.1 Framework and definitions

We consider methods in which iterates x_k are updated as follows:

$$x_{k+1} = x_k + \alpha_k d_k, \quad d_k = -B_k^{-1}g(x_k, \xi_k), \quad (k \geq 0), \quad (4)$$

where $\alpha_k > 0$ is the step size, d_k is the search direction, $B_k = B_k^T \approx \nabla^2 f(x_k)$ is positive definite, and we set $H_k = B_k^{-1}$. Setting B_k to a multiple of I results in a first-order method.

A common practice in machine learning is to divide the dataset into a *training* set, that is used to train the chosen model, and a *test* set, that is used to evaluate the performance of the trained model on unseen data. The training and test sets are usually partitioned into *batches*, which are subsets of these sets. The term *epoch* refers to an optimization cycle over the entire training set, which means that all batches are used during this cycle.

We use the expression *semi-deterministic approach* to refer to an optimization algorithm that uses exact function values and inexact gradient values.

1.2 Contributions and organization

The paper is naturally divided in two parts.

Section 2 presents our contributions to first-order optimization methods wherein we define a framework to transform a semi-deterministic quadratic regularization method, i.e., one that uses exact function values and inexact gradient estimates, to a fully stochastic optimization algorithm suitable for machine learning problems. We go a step further to propose a novel first-order optimization algorithm with adaptive step size and sampling that exploits the two-phase nature of stochastic optimization. The result is a generic framework to adapt quadratic regularization methods to the stochastic setting. We demonstrate that our implementation is competitive on machine learning problems.

Section 3 addresses second-order methods. We propose a novel stochastic damped L-BFGS method. An eigenvalue control mechanism ensures that our Hessian approximations are uniformly bounded

using less restrictive assumptions than related works. We establish almost sure convergence to a stationary point, and demonstrate that our implementation is more robust than a previously proposed stochastic damped L-BFGS algorithm in the highly nonconvex case of deep learning.

Finally, [Section 4](#) draws conclusions and outlines potential research directions.

2 Stochastic adaptive regularization with dynamic sampling

This section is devoted to first-order methods and provides a framework of transition from deterministic or semi-deterministic to fully stochastic quadratic regularization methods. In [Section 2.1](#), we consider the generic unconstrained problem

$$\min_{x \in \mathbb{R}^n} f(x), \quad (5)$$

where $f : \mathbb{R}^n \rightarrow \mathbb{R}$ has Lipschitz-continuous gradient. We assume that the gradient estimate accuracy can be selected by the user, but do not require errors to converge to zero. We propose our Adaptive Regularization with Inexact Gradients (ARIG) algorithm and establish global convergence and optimal iteration complexity.

In [Section 2.2](#), we specialize our findings to a context appropriate for machine learning, and consider online optimization problems in the form (1) and (2). We prove that we can use inexact function values in ARIG without losing its convergence guarantees.

We present a stochastic variant of ARIG, as well as a new stochastic adaptive regularization algorithm with dynamic sampling that leverages Pflug’s diagnostic in [Section 2.3](#). We evaluate our implementation on machine learning problems in [Section 2.4](#).

2.1 Adaptive regularization with inexact gradients (ARIG)

We consider (5) and assume that it is possible to obtain $\nabla f(x, \omega_g) \approx \nabla f(x)$ using a user-specified relative error threshold $\omega_g > 0$, i.e.,

$$\|\nabla f(x, \omega_g) - \nabla f(x, 0)\| \leq \omega_g \|\nabla f(x, \omega_g)\|, \quad \text{with } \nabla f(x, 0) = \nabla f(x). \quad (6)$$

Geometrically, (6) requires the exact gradient to lie within a ball centered at the gradient estimate.

We define our assumptions as follows:

Assumption 1. The function f is continuously differentiable over \mathbb{R}^n .

Assumption 2. The gradient of f is Lipschitz continuous, i.e., there exists $L > 0$ such that for all $x, y \in \mathbb{R}^n$, $\|\nabla f(x) - \nabla f(y)\| \leq L \|x - y\|$.

2.1.1 Formulation for ARIG

We define $T(x, s)$ as follows:

$$T(x, s) := f(x) + \nabla f(x)^\top s \approx f(x + s).$$

At iteration k , we consider the approximate Taylor model using the inexact gradient $\nabla f(x, \omega_g)$ defined in (6):

$$\bar{T}_k(s) := f(x_k) + \nabla f(x_k, \omega_g^k)^\top s \approx T(x_k, s). \quad (7)$$

Finally, we define the approximate regularized Taylor model

$$l_k(s) := \bar{T}_k(s) + \frac{1}{2} \sigma_k \|s\|^2, \quad (8)$$

whose gradient is

$$\nabla_s l_k(s) = \nabla_s \bar{T}_k(s) + \sigma_k s = \nabla f(x_k, \omega_g^k) + \sigma_k s,$$

where $\sigma_k > 0$ is the iteration-dependent regularization parameter.

Algorithm 1 summarizes our Adaptive Regularization with Inexact Gradients method. The approach is strongly inspired by the framework laid out by Birgin et al. [3] and the use of inexact gradients in the context of trust-region methods proposed by Carter [13] and described by Conn et al. [18, §8.4]. The adaptive regularization property is apparent in step 5 of **Algorithm 1**, where σ_k is updated based on the quality of the model l_k . σ_k can be interpreted as the inverse of a trust-region radius or as an approximation of L . The main insight is that the trial step s_k is accepted if the ratio

$$\rho_k = \frac{f(x_k) - f(x_k + s_k)}{\bar{T}_k(0) - \bar{T}_k(s_k)} = \frac{f(x_k) - f(x_k + s_k)}{\frac{1}{\sigma_k} \|\nabla f(x_k, \omega_g^k)\|^2},$$

of the actual objective decrease to that predicted by the model is sufficiently large. The regularization parameter is then updated using the following scheme:

- If $\rho_k \geq \eta_2$, then the model quality is sufficiently high that we decrease σ_k so as to promote a larger step at the next iteration;
- if $\eta_1 \leq \rho_k < \eta_2$, then the model quality is sufficient to accept the trial step s_k , but insufficient to decrease σ_k (in our implementation, we increase σ_k);
- if $\rho_k < \eta_1$, then the model quality is low and s_k is rejected. We increase σ_k more aggressively than in the previous case so as to compute a conservative step at the next iteration.

Algorithm 1 Adaptive Regularization with Inexact Gradients—ARIG

1: Choose a tolerance $\epsilon > 0$, the initial regularization parameter $\sigma_0 > 0$, and the constants

$$0 < \sigma_{\min} \leq \sigma_0, \quad 0 < \eta_1 \leq \eta_2 < 1 \quad \text{and} \quad 0 < \gamma_1 < 1 < \gamma_2 < \gamma_3. \quad (9)$$

Set $k = 0$.

2: Choose ω_g^k such that $0 < \omega_g^k \leq 1/\sigma_k$ and compute $\nabla f(x_k, \omega_g^k)$ such that (6) holds. If $\|\nabla f(x_k, \omega_g^k)\| \leq \epsilon/(1 + \omega_g^k)$, terminate with the approximate solution $x_\epsilon = x_k$.

3: Compute the step $s_k = -\frac{1}{\sigma_k} \nabla f(x_k, \omega_g^k)$.

4: Evaluate $f(x_k + s_k)$ and define

$$\rho_k = \frac{f(x_k) - f(x_k + s_k)}{\bar{T}_k(0) - \bar{T}_k(s_k)} = \frac{f(x_k) - f(x_k + s_k)}{\frac{1}{\sigma_k} \|\nabla f(x_k, \omega_g^k)\|^2}. \quad (10)$$

If $\rho_k \geq \eta_1$, set $x_{k+1} = x_k + s_k$. Otherwise, set $x_{k+1} = x_k$.

5: Set

$$\sigma_{k+1} \in \begin{cases} [\max(\sigma_{\min}, \gamma_1 \sigma_k), \sigma_k] & \text{if } \rho_k \geq \eta_2, \\ [\sigma_k, \gamma_2 \sigma_k] & \text{if } \eta_1 \leq \rho_k < \eta_2, \\ [\gamma_2 \sigma_k, \gamma_3 \sigma_k] & \text{if } \rho_k < \eta_1. \end{cases} \quad (11)$$

Increment k by one and go to step 2 if $\rho_k \geq \eta_1$ or to step 3 otherwise.

2.1.2 Convergence and complexity analysis

The condition (6) ensures that at each iteration k ,

$$\|\nabla f(x_k)\| \leq \|\nabla f(x_k) - \nabla f(x_k, \omega_g^k)\| + \|\nabla f(x_k, \omega_g^k)\| \leq (1 + \omega_g^k) \|\nabla f(x_k, \omega_g^k)\|.$$

Therefore, when the termination occurs, the condition $\|\nabla f(x_k, \omega_g^k)\| \leq \epsilon/(1 + \omega_g^k)$ implies $\|\nabla f(x_\epsilon)\| \leq \epsilon$ and x_ϵ is an approximate first-order critical point to within the desired stopping condition.

The following analysis represents the adaptation of the general properties presented by [3] to a first-order model with inexact gradients.

We first obtain an upper bound on σ_k . The following result is similar to Birgin et al. [3, Lemma 2.2]. All proofs may be found in [Appendix A](#).

Lemma 1. For all $k \geq 0$,

$$\sigma_k \leq \sigma_{\max} := \max \left[\sigma_0, \frac{\gamma_3(\frac{1}{2}L + 1)}{1 - \eta_2} \right]. \quad (12)$$

We now bound the total number of iterations in terms of the number of successful iterations and state our main evaluation complexity result.

Theorem 1 (3, Theorem 2.5). Let [Assumptions 1](#) and [2](#) be satisfied. Assume that there exists κ_{low} such that $f(x_k) \geq f_{\text{low}}$ for all $k \geq 0$. Assume also that $\omega_g^k \leq 1/\sigma_k$ for all $k \geq 0$. Then, given $\epsilon > 0$, [Algorithm 1](#) needs at most

$$\left\lceil \kappa_s \frac{f(x_0) - f_{\text{low}}}{\epsilon^2} \right\rceil$$

successful iterations (each involving one evaluation of f and its approximate gradient) and at most

$$\left\lceil \kappa_s \frac{f(x_0) - f_{\text{low}}}{\epsilon^2} \right\rceil \left(1 + \frac{|\log \gamma_1|}{\log \gamma_2} \right) + \frac{1}{\log \gamma_2} \log \left(\frac{\sigma_{\max}}{\sigma_0} \right)$$

iterations in total to produce an iterate x_ϵ such that $\|\nabla f(x_\epsilon)\| \leq \epsilon$, where σ_{\max} is given by [Lemma 1](#) and where

$$\kappa_s = \frac{(1 + \sigma_{\max})^2}{\eta_1 \sigma_{\min}}.$$

[Theorem 1](#) implies that $\liminf \|\nabla f(x_k)\| = 0$ provided $\{f(x_k)\}$ is bounded below. The following corollary results immediately.

Corollary 1. Under the assumptions of [Theorem 1](#), either $f(x_k) \rightarrow -\infty$ or $\liminf \|\nabla f(x_k)\| = 0$.

2.2 Stochastic adaptive regularization with dynamic sampling

Although ARIG enjoys strong theoretical guarantees under weak assumptions, it is not adapted to the context of machine learning, where we do not have access to exact objective values. Therefore, we propose to build upon ARIG in order to design a stochastic first-order optimization with adaptive regularization with similar convergence guarantees.

2.2.1 Adaptive regularization with inexact gradient and function values

Conn et al. [[18](#), §10.6] provide guidelines to allow for inexact objective values in trust-region methods while preserving convergence properties. This section adapts their strategy to the context of [Algorithm 1](#). Assume that we do not have access to direct evaluations of $f(x)$, but we can obtain an approximation $f(x_k, \omega_f^k)$ depending on a user-specified error threshold ω_f^k , such that

$$|f(x, \omega_f^k) - f(x, 0)| \leq \omega_f^k, \quad \text{with} \quad f(x, 0) = f(x). \quad (13)$$

We redefine the approximate Taylor series ([7](#)) using inexact gradient and function values:

$$\bar{T}_k(s) = f(x, \omega_f^k) + \nabla f(x_k, \omega_g^k)^\top s. \quad (14)$$

If the approximations of $f(x_k)$ and $f(x_k + s_k)$ are not sufficiently accurate, the ratio ρ_k defined in ([10](#)) loses its meaning, since it is supposed to quantify the ratio of the *actual decrease* of the function value to the *predicted decrease* by the model. To circumvent the difficulty, we have the following result whose proof is elementary.

Lemma 2. Let the constant η_0 be such that $\eta_0 < \frac{1}{2}\eta_1$ and let s_k be a step taken from x_k such that $\bar{T}_k(s_k) < \bar{T}_k(0)$. If ω_f^k and $\hat{\omega}_f^k$ satisfy,

$$\max(\omega_f^k, \hat{\omega}_f^k) \leq \eta_0 [\bar{T}_k(0) - \bar{T}_k(s_k)], \quad (15)$$

then, a sufficient decrease in the inexact function values,

$$\bar{\rho}_k := \frac{\bar{f}(x_k, \omega_f^k) - \bar{f}(x_k + s_k, \hat{\omega}_f^k)}{\bar{T}_k(0) - \bar{T}_k(s_k)} \geq \eta_1,$$

implies a sufficient decrease in the actual function values,

$$\rho_k = \frac{f(x_k) - f(x_k + s_k)}{\bar{T}_k(0) - \bar{T}_k(s_k)} \geq \bar{\eta}_1,$$

where $\bar{\eta}_1 = \eta_1 - 2\eta_0 > 0$.

Lemma 2 states that whenever (15) is satisfied, a sufficient decrease in the *stochastic* objective, i.e., $\bar{\rho}_k \geq \eta_1$, implies that the *actual* decrease in the exact objective is at least a fraction $\bar{\eta}_1$ of the decrease predicted by the model. It also means that, as long as we can guarantee that (15) is satisfied at each iteration, the convergence and complexity results from the previous section continue to apply.

In the next section, we use adaptive sampling to satisfy (6) in expectation with respect to the batch used to obtain $\nabla f(x, \omega_g)$ without evaluation of the full gradient.

2.2.2 Stochastic ARIG with adaptive sampling

To use a stochastic version of ARIG in machine learning, we must compute a gradient approximation $g(x_k, \xi_k)$ for mini-batch ξ_k ,

$$g(x_k, \xi_k) = \frac{1}{m_k} \sum_{i=1}^{m_k} g(x_k, \xi_{k,i}), \quad (16)$$

that satisfies

$$\|g(x_k, \xi_k) - \nabla f(x_k)\| \leq \omega_g^k \|g(x_k, \xi_k)\|, \quad (17)$$

where ω_g^k is an iteration-dependent error threshold.

Condition (17) is sometimes called the *norm test*. We can use the adaptive sampling strategy of Byrd et al. [11] to maintain this condition satisfied along the iterations. Let us introduce their adaptive sampling strategy. Since ω_g^k should verify $0 < \omega_g^k \leq 1/\sigma_k$, let us fix it to $\omega_g^k := 1/\sigma_k$. This choice is natural since a higher error-threshold would require the use of fewer samples to compute $g(x_k, \xi_k)$. Since $g(x_k, \xi_k)$ is obtained as a sample average (16), it is an unbiased estimate of $\nabla f(x_k)$. Therefore,

$$\mathbb{E}_{\xi_k} [\|g(x_k, \xi_k) - \nabla f(x_k)\|^2] = \|\mathbb{V}_{\xi_k} [g(x_k, \xi_k)]\|_1, \quad (18)$$

where $\mathbb{V}_{\xi_k} [g(x_k, \xi_k)]$ is a vector with components $\mathbb{V}_{\xi_k} [g_j(x_k, \xi_k)]$, $j = 1, \dots, n$. Freund and Walpole [26, p.183] establish that,

$$\mathbb{V}_{\xi_k} [g(x_k, \xi_k)] = \frac{\mathbb{V}_{i \in \xi_k} (\nabla f_i(x_k))}{m_k} \frac{(N - m_k)}{N - 1}. \quad (19)$$

The direct computation of the population variance $\mathbb{V}_{i \in \xi_k} (\nabla f_i(x_k))$ is expensive. Hence, Byrd et al. [11] suggest to estimate it with the sample variance

$$\hat{\mathbb{V}}_{i \in \xi_k} (\nabla f_i(x_k)) = \frac{1}{m_k - 1} \sum_{i \in \xi_k} (\nabla f_i(x_k) - g(x_k, \xi_k))^2, \quad (20)$$

where the square is applied elementwise. It follows from (18) and (19) that

$$\mathbb{E}_{\xi_k} [\|g(x_k, \xi_k) - \nabla f(x_k)\|^2] = \frac{\|\hat{\mathbb{V}}_{i \in \xi_k}(\nabla f_i(x_k))\|_1}{m_k} \frac{(N - m_k)}{N - 1}. \quad (21)$$

In the context of large-scale optimization, we let $N \rightarrow \infty$, and (17) implies

$$\frac{\|\hat{\mathbb{V}}_{i \in \xi_k}(\nabla f_i(x_k))\|_1}{m_k} \leq \frac{1}{\sigma_k^2} \|g(x_k, \xi_k)\|^2. \quad (22)$$

Finally, our strategy to obtain $g(x_k, \xi_k)$ that satisfies (17), with $\omega_g^k = 1/\sigma_k$, is as follows:

1. at iteration k , sample a new batch of size $m_k = m_{k-1}$ and compute the sample variance $\hat{\mathbb{V}}_{i \in \xi_k}$ and mini-batch gradient $g(x_k, \xi_k)$;
2. if (22) holds, use mini-batch m_k at iteration k . Otherwise, sample a new batch $\tilde{\xi}$ of size

$$\tilde{m}_k := \left\lceil \frac{\sigma_k^2 \|\hat{\mathbb{V}}_{i \in \xi_k}(\nabla f_i(x_k))\|_1}{\|g(x_k, \xi_k)\|^2} \right\rceil. \quad (23)$$

It is important to note that the strategy above relies on the assumption that the change in the batch size is gradual. Therefore, we assume that

$$\|\hat{\mathbb{V}}_{i \in \tilde{\xi}_k}(\nabla f_i(x_k))\|_1 \approx \|\hat{\mathbb{V}}_{i \in \xi_k}(\nabla f_i(x_k))\|_1 \quad \text{and} \quad \|g(x_k, \tilde{\xi}_k)\| \approx \|g(x_k, \xi_k)\|. \quad (24)$$

Lotfi [38] provides a comparison between stochastic ARIG with adaptive sampling and SGD on the MNIST dataset [35].

Theoretically speaking, convergence requires that either the variance of the gradient estimates or the step size decays to zero. For the strategy defined above, the step size does not converge to zero. Moreover, the percentage batch size would need to be close to 100% for the gradient variance to converge to zero. This discussion motivates our work in Section 2.3, where we propose a new algorithm that leverages both adaptive regularization and adaptive sampling, but not simultaneously. Although we do not provide theoretical guarantees, it is conceivable that the new algorithm globally converges to a stationary points since we choose a step size that converges to zero in the second phase of the optimization [10].

2.3 Stochastic adaptive regularization with dynamic sampling and convergence diagnostics

In this section, we present a novel algorithm that incorporates both adaptive regularization and adaptive sampling by leveraging the two-phase nature of stochastic optimization.

2.3.1 Statistical testing for stochastic optimization

Stochastic optimization algorithms exhibit two phases: a transient phase and a stationary phase [42]. The transient phase is characterized by the convergence of the algorithm to a region of interest that contains potentially a good local or global optimum. In the stationary phase, the algorithm oscillates around that optimum in the region of interest.

Several authors aimed to derive a stationary condition to indicate when stochastic procedures leave the transient phase and enter the stationary phase. Some of the early works in this line of research are Pflug's works [46–48]. In particular, Pflug's procedure was introduced in Pflug [47], and it consists of keeping a running average of the inner product of successive gradients in order to detect the end of

the transient phase. Pflug’s idea is simple: when the running average becomes negative, that suggests that the consecutive stochastic gradients are pointing in different directions. This should signal that the algorithm entered the stationary phase where it oscillates around a local optimum. More recently, Chee and Toulis [14] developed a statistical diagnostic test in order to detect when Stochastic Gradient Descent (SGD) enters the stationary phase, by combining the statistical test from Pflug [47] with the SGD update. They presented theoretical arguments and practical experiments to prove that the test activation occurs in the phase-transition for SGD. Algorithm 2 describes Pflug’s diagnostic proposed by Chee and Toulis [14] for convergence of SGD applied to (2), where $g(x_k)$ represents an approximation to the full gradient $\nabla f(x_k)$. Note that in this case, ξ_k contains a single sample. Chee and Toulis prove that the convergence diagnostic in Algorithm 2 satisfies $\mathbb{E}(S_k - S_{k-1}) < 0$ as $k \rightarrow +\infty$, which implies that the algorithm terminates almost surely.

Algorithm 2 Pflug diagnostic for convergence of SGD

Require: $x_0 \in \mathbb{R}^n$
1: Choose $\gamma > 0$ and a positive integer named “burn-in”.
2: Set $S_0 = 0$. Compute $g(x_0)$ and $x_1 = x_0 - \gamma g(x_0)$.
3: **for** $k = 0, 1, \dots$ **do**
4: Compute $g(x_k)$ and $x_{k+1} = x_k - \gamma g(x_k)$.
5: Define $S_k = S_{k-1} + g(x_k)^\top g(x_{k-1})$.
6: **if** $k > \text{burn-in}$ and $S_k < 0$ **then**
7: return k .
8: **end if**
9: **end for**

2.3.2 Adaptive regularization and sampling using Pflug diagnostic

To harness Pflug’s diagnostic, we discuss the bias-variance trade-off in machine learning. The bias of a model is the extent to which it approximates a global optimum, whereas its variance is the amount of variation in the solution if the training set changes. The generalization error in machine learning is the sum of the bias and the variance, hence the trade-off between the two quantities. Now notice that

- in the transient phase, we would like to reduce the bias of the model by converging to a promising region of interest. Therefore, we propose to use adaptive regularization in order to adapt the step size automatically, taking into account the ratio ρ_k . We use a fixed batch size during this phase.
- In the stationary phase, we would like to reduce the variance of the gradient estimates. Therefore, we use adaptive sampling as a variance reduction technique. Additionally, we choose a step size sequence that converges to zero to ensure global convergence to a stationary point.

We call this algorithm ARAS for Adaptive Regularization and Adaptive Sampling algorithm. Algorithm 3 states the complete procedure. Note that we reduce the number of parameters that were initially introduced in ARIG by eliminating γ_3 and η_2 , which simplifies the algorithm. Inspired by Curtis et al. [19], we do not require the decrease to be sufficient to accept a step. Instead, we accept the step in all cases. We also use a constant maximum batch size, m_{\max} to ensure that the batch sizes stay reasonable.

2.4 Experimental evaluation

We study the performance of ARAS in both convex and nonconvex settings. First, we compare it to SGD and SGD with momentum in the convex setting, where we consider a logistic regression problem. Then, we compare two versions of ARAS to SGD in the nonconvex setting, where we consider a nonconvex support vector machine problem with a sigmoid loss function.

Algorithm 3 Adaptive Regularization and Sampling using Pflug diagnostic—ARAS

1: Choose the initial regularization parameter $\sigma_0 > 0$, the initial batch size m_0 , the maximum batch size m_{\max} , the total number of epochs N_{epochs} , burn-in > 0 and the constants $0 < \sigma_{\min} \leq \sigma_0$, $0 < \eta < 1$, and $0 < \gamma_1 < 1 < \gamma_2$. Set $k = 0$, $t = 2$, $S_0 = 0$, and Transient = True.

2: **for** $n = 1, \dots, N_{\text{epochs}}$ **do**

3: **if** Transient is True **then**

4: Sample ξ_k and compute $g(x_k, \xi_k)$.

5: Compute $s_k = -\frac{1}{\sigma_k}g(x_k, \xi_k)$ and define $x_{k+1} = x_k + s_k$.

6: Evaluate $f(x_{k+1}, \xi_k)$ and $g(x_{k+1}, \xi_k)$, and define

$$\bar{\rho}_k = \frac{f(x_k, \xi_k) - f(x_{k+1}, \xi_k)}{\bar{T}_k(0) - \bar{T}_k(s_k)} = \frac{f(x_k, \xi_k) - f(x_{k+1}, \xi_k)}{\frac{1}{\sigma_k} \|g(x_k, \xi_k)\|^2}. \quad (25)$$

7: Set

$$\sigma_{k+1} \in \begin{cases} [\max(\sigma_{\min}, \gamma_1 \sigma_k), \sigma_k] & \text{if } \bar{\rho}_k \geq \eta, \\ [\sigma_k, \gamma_2 \sigma_k] & \text{if } \bar{\rho}_k < \eta. \end{cases} \quad (26)$$

8: Compute $S_{k+1} = S_k + g(x_{k+1}, \xi_k)^\top g(x_k, \xi_k)$ and set $m_{k+1} = m_k$.

9: **If** ($k > \text{burn-in}$ and $S_{k+1} < 0$), then set Transient = False.

10: **else**

11: Sample ξ_k . Compute $g(x_k, \xi_k)$ and $\|V_{i \in \xi_k}(\nabla f_i(x_k))\|_1$.

12: **if** (22) holds **then**

13: Go to step 18.

14: **else**

15: Set

$$m_k = \min \left(\left\lceil \frac{\sigma_k^2 \|V_{i \in \xi_k}(\nabla f_i(x_k))\|_1}{\|g(x_k, \xi_k)\|^2} \right\rceil, m_{\max} \right)$$

16: Sample a new ξ_k of new size m_k and compute $g(x_k, \xi_k)$.

17: **end if**

18: Compute the step $s_k = -\frac{1}{\sigma_k}g(x_k, \xi_k)$ and define $x_{k+1} = x_k + s_k$.

19: Set $m_{k+1} = m_k$, $\sigma_{k+1} = \sigma_k \frac{t}{t-1}$, and increment t by one.

20: **end if**

21: Increment k by one and go to step 1.

22: **end for**

2.4.1 Logistic regression with MNIST

To evaluate the performance of our algorithm in the convex setting, we compare it to SGD and SGD with momentum for solving a logistic regression problem on the MNIST dataset. All parameters were set using a grid-search and all algorithms were trained for 10 epochs each.

Figure 1 shows the performance of the three algorithms. ARAS outperforms SGD and SGD with momentum on both the training and test sets.

2.4.2 Nonconvex support vector machine with RCV1

In the nonconvex setting, we compare the performance of ARAS and SGD for solving the finite-sum version of a nonconvex support vector machine problem, with a sigmoid loss function, a problem that was considered in Wang et al. [55]:

$$\min_{x \in \mathbb{R}^n} \mathbb{E}_{u,v} [1 - \tanh(v x^\top u)] + \lambda \|x\|^2, \quad (27)$$

where $u \in \mathbb{R}^n$ represents the feature vector, $v \in \{-1, 1\}$ represents the label value, and λ is the regularization parameter.

The finite-sum version of (27) can be written as

$$\min_{x \in \mathbb{R}^n} \frac{1}{N} \sum_{i=1}^N f_i(x) + \lambda \|x\|^2, \quad (28)$$

where $f_i(x) = 1 - \tanh(v_i x^\top u_i)$ for $i = 1, \dots, N$.

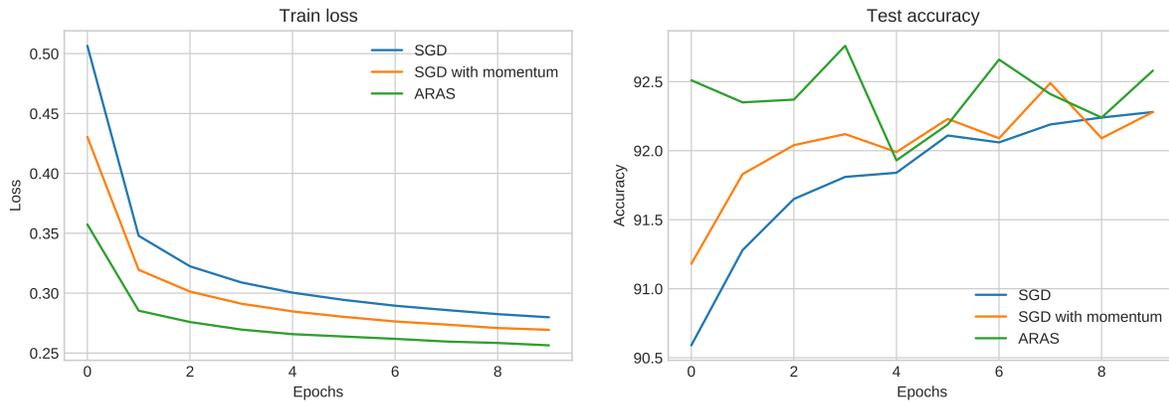


Figure 1: Evolution of the training loss (left) and the validation accuracy (right) for ARAS, SGD and SGD with momentum, solving a logistic regression problem on MNIST.

We use a subset¹ of the dataset Reuters Corpus Volume I (RCV1) [36], which is a collection of manually categorized newswire stories. The subset contains 9625 stories with 29992 distinct words. There are four categories that we reduce to the binary categories $\{-1, 1\}$.² We then solve this binary classification problem.

Figure 2 reports our results; ARAS largely outperforms SGD both in terms of the training loss and test accuracy.

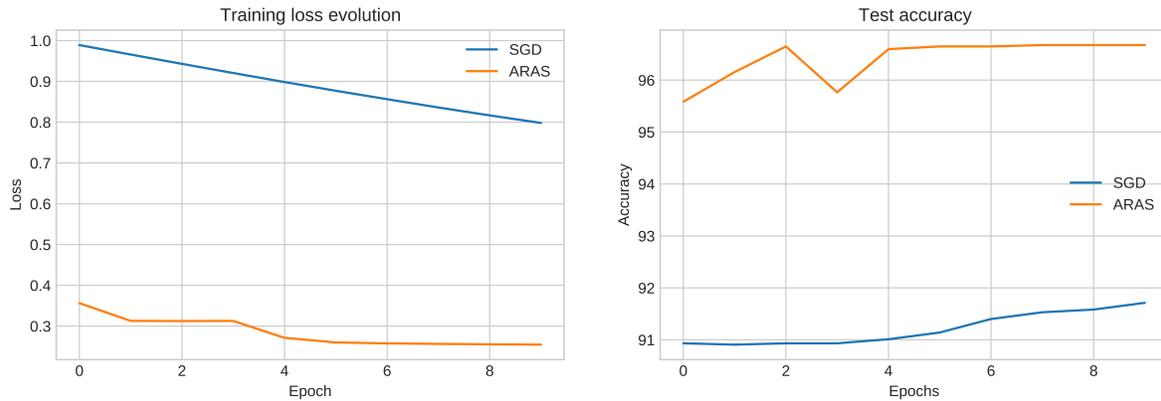


Figure 2: Evolution of the training loss (left) and the test accuracy (right) for SGD and ARAS solving a nonconvex SVM problem on RCV1.

3 Stochastic damped L-BFGS with controlled norm of the Hessian approximation

In this section, we consider second-order methods and we assume that, at iteration k , we can obtain a stochastic approximation (3) of $\nabla f(x_k)$. Iterates are updated according to (4), where this time $B_k \neq I$.

The stochastic BFGS method [53] computes an updated approximation H_{k+1} according to

$$H_{k+1} = V_k H_k V_k^\top + \rho_k s_k s_k^\top, \quad \text{where} \quad V_k = I - \rho_k s_k y_k^\top \quad \text{and} \quad \rho_k = 1/s_k^\top y_k, \quad (29)$$

¹We downloaded the subset from <http://www.cad.zju.edu.cn/home/dengcai/Data/TextData.html>

²The two categories “MCAT” and “ECAT” correspond to label value 1, and the two categories “C15” and “GCAT” correspond to label value -1 .

which ensures that the secant equation $B_{k+1}s_k = y_k$ is satisfied, where

$$s_k = x_{k+1} - x_k, \quad \text{and} \quad y_k = g(x_{k+1}, \xi_k) - g(x_k, \xi_k). \quad (30)$$

If $H_k \succ 0$ and the curvature condition $s_k^\top y_k > 0$ holds, then $H_{k+1} \succ 0$ [25].

Because storing H_k and performing matrix-vector products is costly for large-scale problems, we use the limited-memory version of BFGS (L-BFGS) [37, 45], in which H_k only depends on the most recent p iterations and an initial $H_k^0 \succ 0$. The parameter p is the *memory*. The inverse Hessian update can be written as

$$H_k = (V_{k-1}^\top \cdots V_{k-p}^\top) H_k^0 (V_{k-p} \cdots V_{k-1}) + \rho_{k-p} (V_{k-1}^\top \cdots V_{k-p+1}^\top) s_{k-p} s_{k-p}^\top (V_{k-p+1} \cdots V_{k-1}) + \cdots + \rho_{k-1} s_{k-1} s_{k-1}^\top. \quad (31)$$

When α_k in (4) is not computed using a Wolfe line search [57, 58], there is no guarantee that the curvature condition holds. A common strategy is to simply skip the update. By contrast, Powell [50] proposed *damping*, which consists in updating H_k using a modified y_k , denoted by \hat{y}_k , to benefit from information discovered at iteration k while ensuring sufficient positive definiteness. We use

$$\hat{y}_k := \theta_k y_k + (1 - \theta_k) B_{k+1}^0 s_k, \quad (32)$$

which is inspired by Wang et al. [55], and differs from the original proposal of Powell [50], where

$$\theta_k = 1 \quad \text{if} \quad s_k^\top y_k \geq \eta s_k^\top B_{k+1}^0 s_k, \quad \text{and} \quad (1 - \eta) \frac{s_k^\top B_{k+1}^0 s_k}{s_k^\top B_{k+1}^0 s_k - s_k^\top y_k} \quad \text{otherwise}, \quad (33)$$

with $\eta \in (0, 1)$ and $B_{k+1}^0 := (H_{k+1}^0)^{-1}$. The choice (32) ensures that the curvature condition

$$s_k^\top \hat{y}_k \geq \eta s_k^\top B_{k+1}^0 s_k \geq \eta \lambda_{\min}(B_{k+1}^0) \|s_k\|^2 > 0, \quad (34)$$

is always satisfied since $H_{k+1}^0 \succ 0$. We obtain the damped L-BFGS update, which is (31) with each V_i and ρ_i replaced with $\hat{V}_i = I - \hat{\rho}_i s_i^\top \hat{y}_i$ and $\hat{\rho}_i = 1/s_i^\top \hat{y}_i$.

In the next section, we propose a new version of stochastic damped L-BFGS that maintains estimates of the smallest and largest eigenvalues of H_k . We show that the new algorithm requires less restrictive assumptions with respect to current literature, while being more robust to ill-conditioning and maintaining favorable convergence properties and complexity as established in Section 3.2. Finally, Section 3.3 provides a detailed experimental evaluation.

3.1 A new stochastic damped L-BFGS with controlled Hessian norm

Our working assumptions are Assumptions 1 and 2. We begin by deriving bounds on the smallest and largest eigenvalues of H_{k+1} as functions of bounds on those of H_k . Proofs can be found in Appendix A.

Lemma 3. Let s and $y \in \mathbb{R}^n$ such that $s^\top y \geq \gamma \|s\|^2$ with $\gamma > 0$, and such that $\|y\| \leq L_y \|s\|$, with $L_y > 0$. Let $A = \mu V V^\top + \rho s s^\top$, where $\rho = 1/s^\top y$, $\mu > 0$, and $V = I - \rho s y^\top$. Then,

$$0 < \min \left(\frac{1}{L_y}, \frac{\mu}{1 + \frac{\mu}{\gamma} L_y^2} \right) \leq \lambda_{\min}(A) \leq \lambda_{\max}(A) \leq \frac{1}{\gamma} + \max \left(0, \frac{\mu}{\gamma^2} L_y^2 - \frac{\mu}{1 + \frac{\mu}{\gamma} L_y^2} \right).$$

In order to use Lemma 3 to obtain bounds on the eigenvalues of H_{k+1} , we make the following assumption:

Assumption 3. There is $L_g > 0$ such that for all $x, y \in \mathbb{R}^n$, $\|g(x, \xi) - g(y, \xi)\| \leq L_g \|x - y\|$.

Assumption 3 is required to prove convergence and convergence rates for most recent stochastic quasi-Newton methods [59]. It is less restrictive than requiring $f(x, \xi)$ to be twice differentiable with respect to x , and the Hessian $\nabla_{xx}^2 f(x, \xi)$ to exist and be bounded for all x and ξ , as in Wang et al. [55].

The next theorem shows that the eigenvalues of H_{k+1} are bounded and bounded away from zero.

Theorem 2. Let **Assumptions 1 to 3** hold. Let $H_{k+1}^0 \succ 0$ and $p > 0$. If H_{k+1} is obtained by applying p times the damped BFGS update formula with inexact gradient to H_{k+1}^0 , there exist easily computable constants λ_{k+1} and Λ_{k+1} that depend on L_g and H_{k+1}^0 such that $0 < \lambda_{k+1} \leq \lambda_{\min}(H_{k+1}) \leq \lambda_{\max}(H_{k+1}) \leq \Lambda_{k+1}$.

The precise form of λ_{k+1} and Λ_{k+1} is given in (48) and (50) in **Appendix A**.

A common choice for H_{k+1}^0 is $H_{k+1}^0 = \gamma_{k+1}^{-1} I$ where $\gamma_{k+1} = y_k^\top y_k / s_k^\top y_k$, is the *scaling parameter*. This choice ensures that the search direction is well scaled, which promotes large steps. To keep H_{k+1}^0 from becoming nearly singular or non positive definite, we define

$$H_{k+1}^0 = \left(\max(\underline{\gamma}_{k+1}, \min(\gamma_{k+1}, \bar{\gamma}_{k+1})) \right) I, \quad (35)$$

where $0 < \underline{\gamma}_{k+1} < \bar{\gamma}_{k+1}$ can be constants or iteration dependent.

The Hessian-gradient product used to compute $d_k = -H_k g(x_k, \xi_k)$ can be obtained cheaply by exploiting a recursive algorithm [39, 45].

Motivated by the success of recent methods combining variance reduction with stochastic L-BFGS [28, 41, 55], we apply an SVRG-like type of variance reduction [31] to the update. Not only does this accelerate the convergence, since we can choose a constant step size, but it also improves the quality of the curvature approximation.

We summarize our complete algorithm, VARIance-Reduced stochastic damped L-BFGS with Controlled Hessian Norm (VARCHEN), as **Algorithm 4**.

Algorithm 4 Variance-Reduced Stochastic Damped L-BFGS with Controlled Hessian Norm

- 1: Choose $x_0 \in \mathbb{R}^n$, step size sequence $\{\alpha_k > 0\}_{k \geq 0}$, batch size sequence $\{m_k > 0\}_{k \geq 0}$, eigenvalue limits $\lambda_{\max} > \lambda_{\min} > 0$, memory parameter p , total number of epochs N_{epochs} , and sequences $\{\underline{\gamma}_k > 0\}_{k \geq 0}$ and $\{\bar{\gamma}_{k+1} > 0\}_{k \geq 0}$, such that $0 < \lambda_{\min} < \underline{\gamma}_k < \bar{\gamma}_k < \lambda_{\max}$, for every $k \geq 0$. Set $k = 0$ and $H_0 = I$.
 - 2: **for** $t = 1, \dots, N_{\text{epochs}}$ **do**
 - 3: Define $x_k^t = x_k$ and compute the full gradient $\nabla f(x_k^t)$. Set $M = 0$.
 - 4: **while** $M < N$ **do**
 - 5: Sample batch ξ_k of size $m_k \leq N - M$ and compute $g(x_k, \xi_k)$ and $g(x_k^t, \xi_k)$.
 - 6: Define $\tilde{g}(x_k, \xi_k) = g(x_k, \xi_k) - g(x_k^t, \xi_k) + \nabla f(x_k^t)$.
 - 7: Estimate Λ_k and λ_k in **Theorem 2**. If $\Lambda_k > \lambda_{\max}$ or $\lambda_k < \lambda_{\min}$, delete s_i, y_i and \hat{y}_i for $i = k - p + 1, \dots, k - 2$.
 - 8: Compute $d_k = -H_k \tilde{g}(x_k, \xi_k)$.
 - 9: Define $x_{k+1} = x_k + \alpha_k d_k$, and compute s_k, y_k as in (30), and \hat{y}_k as in (32).
 - 10: Increment k by one and update $M \leftarrow M + m_k$.
 - 11: **end while**
 - 12: **end for**
-

In step 7 of **Algorithm 4**, we compute an estimate of the upper and lower bounds on $\lambda_{\max}(H_k)$ and $\lambda_{\min}(H_k)$, respectively. The only unknown quantity in the expressions of Λ_k and λ_k in **Theorem 2** is L_g , which we estimate as $L_g \approx L_{g,k} := \|y_k\| / \|s_k\|$. When the estimates are not within the limits $[\lambda_{\min}, \lambda_{\max}]$, we delete s_i, y_i and $\hat{y}_i, i \in \{k - p + 1, \dots, k - 2\}$ from storage, such that $H_k g(x_k, \xi_k)$ is computed using the most recent pair (s_{k-1}, \hat{y}_{k-1}) only and $d_k = -H_k g(x_k, \xi_k)$. Although it is not theoretically guaranteed that the eigenvalues of the new Hessian approximation using the most recent pair are within the bounds, we found that it is the case in practice and that this strategy yields better results than setting the approximation to the identity. Finally, a full gradient is computed once in every epoch in step 3. The term $g(x_k, \xi_k) - \nabla f(x_k^t)$ can be seen as the bias in the gradient estimation $g(x_k, \xi_k)$, and it is used here to correct the gradient approximation in step 6.

3.2 Convergence and complexity analysis

We show that [Algorithm 4](#) satisfies the assumptions of the convergence analysis and iteration complexity of Wang et al. [55] for stochastic quasi-Newton methods. We make an additional assumption used by Wang et al. [55] to establish global convergence.

Assumption 4. For all k , ξ_k is independent of $\{x_1, \dots, x_k\}$, $\mathbb{E}_{\xi_k} [g(x_k, \xi_k)] = \nabla f(x_k)$, and there exists $\sigma > 0$ such that $\mathbb{E}_{\xi_k} [\|g(x_k, \xi_k) - \nabla f(x_k)\|^2] \leq \sigma^2$.

Our first result follows from Wang et al. [55, Theorem 2.6], whose remaining assumptions are satisfied as a consequence of (3), [Theorem 2](#), the mechanism of [Algorithm 4](#), (29) and our choice of α_k below.

Theorem 3. Assume $m_k = m$ for all k , that [Assumptions 1 to 4](#) hold for $\{x_k\}$ generated by [Algorithm 4](#), and that $\alpha_k := c/(k+1)$ where $0 < c \leq \lambda_{\min}/(L\lambda_{\max})$. Then, $\liminf \|\nabla f(x_k)\| = 0$ with probability 1. Moreover, there is $M_f > 0$ such that $\mathbb{E}[f(x_k)] \leq M_f$ for all k . If we additionally assume that there exists $M_g > 0$ such that $\mathbb{E}_{\xi_k} [\|g(x_k, \xi_k)\|^2] \leq M_g$, then $\lim \|\nabla f(x_k)\| = 0$ with probability 1.

Our next result follows in the same way from Wang et al. [55, Theorem 2.8].

Theorem 4. Let assumptions of [Theorem 3](#) be satisfied. Assume in addition that there exists f_{low} such that $f(x_k) \geq f_{\text{low}}$ for all k . Let $\alpha_k := \lambda_{\min}/(L\lambda_{\max}^2)k^{-\beta}$ for all $k > 0$, with $\beta \in (\frac{1}{2}, 1)$. Then, for all $T > 0$,

$$\frac{1}{T} \sum_{k=1}^T \mathbb{E} [\|\nabla f(x_k)\|^2] \leq \frac{2L(M_f - f_{\text{low}})\Lambda^2}{\lambda^2} T^{\beta-1} + \frac{\sigma^2}{(1-\beta)m} (T^{-\beta} - T^{-1}). \quad (36)$$

Moreover, for any $\epsilon \in (0, 1)$, we achieve

$$\frac{1}{T} \sum_{k=1}^T \mathbb{E} [\|\nabla f(x_k)\|^2] \leq \epsilon.$$

after at most $T = \mathcal{O}(\epsilon^{-1/(1-\beta)})$ iterations.

3.3 Experimental evaluation

We compare VARCHEN to SdLBFGS-VR [55] and to SVRG [31] for solving a multi-class classification problem. We train a modified version³ of the deep neural network model DavidNet⁴ proposed by David C. Page, on CIFAR-10 [33] for 20 epochs. Note that we also used VARCHEN and SdLBFGS-VR to solve a logistic regression problem using the MNIST dataset [35] and a nonconvex support-vector machine problem with a sigmoid loss function using the RCV1 dataset [36]. The performance of both algorithms are on par on those problems because, in contrast with DavidNet on CIFAR-10, they are not highly nonconvex or ill conditioned.

[Figure 3](#) shows that VARCHEN outperforms SdLBFGS-VR for the training loss minimization task, and both outperform SVRG. VARCHEN has an edge over SdLBFGS-VR in terms of the validation accuracy, and both outperform SVRG. More importantly, the performance of VARCHEN is more consistent than that of SdLBFGS-VR, displaying a smoother, less oscillatory behaviour. To further investigate this observation, we plot the evolution of Λ_k and λ_k as shown in [Figure 4](#). We see that the estimate of the lower bound on the smallest eigenvalue is smaller for SdLBFGS-VR compared to VARCHEN. We also notice that the estimate of the upper bound of the largest eigenvalue of H_k takes even more extreme values for SdLBFGS-VR compared to VARCHEN. The extreme values λ_k and Λ_k reflect an ill-conditioning problem encountered when using SdLBFGS-VR and we believe that it explains the extreme oscillations in the performance of SdLBFGS-VR.

³FastResNet Hyperparameters tuning with Ax on CIFAR10

⁴<https://myrtle.ai/learn/how-to-train-your-resnet-4-architecture/>

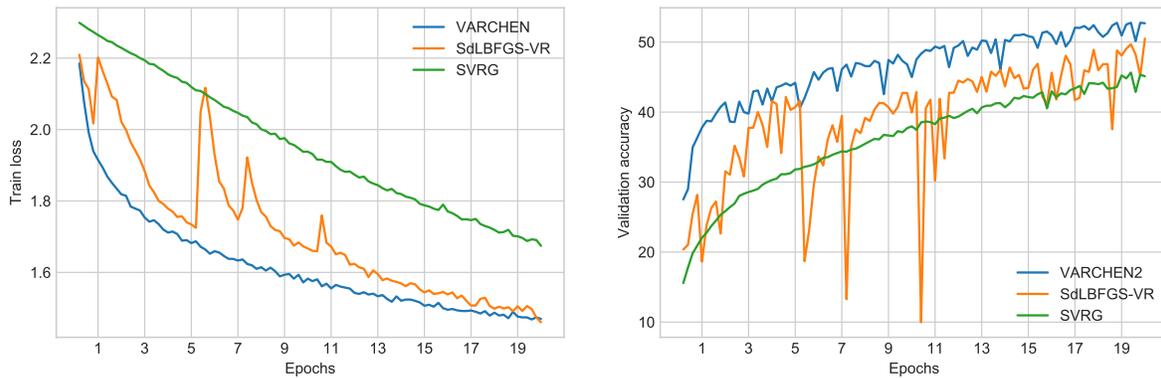


Figure 3: Evolution of the training loss (left) and the validation accuracy (right) for training a modified DavidNet on CIFAR-10.

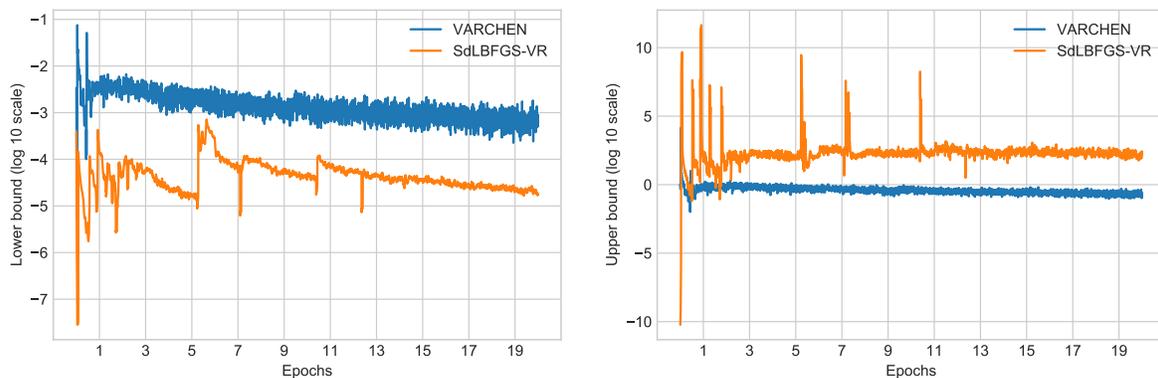


Figure 4: Evolution of the lower bound on the smallest eigenvalue λ_k (left) and the upper bound on the largest eigenvalue Λ_k (right) on a base 10 logarithmic scale for training a modified DavidNet on CIFAR-10.

4 Conclusions and future work

In this paper, we adapt deterministic first- and second-order methods for large-scale optimization to the stochastic context of machine learning, while preserving convergence guarantees and iteration complexity.

For first-order methods, we propose an adaptation of ARIG to the stochastic setting that maintains the same theoretical guarantees. However, this new version is still not practical for solving machine learning problems since it requires evaluating the full gradient to test the accuracy of the gradient approximation. Additionally, objective approximations are required to be increasingly accurate as convergence occurs. Therefore, we propose to leverage adaptive sampling to satisfy the accuracy condition on the gradient approximation. Lotfi [38] shows that the resulting algorithm compares favorably with SGD while maintaining batch size reasonable. Finally, we propose ARAS, which addresses both the practicality and the overall convergence concerns. The computational experiments show that ARAS outperforms SGD and enjoys a conceivable overall convergence property. Nevertheless, its convergence rate and iteration complexity remain an open question.

Concerning second-order methods, we use the stochastic damped L-BFGS algorithm in a nonconvex setting, where there are no guarantees that H_k remains well-conditioned and numerically nonsingular throughout. We introduce a new stochastic damped L-BFGS algorithm that monitors the quality of H_k during the optimization by maintaining bounds on its largest and smallest eigenvalues. Our work is the first to address the Hessian singularity problem by approximating and leveraging such bounds.

Moreover, we propose a new initial inverse Hessian approximation that results in a smoother, less oscillatory training loss and validation accuracy evolution. Additionally, we use variance reduction in order to improve the quality of the curvature approximation and accelerate convergence. Our algorithm converges almost-surely to a stationary point and numerical experiments show that it is more robust to ill-conditioning and more suitable to the highly nonconvex context of deep learning than SdLBFGS-VR. We consider this work to be a first step towards the use of bound estimates to control the quality of the Hessian approximation in approximate second-order algorithms. Future work should aim to improve the quality of these bounds and explore another form of variance reduction that consists of adaptive sampling [5, 7, 30].

The first- and second-order algorithms that we propose in this paper fall under the umbrella of adaptive methods that are crucial today to reduce the alarming engineering energy and computational costs of hyperparameter tuning of deep neural networks. Such algorithms will not only have a profound impact on the efficiency of machine learning methods, but they will also allow entities with limited computational resources to benefit from these methods.

Appendix A Proofs

Lemma 1. By definition of L , Taylor’s theorem yields $|f(x + s) - T(x, s)| \leq \frac{1}{2}L\|s\|^2$ for all x and s . We deduce from (6) that $|T(x_k, s_k) - \bar{T}_k(s_k)| \leq \omega_g^k \|\nabla f(x_k, \omega_g^k)\| \|s_k\|$ for all k . Together, those two observations, the triangle inequality, the definition of s_k in Step 3 of Algorithm 1 and the bound $\omega_g^k \leq \sigma_k^{-1}$ yield

$$|\rho_k - 1| = \frac{|f(x_k + s_k) - \bar{T}_k(s_k)|}{\sigma_k^{-1} \|\nabla f(x_k, \omega_g^k)\|^2} \leq \frac{\frac{1}{2}L\sigma_k^{-2} + \omega_g^k \sigma_k^{-1}}{\sigma_k^{-1}} \leq (\frac{1}{2}L + 1)\sigma_k^{-1}.$$

Thus, $\sigma_k \geq (\frac{1}{2}L + 1)/(1 - \eta_2) \implies \rho_k \geq \eta_2$, which means that iteration k is very successful and $\sigma_{k+1} \leq \sigma_k$. In view of (11), σ_k can never exceed the previous threshold by a factor larger than γ_3 , except if σ_0 were larger than the latter bound. Overall, we obtain (12). \square

Lemma 2. Suppose two values of the error threshold ω_f^k and $\hat{\omega}_f^k$, such that,

$$|f(x_k, \omega_f^k) - f(x_k)| \leq \omega_f^k, \quad (37)$$

$$|f(x_k + s_k, \hat{\omega}_f^k) - f(x_k + s_k)| \leq \hat{\omega}_f^k. \quad (38)$$

We also suppose a constant $\eta_0 < \frac{1}{2}\eta_1$ and we require that

$$\max(\omega_f^k, \hat{\omega}_f^k) \leq \eta_0 [\bar{T}_k(0) - \bar{T}_k(s_k)]. \quad (39)$$

Using (37) and (39), we deduce that

$$f(x_k, \omega_f^k) - f(x_k) \leq |f(x_k, \omega_f^k) - f(x_k)| \leq \omega_f^k \leq \eta_0 [\bar{T}_k(0) - \bar{T}_k(s_k)]. \quad (40)$$

Using (38) and (39), we have that

$$f(x_k + s_k, \hat{\omega}_f^k) - f(x_k + s_k) \leq |f(x_k + s_k, \hat{\omega}_f^k) - f(x_k + s_k)|, \quad (41)$$

$$\leq \hat{\omega}_f^k \leq \eta_0 [\bar{T}_k(0) - \bar{T}_k(s_k)]. \quad (42)$$

From (40) and (41), we have that

$$\beta_k := \frac{\left[f(x_k, \omega_f^k) - f(x_k) \right] - \left[f(x_k + s_k, \hat{\omega}_f^k) - f(x_k + s_k) \right]}{\bar{T}_k(0) - \bar{T}_k(s_k)} \leq \frac{\omega_f^k + \hat{\omega}_f^k}{\bar{T}_k(0) - \bar{T}_k(s_k)} \leq 2\eta_0. \quad (43)$$

Notice that in this case, we have

$$\bar{\rho}_k := \frac{f(x_k, \omega_f^k) - f(x_k + s_k, \hat{\omega}_f^k)}{\bar{T}_k(0) - \bar{T}_k(s_k)} = \frac{f(x_k) - f(x_k + s_k)}{\bar{T}_k(0) - \bar{T}_k(s_k)} + \beta_k = \rho_k + \beta_k. \quad (44)$$

Combining (43) and (44), we obtain the following implication:

$$\bar{\rho}_k \geq \eta_1 \implies \rho_k \geq \eta_1 - 2\eta_0 \geq \bar{\eta}_1 > 0, \quad (45)$$

where $\bar{\eta}_1$ is by definition equal to $\eta_1 - 2\eta_0$. □

Lemma 3. First notice that

$$\gamma \|s\|^2 \leq s^\top y \leq \|s\| \|y\|, \quad \text{and thus} \quad \|s\| \leq \frac{1}{\gamma} \|y\|.$$

Therefore,

$$\frac{1}{\|s\| \|y\|} \leq \rho \leq \frac{1}{\gamma} \frac{1}{\|s\|^2}. \quad (46)$$

Since A is a real symmetric matrix, the spectral theorem states that its eigenvalues are real and it can be diagonalized by an orthogonal matrix. That means that we can find n orthogonal eigenvectors and n eigenvalues counted with multiplicity.

Consider first the special case where s and y are collinear, i.e., there exists $\theta > 0$ such that $y = \theta s$. Any vector such that $u \in s^\perp$, where $s^\perp = \{x \in \mathbb{R}^n : x^\top s = 0\}$, is an eigenvector of A associated with the eigenvalue μ of multiplicity $n - 1$. Moreover, $s^\top y = \theta \|s\|^2 = \|s\| \|y\|$, $\rho = 1/(\theta \|s\|^2)$ and we have

$$As = \left[\mu(I - \rho\theta ss^\top) + \rho ss^\top \right] s = \left[\mu(1 - \rho\theta \|s\|^2) + \rho \|s\|^2 \right] s = \rho \|s\|^2 s.$$

Let us call $\lambda = \rho \|s\|^2$, the eigenvalue associated with eigenvector s . From (46) and $\|y\| \leq L_y \|s\|$, we deduce that

$$\frac{1}{L_y} \leq \lambda \leq \frac{1}{\gamma}.$$

Suppose now that s and y are linearly independent. Any u such that $u^\top s = 0 = u^\top y$ satisfies $Au = \mu u$. This provides us with a $(n - 2)$ -dimensional eigenspace S , associated to the eigenvalue μ of multiplicity $n - 2$. Note that

$$\begin{aligned} As &= \rho \|s\|^2 (1 + \mu\rho \|y\|^2) s - \rho \|s\|^2 \mu y, \\ Ay &= s. \end{aligned}$$

Thus neither s nor y is an eigenvector of A . Now consider an eigenvalue λ associated with an eigenvector u , such that $u \in S^\perp$. Since s and y are linearly-independent, we can search for u of the form $u = s + \beta y$ with $\beta > 0$. The condition $Au = \lambda u$ yields

$$\begin{aligned} \rho \|s\|^2 (1 + \mu\rho \|y\|^2) + \beta &= \lambda, \\ -\rho \|s\|^2 \mu &= \lambda\beta. \end{aligned}$$

We eliminate $\beta = \lambda - \rho \|s\|^2 (1 + \mu\rho \|y\|^2)$ and obtain

$$p(\lambda) = 0,$$

where

$$p(\lambda) = \lambda^2 - \lambda\rho \|s\|^2 (1 + \mu\rho \|y\|^2) + \rho \|s\|^2 \mu.$$

The roots of p must be the two remaining eigenvalues $\lambda_1 \leq \lambda_2$ that we are looking for. In order to establish the lower bound, we need a lower bound on λ_1 whereas to establish the upper bound, we need an upper bound on λ_2 .

On the one hand, let l be the tangent to the graph of p at $\lambda = 0$, defined by

$$l(\lambda) = p(0) + p'(0)\lambda = \mu\rho\|s\|^2 - \lambda\rho\|s\|^2(1 + \mu\rho\|y\|^2).$$

Its unique root is

$$\bar{\lambda} = \frac{\mu}{1 + \mu\rho\|y\|^2}.$$

From (46) and since $\|y\| \leq L_y\|s\|$, we deduce that

$$\bar{\lambda} \geq \frac{\mu}{1 + \frac{\mu}{\gamma} \frac{\|y\|^2}{\|s\|^2}} \geq \frac{\mu}{1 + \frac{\mu}{\gamma} L_y^2}.$$

Since p is convex, it remains above its tangent, and $\bar{\lambda} \leq \lambda_1$.

Finally,

$$\lambda_{\min}(A) \geq \min\left(\frac{1}{L_y}, \frac{\mu}{1 + \frac{\mu}{\gamma} L_y^2}\right) > 0.$$

This establishes the lower bound.

On the other hand, the discriminant $\Delta = \rho^2\|s\|^4(1 + \mu\rho\|y\|^2)^2 - 4\rho\|s\|^2\mu$ must be nonnegative since A is real symmetric, and its eigenvalues are real. We have

$$\lambda_2 = \frac{\rho\|s\|^2(1 + \mu\rho\|y\|^2) + \sqrt{\rho^2\|s\|^4(1 + \mu\rho\|y\|^2)^2 - 4\rho\|s\|^2\mu}}{2}.$$

For any positive a and b such that $a^2 - b > 0$, we have $\sqrt{a^2 - b} \leq a - \frac{b}{2a}$. Thus,

$$\lambda_2 \leq \rho\|s\|^2(1 + \mu\rho\|y\|^2) - \frac{\mu}{1 + \mu\rho\|y\|^2}.$$

From (46), we deduce that

$$\lambda_2 \leq \frac{1}{\gamma} + \frac{\mu}{\gamma^2} \frac{\|y\|^2}{\|s\|^2} - \frac{\mu}{1 + \frac{\mu}{\gamma} \frac{\|y\|^2}{\|s\|^2}}.$$

And since $\|y\| \leq L_y\|s\|$, it follows

$$\lambda_2 \leq \frac{1}{\gamma} + \frac{\mu}{\gamma^2} L_y^2 - \frac{\mu}{1 + \frac{\mu}{\gamma} L_y^2}.$$

Finally,

$$\lambda_{\max}(A) \leq \max\left(\frac{1}{\gamma}, \frac{1}{\gamma} + \frac{\mu}{\gamma^2} L_y^2 - \frac{\mu}{1 + \frac{\mu}{\gamma} L_y^2}\right),$$

which establishes the upper bound. \square

Theorem 2. Consider one damped BFGS update using s and y defined in (30) and \hat{y}_k defined in (32), i.e, $p = 1$,

$$H_{k+1} = \hat{V}_k H_{k+1}^0 \hat{V}_k^\top + \hat{\rho}_k s_k s_k^\top, \quad \text{where} \quad \hat{\rho}_k = 1/s_k^\top \hat{y}_k, \quad \hat{V}_k = I - \hat{\rho}_k s_k \hat{y}_k^\top.$$

Let $0 < \mu_1 := \lambda_{\min}(H_{k+1}^0) \leq \mu_2 := \lambda_{\max}(H_{k+1}^0)$. We have

$$\lambda_{\min}(\mu_1 \hat{V}_k \hat{V}_k^\top + \hat{\rho}_k s_k s_k^\top) \leq \lambda_{\min}(H_{k+1}) \leq \lambda_{\max}(H_{k+1}) \leq \lambda_{\max}(\mu_2 \hat{V}_k \hat{V}_k^\top + \hat{\rho}_k s_k s_k^\top). \quad (47)$$

Let us show that we can apply [Lemma 3](#) to

$$A_1 := \mu_1 \hat{V}_k \hat{V}_k^\top + \hat{\rho}_k s_k s_k^\top \quad \text{and} \quad A_2 := \mu_2 \hat{V}_k \hat{V}_k^\top + \hat{\rho}_k s_k s_k^\top.$$

From [\(34\)](#), we obtain

$$s_k^\top \hat{y}_k \geq \eta \lambda_{\min}(B_{k+1}^0) \|s_k\|^2 = \frac{\eta}{\lambda_{\max}(H_{k+1}^0)} \|s_k\|^2 = \frac{\eta}{\mu_2} \|s_k\|^2.$$

[Assumption 3](#) yields

$$\|\hat{y}_k\| = \|\theta_k y_k + (1 - \theta_k) B_{k+1}^0 s_k\| \leq \|y_k\| + \|B_{k+1}^0 s_k\| \leq (L_g + 1/\mu_1) \|s_k\|.$$

Therefore, we can first apply [Lemma 3](#) with s_k , \hat{y}_k , $\gamma = \eta/\mu_2 > 0$, $L_y = L_g + 1/\mu_1 > 0$ and $\mu = \mu_1 > 0$ for A_1 , and apply it again with $\mu = \mu_2 > 0$ for A_2 . Let $L_1 := L_g + 1/\mu_1$. [Lemma 3](#) and [\(47\)](#) yield

$$\begin{aligned} \lambda_{\min}(H_{k+1}) &\geq \min\left(\frac{1}{L_1}, \frac{\mu_1}{1 + \frac{\mu_1 \mu_2}{\eta} L_1^2}\right) > 0, \\ \lambda_{\max}(H_{k+1}) &\leq \frac{\mu_2}{\eta} + \max\left(0, \frac{\mu_2^3}{\eta^2} L_1^2 - \frac{\mu_2}{1 + \frac{\mu_2^2}{\eta} L_1^2}\right). \end{aligned}$$

Now, consider the case where $p > 1$ and let

$$H_{k+1}^{(h+1)} := \hat{V}_{k-h} H_{k+1}^{(h)} \hat{V}_{k-h}^\top + \hat{\rho}_{k-h} s_{k-h} s_{k-h}^\top, \quad 0 \leq h \leq p-1,$$

where

$$H_{k+1}^{(p)} := H_{k+1}, \quad \hat{\rho}_{k-h} = 1/s_{k-h}^\top \hat{y}_{k-h}, \quad \hat{V}_k = I - \hat{\rho}_{k-h} s_{k-h} \hat{y}_{k-h}^\top.$$

Similarly to the case $p = 1$, we may write

$$\begin{aligned} \lambda_{\min}(H_{k+1}^{(h+1)}) &\geq \lambda_{\min}(\mu_1^{(h)} \hat{V}_{k-h} \hat{V}_{k-h}^\top + \hat{\rho}_{k-h} s_{k-h} s_{k-h}^\top), & \mu_1^{(h)} &:= \lambda_{\min}(H_{k+1}^{(h)}), \\ \lambda_{\max}(H_{k+1}^{(h+1)}) &\leq \lambda_{\max}(\mu_2^{(h)} \hat{V}_{k-h} \hat{V}_{k-h}^\top + \hat{\rho}_{k-h} s_{k-h} s_{k-h}^\top), & \mu_2^{(h)} &:= \lambda_{\max}(H_{k+1}^{(h)}). \end{aligned}$$

Assume by recurrence that $0 < \mu_1^{(h)} \leq \mu_2^{(h)}$. We show that we can apply [Lemma 3](#) to

$$A_1^{(h)} := \mu_1^{(h)} \hat{V}_{k-h} \hat{V}_{k-h}^\top + \hat{\rho}_{k-h} s_{k-h} s_{k-h}^\top \quad \text{and} \quad A_2^{(h)} := \mu_2^{(h)} \hat{V}_{k-h} \hat{V}_{k-h}^\top + \hat{\rho}_{k-h} s_{k-h} s_{k-h}^\top.$$

From [\(34\)](#), we have

$$s_{k-h}^\top \hat{y}_{k-h} \geq \eta \lambda_{\min}(B_{k-h+1}^0) \|s_{k-h}\|^2 = \frac{\eta}{\lambda_{\max}(H_{k-h+1}^0)} \|s_{k-h}\|^2.$$

Using [Assumption 3](#),

$$\|\hat{y}_{k-h}\| = \|\theta_{k-h} y_{k-h} + (1 - \theta_{k-h}) B_{k-h+1}^0 s_{k-h}\| \leq L_g \|s_{k-h}\| + \|B_{k-h+1}^0 s_{k-h}\|,$$

so that

$$\|\hat{y}_{k-h}\| \leq (L_g + \frac{1}{\lambda_{\min}(H_{k-h+1}^0)}) \|s_{k-h}\|.$$

We first apply [Lemma 3](#) with $s = s_{k-h}$, $y = \hat{y}_{k-h}$, $\gamma = \eta/\lambda_{\max}(H_{k-h+1}^0) > 0$, $L_y = L_g + 1/\lambda_{\min}(H_{k-h+1}^0) > 0$ and $\mu = \mu_1^{(h)} > 0$ for $A_1^{(h)}$, and apply it a second time with $\mu = \mu_2^{(h)} > 0$ for $A_2^{(h)}$. Let $L_{k-h+1} := L_g + 1/\lambda_{\min}(H_{k-h+1}^0)$ and $\gamma_{k-h+1} = \eta/\lambda_{\max}(H_{k-h+1}^0)$. Then we have

$$\lambda_{\min}(H_{k+1}^{(h+1)}) \geq \min \left(\frac{1}{L_{k-h+1}}, \frac{\lambda_{\min}(H_{k+1}^{(h)})}{1 + \frac{\lambda_{\min}(H_{k+1}^{(h)})}{\gamma_{k-h+1}} L_{k-h+1}^2} \right), \quad (48)$$

and

$$\lambda_{\max}(H_{k+1}^{(h+1)}) \leq \frac{1}{\gamma_{k-h+1}} + \max \left(0, \frac{\lambda_{\max}(H_{k+1}^{(h)})}{\gamma_{k-h+1}^2} L_{k-h+1}^2 - \frac{\lambda_{\max}(H_{k+1}^{(h)})}{1 + \frac{\lambda_{\max}(H_{k+1}^{(h)})}{\gamma_{k-h+1}} L_{k-h+1}^2} \right). \quad (49)$$

It is clear that we can obtain the lower bound on $\lambda_{\min}(H_{k+1})$ recursively using (48). Obtaining the upper bound on $\lambda_{\max}(H_{k+1})$ using (49) is trickier. However, we notice that inequality (49) implies

$$\lambda_{\max}(H_{k+1}^{(h+1)}) \leq \frac{1}{\gamma_{k-h+1}} + \max \left(0, \frac{\lambda_{\max}(H_{k+1}^{(h)})}{\gamma_{k-h+1}^2} L_{k-h+1}^2 - \frac{\lambda_{\min}(H_{k+1}^{(h)})}{1 + \frac{\lambda_{\max}(H_{k+1}^{(h)})}{\gamma_{k-h+1}} L_{k-h+1}^2} \right). \quad (50)$$

This upper bound is less tight but it allows us to bound $\lambda_{\max}(H_{k+1})$ recursively. \square

Appendix B Experimental details

For second-order methods:

- We apply our Hessian norm control using the bound on the maximum and the minimum eigenvalues of H_{k+1} , where the latter is equivalent to controlling $\|B_{k+1}\|$;
- In the definition of H_{k+1}^0 in (35), we choose $\underline{\gamma}_{k+1}$ and $\bar{\gamma}_{k+1}$ constant;
- The numerical values for all algorithms are the ones that yielded the best results among all sets of values that we experimented with.

Numerical values:

- For all algorithms: we train the network for 20 epochs and use a batch size of 256 samples;
- For SVRG, we choose a step size equal to 0.001;
- For both SdLBFGS-VR and [Algorithm 4](#), the memory parameter $p = 10$, the minimal scaling parameter $\underline{\gamma}_{k+1} = 0.1$ for all k , the constant step size $\alpha_k = 0.1$ and $\eta = 0.25$;
- For [Algorithm 4](#), we use a maximal scaling parameter $\bar{\gamma}_{k+1} = 10^5$ for all k , a lower bound limit $\lambda_{\min} = 10^{-5}$ and an upper bound limit $\lambda_{\max} = 10^5$.

References

- [1] S.-I. Amari. Natural gradient works efficiently in learning. *Neural computation*, 10(2):251–276, 1998.
- [2] H. Asi and J. C. Duchi. The importance of better models in stochastic optimization. *Proceedings of the National Academy of Sciences*, 116(46):22924–22930, 2019.
- [3] E. G. Birgin, J. Gardenghi, J. M. Martínez, S. A. Santos, and Ph. L. Toint. Worst-case evaluation complexity for unconstrained nonlinear optimization using high-order regularized models. *Mathematical Programming*, 163(1–2):359–368, 2017.
- [4] J. Blanchet, C. Cartis, M. Menickelly, and K. Scheinberg. Convergence rate analysis of a stochastic trust-region method via supermartingales. *INFORMS journal on optimization*, 1(2):92–119, 2019.

- [5] R. Bollapragada and S. M. Wild. Adaptive sampling quasi-Newton methods for derivative-free stochastic optimization. arXiv preprint arXiv:1910.13516, 2019.
- [6] R. Bollapragada, R. Byrd, and J. Nocedal. Adaptive sampling strategies for stochastic optimization. *SIAM Journal on Optimization*, 28(4):3312–3343, 2018.
- [7] R. Bollapragada, D. Mudigere, J. Nocedal, H.-J. M. Shi, and P. T. P. Tang. A progressive batching l-bfgs method for machine learning. arXiv preprint arXiv:1802.05374, 2018.
- [8] A. Bordes, L. Bottou, and P. Gallinari. SGD-QN: Careful quasi-Newton stochastic gradient descent. *Journal of Machine Learning Research*, 10(Jul):1737–1754, 2009.
- [9] L. Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT’2010*, pages 177–186. Springer, 2010.
- [10] L. Bottou, F. E. Curtis, and J. Nocedal. Optimization methods for large-scale machine learning. *Siam Review*, 60(2):223–311, 2018.
- [11] R. H. Byrd, G. M. Chin, J. Nocedal, and Y. Wu. Sample size selection in optimization methods for machine learning. *Mathematical programming*, 134(1):127–155, 2012.
- [12] R. H. Byrd, S. L. Hansen, J. Nocedal, and Y. Singer. A stochastic quasi-Newton method for large-scale optimization. *SIAM Journal on Optimization*, 26(2):1008–1031, 2016.
- [13] R. G. Carter. On the global convergence of trust region algorithms using inexact gradient information. *SIAM Journal on Numerical Analysis*, 28(1):251–265, 1991. DOI: [10.1137/0728014](https://doi.org/10.1137/0728014).
- [14] J. Chee and P. Toulis. Convergence diagnostics for stochastic gradient descent with constant learning rate. In *International Conference on Artificial Intelligence and Statistics*, pages 1476–1485, 2018.
- [15] H. Chen, H.-C. Wu, S.-C. Chan, and W.-H. Lam. A stochastic quasi-Newton method for large-scale nonconvex optimization with applications. *IEEE Transactions on Neural Networks and Learning Systems*, 2019.
- [16] R. Chen, M. Menickelly, and K. Scheinberg. Stochastic optimization using a trust-region method and random models. *Mathematical Programming*, 169(2):447–487, 2018.
- [17] J. Collins, J. Sohl-Dickstein, and D. Sussillo. Capacity and trainability in recurrent neural networks. arXiv preprint arXiv:1611.09913, 2016.
- [18] A. R. Conn, N. I. M. Gould, and Ph. L. Toint. *Trust-Region Methods*. Number 1 in MOS-SIAM Series on Optimization. SIAM, Philadelphia, USA, 2000. DOI: [10.1137/1.9780898719857](https://doi.org/10.1137/1.9780898719857).
- [19] F. E. Curtis, K. Scheinberg, and R. Shi. A stochastic trust region algorithm based on careful step normalization. *Informatics Journal on Optimization*, 1(3):200–220, 2019.
- [20] R. S. Dembo, S. C. Eisenstat, and T. Steihaug. Inexact Newton methods. *SIAM Journal on Numerical analysis*, 19(2):400–408, 1982.
- [21] J. E. Dennis Jr and J. J. Moré. A characterization of superlinear convergence and its application to quasi-Newton methods. *Mathematics of computation*, 28(126):549–560, 1974.
- [22] J. E. Dennis Jr and R. B. Schnabel. *Numerical methods for unconstrained optimization and nonlinear equations*, volume 16. SIAM, Philadelphia, PA, 1996.
- [23] J. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7), 2011.
- [24] C. Fang, C. J. Li, Z. Lin, and T. Zhang. Spider: Near-optimal non-convex optimization via stochastic path-integrated differential estimator. In *Advances in Neural Information Processing Systems*, pages 689–699, 2018.
- [25] R. Fletcher. A new approach to variable metric algorithms. *The computer journal*, 13(3):317–322, 1970.
- [26] J. Freund and R. Walpole. *Mathematical statistics*. prentice-hall. Inc., New Jersey, USA, 1971.
- [27] M. P. Friedlander and M. Schmidt. Hybrid deterministic-stochastic methods for data fitting. *SIAM Journal on Scientific Computing*, 34(3):A1380–A1405, 2012.
- [28] R. Gower, D. Goldfarb, and P. Richtárik. Stochastic block BFGS: Squeezing more curvature out of data. In *International Conference on Machine Learning*, pages 1869–1878, 2016.
- [29] F. S. Hashemi, S. Ghosh, and R. Pasupathy. On adaptive sampling rules for stochastic recursions. In *Proceedings of the Winter Simulation Conference 2014*, pages 3959–3970. IEEE, 2014.
- [30] A. Jalilzadeh, A. Nedić, U. V. Shanbhag, and F. Yousefian. A variable sample-size stochastic quasi-Newton method for smooth and nonsmooth stochastic convex optimization. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 4097–4102. IEEE, 2018.

- [31] R. Johnson and T. Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in neural information processing systems*, pages 315–323, 2013.
- [32] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- [33] A. Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- [34] J. Larson and S. C. Billups. Stochastic derivative-free optimization using a trust region framework. *Computational Optimization and applications*, 64(3):619–645, 2016.
- [35] Y. LeCun, C. Cortes, and C. Burges. MNIST handwritten digit database. ATT Labs [Online]. Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010.
- [36] D. D. Lewis, Y. Yang, T. G. Rose, and F. Li. Rcv1: A new benchmark collection for text categorization research. *Journal of machine learning research*, 5(Apr):361–397, 2004.
- [37] D. C. Liu and J. Nocedal. On the limited memory BFGS method for large scale optimization. *Mathematical programming*, 45(1–3):503–528, 1989.
- [38] S. Lotfi. *Stochastic First and Second Order Optimization Methods for Machine Learning*. PhD thesis, Polytechnique Montréal, 2020.
- [39] S. Lotfi, T. B. de Ruisselet, D. Orban, and A. Lodi. Stochastic damped l-bfgs with controlled norm of the hessian approximation. *arXiv preprint arXiv:2012.05783*, 2020.
- [40] A. Mokhtari and A. Ribeiro. Res: Regularized stochastic BFGS algorithm. *IEEE Transactions on Signal Processing*, 62(23):6089–6104, 2014.
- [41] P. Moritz, R. Nishihara, and M. Jordan. A linearly-convergent stochastic L-BFGS algorithm. In *Artificial Intelligence and Statistics*, pages 249–258, 2016.
- [42] N. Murata. A statistical study of on-line learning. *Online Learning and Neural Networks*. Cambridge University Press, Cambridge, UK, pages 63–92, 1998.
- [43] Y. E. Nesterov. A method for solving the convex programming problem with convergence rate $O(1/k^2)$. In *Dokl. akad. nauk SSSR*, volume 269, pages 543–547, 1983.
- [44] L. M. Nguyen, J. Liu, K. Scheinberg, and M. Takáč. Sarah: A novel method for machine learning problems using stochastic recursive gradient. *arXiv preprint arXiv:1703.00102*, 2017.
- [45] J. Nocedal. Updating quasi-Newton matrices with limited storage. *Mathematics of computation*, 35(151):773–782, 1980.
- [46] G. C. Pflug. On the determination of the step size in stochastic quasigradient methods. 1983.
- [47] G. C. Pflug. Non-asymptotic confidence bounds for stochastic approximation algorithms with constant step size. *Monatshefte für Mathematik*, 110(3–4):297–314, 1990.
- [48] G. C. Pflug. Gradient estimates for the performance of markov chains and discrete event processes. *Annals of Operations Research*, 39(1):173–194, 1992.
- [49] B. T. Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5):1–17, 1964.
- [50] M. J. Powell. Algorithms for nonlinear constraints that use Lagrangian functions. *Mathematical programming*, 14(1):224–248, 1978.
- [51] E. Real, A. Aggarwal, Y. Huang, and Q. V. Le. Regularized evolution for image classifier architecture search. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 4780–4789, 2019.
- [52] H. Robbins and S. Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.
- [53] N. N. Schraudolph, J. Yu, and S. Günter. A stochastic quasi-Newton method for online convex optimization. In *Artificial intelligence and statistics*, pages 436–443, 2007.
- [54] T. Tieleman and G. Hinton. Rmsprop: Divide the gradient by a running average of its recent magnitude. *coursea: Neural networks for machine learning*. COURSERA Neural Networks Mach. Learn, 2012.
- [55] X. Wang, S. Ma, D. Goldfarb, and W. Liu. Stochastic quasi-Newton methods for nonconvex stochastic optimization. *SIAM Journal on Optimization*, 27(2):927–956, 2017.
- [56] Z. Wang, K. Ji, Y. Zhou, Y. Liang, and V. Tarokh. Spiderboost and momentum: Faster variance reduction algorithms. In *Advances in Neural Information Processing Systems*, pages 2406–2416, 2019.
- [57] P. Wolfe. Convergence conditions for ascent methods. *SIAM Review*, 11(2):226–235, 1969.

-
- [58] P. Wolfe. Convergence conditions for ascent methods II: Some corrections. *SIAM Review*, 13(2):185–188, 1971.
 - [59] F. Yousefian, A. Nedić, and U. V. Shanbhag. A smoothing stochastic quasi-Newton method for non-Lipschitzian stochastic optimization problems. In *2017 Winter Simulation Conference (WSC)*, pages 2291–2302. IEEE, 2017.
 - [60] B. Zoph and Q. V. Le. Neural architecture search with reinforcement learning. arXiv preprint arXiv:1611.01578, 2016.