

SLiSeS: Subsampled Line Search Spectral Gradient Method for Finite Sums

Stefania Bellavia ^{*†} Nataša Krejić [‡] Nataša Krklec Jerinkić [§] Marcos Raydan [¶]

October 8, 2024

Dedicated to the memory of Andreas Griewank

Abstract

The spectral gradient method is known to be a powerful low-cost tool for solving large-scale optimization problems. In this paper, our goal is to exploit its advantages in the stochastic optimization framework, especially in the case of mini-batch subsampling that is often used in big data settings. To allow the spectral coefficient to properly explore the underlying approximate Hessian spectrum, we keep the same subsample for a prefixed number of iterations before subsampling again. We analyze the required algorithmic features and the conditions for almost sure convergence, and present initial numerical results that show the advantages of the proposed method.

Key words: finite sum minimization, subsampling, spectral gradient methods, line search

1 Introduction

We are concerned with the minimization of the average of finitely many possibly nonconvex smooth functions

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

where each $f_i : \mathbb{R}^n \rightarrow \mathbb{R}$ for $1 \leq i \leq N$ is bounded from below and f_{low} denotes the objective function's lower bound.

Problem (1) has its origin in large-scale data analysis applications. For instance, a wide variety of problems arising in machine learning involve computing an approximate minimizer of the sum of a loss function over a large number of training examples, where there is a large amount of redundancy between examples. In those cases, it is almost mandatory to employ stochastic iterative methods that

^{*}Department of Industrial Engineering, University of Florence, Viale Morgagni, 40/44, 50134 Florence, Italy, e-mail: stefania.bellavia@unifi.it.

[†]INDAM-GNCS Research group, Roma, Italy.

[‡]Department of Mathematics and Informatics, Faculty of Sciences, University of Novi Sad, Trg Dositeja Obradovića 4, 21000 Novi Sad, Serbia, e-mail: natasak@uns.ac.rs.

[§]Department of Mathematics and Informatics, Faculty of Sciences, University of Novi Sad, Trg Dositeja Obradovića 4, 21000 Novi Sad, Serbia, e-mail: natasa.krklec@dmf.uns.ac.rs.

[¶]Center for Mathematics and Applications (NovaMath), FCT NOVA, 2829-516, Caparica, Portugal, e-mail: m.raydan@fct.unl.pt.

update the prediction model based on a relatively small randomly chosen subset (or sample) of the training data [3, 8].

Let us denote the full sample by $\mathcal{N} = \{1, 2, \dots, N\}$ and, independently of the applied iterative method, let us denote the randomly chosen subsample at iteration k by $\mathcal{N}_k \subseteq \mathcal{N}$ where $|\mathcal{N}_k| = S \ll N$. The function estimator obtained by averaging the functions f_i in \mathcal{N}_k is given by

$$f_{\mathcal{N}_k}(x) = \frac{1}{S} \sum_{i \in \mathcal{N}_k} f_i(x),$$

and the associated gradient estimator is

$$\nabla f_{\mathcal{N}_k}(x) = \frac{1}{S} \sum_{i \in \mathcal{N}_k} \nabla f_i(x).$$

For solving (1), stochastic gradient (SG) algorithms of the form

$$x_{k+1} = x_k - \gamma_k \nabla f_{\mathcal{N}_k}(x_k),$$

exhibit a convenient low computational cost, and as a consequence have become popular and successful for solving large-scale machine learning problems, for which the full sample approach is prohibitively expensive. However, when combined with standard choices of the step length $\gamma_k > 0$ (constant or decreasing), these methods must perform a large number of iterations to observe an adequate reduction in the objective function [4, 8, 10, 29].

Recently, for solving (1), the SG method has been enriched with the use of spectral step lengths to speed up the convergence of the iterative process; see, e.g., [2, 23, 28, 29, 31]. The spectral method (originally introduced by Barzilai and Borwein [1]) and its variants, for the general (full sample) unconstrained minimization problem, are low-cost gradient methods which have proved to be very effective in practice for large-scale optimization. They have received a lot of attention in the last three decades including theoretical understanding, extensions, and adaptations for different scenarios (unconstrained and constrained) and for some specific applications; see, e.g., the review papers [6, 18, 32] and references therein. The recently proposed schemes in [2, 23, 28, 29, 31], which combine the SG approach with spectral step lengths, have a common feature: they all maintain the traditional option of changing the random subsample \mathcal{N}_k (with constant or dynamically increasing size) at each iteration.

From the very beginning it has been recognized, first for quadratics and later for the general unconstrained minimization problem, that the effectiveness of the spectral methods relies on the relationship between the inverse of the step lengths and the eigenvalues of the underlying Hessian of the objective function, near the local minimizer that is being approximated; see [11–14, 16, 17, 19, 21, 25]. The key issue that has been extensively observed in practice is that, in sharp contrast with the classical Cauchy gradient method, the inverse of the step lengths of the spectral methods move all over the spectrum of the mentioned Hessian. As a consequence, these iterative methods exhibit a (highly) non-monotone behavior during the convergence process. That unconventional behavior of the spectral gradient method, and its extensions, has been formally studied and accurately called the *sweeping spectrum behavior*; see [13, 14].

Now, back to solving (1), we note that changing the sample at each iteration, i.e., changing the chosen functions and the gradients, produces a poor approximation of the inverse of the spectral step lengths to the eigenvalues of the underlying Hessian due to the influence of noise. Therefore,

changing the sample at each iteration is counterproductive to the effectiveness of the sweeping spectrum behavior.

In order to promote the speed-up that can be obtained from the sweeping process, our main contribution in this work is to keep the same sample for a prefixed number of iterations before subsampling again. That way, the inverse of the step lengths can be focused for a few iterations on a set of fixed eigenvalues and enable the spectral method to show its power in reducing the value of the objective function. As we will show in our numerical experiments, this change in the algorithm results in a significant improvement in the practical behavior of the method. Nevertheless, this novel strategy brings in new challenges with respect to the convergence analysis.

The rest of this document is organized as follows. In Section 2, we describe in detail the new algorithm combining the new subsampling strategy with an appropriate line search globalization technique aimed at enhancing the practical behaviour of the overall procedure. In Section 3, we establish the convergence properties of the algorithm under some variance assumptions for both uniform and nonuniform sampling. We prove the almost sure convergence assuming some additional properties on the sampling. Moreover, we suggest a modification of the algorithm which employs a combination of predefined and adaptive step sizes and allows to prove the convergence result avoiding both aforementioned additional assumptions.

In Section 4, we report on the obtained numerical results and give further insights into the proposed approach. Finally, in Section 5 we present some concluding remarks.

2 The SLiSeS Framework

At each iteration, our scheme computes a mini-batch stochastic gradient estimate g_k and employs diminishing step sizes based on the Barzilai-Borwein spectral coefficient. A nonmonotone line search is used that employs stochastic function estimates, at the current iterate and at the tentative new iterate, obtained by averaging the objective function in the same sample set used for the gradient estimator.

Our framework algorithm is described in Algorithm 1, while Algorithm 2 details the line search procedure. We will refer to the proposed framework as SLiSeS (Subsampled Line Search Spectral gradient).

Algorithm 1: SLiSeS (Subsampled Line Search Spectral Gradient Method)

- S0 Initialization: $x_0 \in \mathbb{R}^n, \eta \in (0, 1), S \in \{1, 2, \dots, N\}, m \in \mathbb{N}, 0 < \gamma_{min} \leq 1 \leq \gamma_{max} < \infty, \{t_k\} \in \mathbb{R}_+^\infty$ such that $\sum_k t_k \leq \bar{t} < \infty$ and $\mathcal{N}_0 \subseteq \mathcal{N}, |\mathcal{N}_0| = S$. Set $k = 0$.
- S1 Sampling: If $mod(k, m) = 0$ choose $\mathcal{N}_k \subseteq \mathcal{N}$ such that $|\mathcal{N}_k| = S$. Else, set $\mathcal{N}_k = \mathcal{N}_{k-1}$.
- S2 Compute $g_k = \nabla f_{\mathcal{N}_k}(x_k)$.
- S3 Spectral coefficient: If $mod(k, m) = 0$ and $m > 1$ set $c_k = 1/\|g_k\|$. Else, $s_{k-1} = x_k - x_{k-1}, y_{k-1} = g_k - g_{k-1}$ and set $c_k = \|s_{k-1}\|^2 / (s_{k-1}^T y_{k-1})$.
- S4 Set $\gamma_k = \min\{\gamma_{max}, \max\{\gamma_{min}, c_k\}\} / k$.
- S5 Search direction: Set $d_k = -\gamma_k \nabla f_{\mathcal{N}_k}(x_k)$.
- S6 Step size: Find $\alpha_k \in (0, 1]$ such that

$$f_{\mathcal{N}_k}(x_k + \alpha_k d_k) \leq f_{\mathcal{N}_k}(x_k) + \eta \alpha_k d_k^T \nabla f_{\mathcal{N}_k}(x_k) + t_k$$

by employing Algorithm 2.

S7 Set $x_{k+1} = x_k + \alpha_k d_k$, $k = k + 1$ and go to Step S1.

At step S1 we compute the subsample used to form the gradient estimator at Step S2. This subsample is of size S which is assumed to be significantly smaller than N so that the mini-batch approach is employed at each iteration. Step S1 indicates that we keep the same subsample for m iterations. After every m iterations, a new subsample is chosen. One of the main ingredients of the proposed algorithm is the spectral coefficient. At Step S3 of Algorithm 1 we calculate the coefficient c_k . It is based on the spectral coefficient formula at the iterations where the subsample is not changed. On the contrary, after every m iterations - when a new subsample \mathcal{N}_k is introduced, c_k is set to $1/\|g_k\|$, except for the case $m = 1$. In this latter case the subsample is changed at each iteration and we use the Spectral BB step length. This way the method reduces to a stochastic gradient method with BB choice of the steplength similar to one considered in [2].

Furthermore, scaling with the gradient norm is relevant from practical point of view, while it is not crucial in the analysis and some other choices are feasible too, such as $c_k = 1$ for instance. Occasional avoidance of the spectral coefficient is due to the fact that we let the subsample size S to be arbitrary small and thus the noise can be rather large. In that case the spectral coefficient might not offer any useful information on the spectrum of the Hessian as it would be computed on different estimator functions. Hence, we let the c_k go through the spectrum of the approximate Hessian only when the subsample is not changed, hoping that it will represent well its full sample counterpart. However, this may still provide some poor approximations given that we have a mini-batch instead of growing batch procedure. Moreover, we also aim to solve nonconvex problems which can yield negative values of the spectral coefficients as well. Therefore, in step S4, we “tame” the coefficient c_k by projecting it on some positive, arbitrary large interval $[\gamma_{min}, \gamma_{max}]$ and by dividing it by k . Numerical experiments show that this strategy is beneficial since otherwise the algorithm exhibits erratic behavior. After determining the coefficient γ_k , we scale the negative gradient to obtain the search direction and use the line search procedure as described in Algorithm 2.

Some comments on the Algorithm are in order.

- Step S4 still allows enough freedom for the spectral coefficient; in the numerical result section we will show that SLiSeS outperforms stochastic gradient methods that employs predefined sequence of step sizes such as $1/k$.
- The aforementioned noise of the spectral coefficient can be avoided by additional sampling. For instance, one can choose a new sample at each iteration and then calculate $\nabla f_{\mathcal{N}_k}(x_{k-1})$ to obtain $y_{k-1} = \nabla f_{\mathcal{N}_k}(x_k) - \nabla f_{\mathcal{N}_k}(x_{k-1})$ and the corresponding c_k . Although this strategy is appealing, we believe that the approach presented in Algorithm 1 is more beneficial. The first reason is the computational cost which can be significantly enlarged if additional gradient sampling is applied, especially in high-dimensional problems. Second reason comes from the observed behavior of spectral method to collect the second order information by “walking” through the spectrum which we allow by keeping the same sample for a prefixed number of iterations. The resulting algorithm is rather cheap, efficient and supported by the underlying theoretical foundations presented in Section 3.
- We provide convergence analysis for both uniform and nonuniform sampling strategies, thus allowing biased estimators of the functions and the gradients as well. However, we assume that the approximate gradient is evaluated over the same subsample as the approximate function

(Step S2). This allows us to obtain some valuable results concerning the line search procedure as indicated in Lemma 3.1. Although we allow nonmonotone line search in step S6 and thus any search direction is acceptable, S2 is crucial for proving that the step size is bounded away from zero under the assumption of Lipschitz-continuous gradients. This feature is sufficient for the convergence analysis that we perform in the sense that the proofs are not leaning on the sufficient decrease condition coming from step S6. From theoretical point of view, α_k in Step S6 can even be a constant. However, as we will show in the experiments, the line search that we perform is highly relevant in practical behavior.

The Line Search procedure used to compute α_k in Step S6 of the SLiSeS Algorithm is described in Algorithm 2 and it is denoted as LSP (Line Search procedure). At Step S3 of SLP the quadratic interpolation provides a candidate step size $\tilde{\alpha}_j$; if it becomes too small, in Step S4 LSP switches to the backtracking procedure. The steplength computed by LSP remains bounded away from zero under some standard assumptions as it will be shown in the next section. Although some other approaches are feasible, we choose this strategy motivated by our numerical experience. Notice that it does not require further sampling except for the evaluation of the function estimator at new trial points needed to check the Armijo-like condition.

Algorithm 2: LSP (Line Search Procedure)

S0 Input parameters: $x_k \in \mathbb{R}^n, \eta \in (0, 1), \mathcal{N}_k, \gamma_k, t_k$.

S1 Initialization: Set $j = 0, dm_k = g_k^T d_k, \alpha_j = 1$.

S2 If

$$f_{\mathcal{N}_k}(x_k + \alpha_j d_k) \leq f_{\mathcal{N}_k}(x_k) + \eta \alpha_j dm_k + t_k \tag{2}$$

go to Step S5. Else, go to Step S3.

S3 If $\alpha_j > 0.1$, compute

$$\tilde{\alpha}_j = \frac{-dm_k \alpha_j^2}{2(f_{\mathcal{N}_k}(x_k + \alpha_j d_k) - f_{\mathcal{N}_k}(x_k) - \alpha_j dm_k)}.$$

If $\tilde{\alpha}_j < 0.1\alpha_j$ or $\tilde{\alpha}_j > 0.9\alpha_j$, set $\tilde{\alpha}_j = \alpha_j/2$. Set $\alpha_{j+1} = \tilde{\alpha}_j, j = j + 1$ and go to Step S2.

S4 If $\alpha_j \leq 0.1$, set $\alpha_{j+1} = \alpha_j/2, j = j + 1$ and go to Step S2.

S5 Set $\alpha_k = \alpha_j$ and STOP.

3 Convergence analysis

Within this section we analyze conditions under which the spectral gradient method SLiSeS converges a.s. in stochastic mini-batch framework. We start our analysis by proving that LSP yields uniformly bounded step sizes under the assumption of Lipschitz-continuous gradients stated in A1. Then, Subsection 3.1 considers uniform sampling which yields unbiased estimators. Assuming the common bounded variance assumption A2, we prove a.s. convergence result under the Assumption A3 on the sampling properties (Subsection 3.1.1). We also prove that the a.s. convergence result can be achieved by avoiding Assumption A3, provided that the proposed algorithm is modified as outlined in Subsection 3.1.2. Subsection 3.2 is devoted to nonuniform sampling as stated in Assumption A4. Under the similarity assumption stated in A5, we prove the a.s. convergence for the modified

algorithm as in Subsection 3.1.2. Finally, we show that stronger convergence result can be achieved for functions satisfying the Polyak-Lojasiewicz (P-L) condition and for strongly convex functions in Subsection 3.3.

Assumption 1 (A1). *The functions $f_i, i = 1, \dots, N$ are bounded from below and twice continuously-differentiable with L -Lipschitz-continuous gradients.*

Notice that Assumption A1 implies that $\|\nabla^2 f_i(x)\| \leq L$. Moreover, for any \mathcal{N}_k the function $f_{\mathcal{N}_k}$ is bounded from below with L -Lipschitz-continuous gradients and bounded Hessians, i.e., $\|\nabla^2 f_{\mathcal{N}_k}(x)\| \leq L$.

As already mentioned, the line search procedure is well defined. This is shown in the following Lemma.

Lemma 3.1. *Suppose that A1 holds. Then the LSP procedure is well defined and there exists a constant $\alpha_{min} > 0$ such that $\alpha_{min} \leq \alpha_k \leq 1$ for every k .*

Proof. Since the function $f_{\mathcal{N}_k}$ is bounded from below, and the direction d_k is descent, there exists an interval of step sizes which satisfy Armijo condition (2).

Let us now show that the step sizes are bounded from below. The step size is computed either in S3 or S4 of LSP. Assume first that the current value of α_j is such that $\alpha_j > 0.1$ so we enter S3 to compute $\tilde{\alpha}_j$. If $\tilde{\alpha}_j < 0.1\alpha_j$ then we have $\alpha_{j+1} = \alpha_j/2 > 0.05$. In all other cases we have $\alpha_{j+1} > 0.1\alpha_j > 0.01$. Thus for all α_j computed in S3 we have $\alpha_j > 0.01$. Thus we can enter S4 with $\alpha_j < 0.1$ but at the same time $\alpha_j > 0.01$, i.e. the value that is first used in S4 can not be arbitrary close to 0. Observe that once we enter S4 we can not go back to S3, and all subsequent values of the step size are computed by S4. Let us denote by α'_k the last step size in LSP which did not satisfy (2), i.e., $\alpha'_k = 2\alpha_k$. Since $t_k \geq 0$ there holds

$$f_{\mathcal{N}_k}(x_k + \alpha'_k d_k) > f_{\mathcal{N}_k}(x_k) + \eta \alpha'_k d_k^T \nabla f_{\mathcal{N}_k}(x_k).$$

On the other hand, using the Lipschitz continuity we obtain

$$f_{\mathcal{N}_k}(x_k + \alpha'_k d_k) \leq f_{\mathcal{N}_k}(x_k) + \alpha'_k \nabla f_{\mathcal{N}_k}(x_k)^T d_k + \frac{L}{2} \alpha'^2_k \|d_k\|^2.$$

Combining the previous two inequalities and $d_k = -\gamma_k \nabla f_{\mathcal{N}_k}(x_k)$, with $\gamma_k \leq \gamma_{max}$ as specified Step S4 of SLiSeS algorithm, we obtain

$$\alpha'_k \geq \frac{-2(1-\eta) \nabla f_{\mathcal{N}_k}(x_k)^T d_k}{L \|d_k\|^2} = \frac{2(1-\eta) \gamma_k \|\nabla f_{\mathcal{N}_k}(x_k)\|^2}{L \gamma_k^2 \|\nabla f_{\mathcal{N}_k}(x_k)\|^2} \geq \frac{2(1-\eta)}{L \gamma_{max}}.$$

Thus, $\alpha_k \geq (1-\eta)/(L \gamma_{max})$ for α_k computed in S4. Recall that we already concluded that $\alpha_k > 0.01$ if it is computed in S3. Therefore the statement holds with

$$\alpha_{min} = \min\left\{\frac{(1-\eta)}{L \gamma_{max}}, 0.01\right\}.$$

□

Denote by \mathcal{F}_k a σ -algebra generated by $\mathcal{N}_0, \dots, \mathcal{N}_{k-1}$. Notice that x_k is \mathcal{F}_k -measurable, i.e., x_k is fully determined (known) under condition of knowing \mathcal{F}_k . On the other hand, α_k, g_k and γ_k are not \mathcal{F}_k -measurable in general. We will be interested only in \mathcal{F}_k for iterations k such that $\text{mod}(k, m) = 0$. We will refer to these iterations as outer iterations. In the remaining iterations, to be referred to as inner iterations, we keep the same sample.

The following theorem from [26] is used for proving the a.s. convergence result.

Theorem 3.1. Let $U_l, \beta_l, \xi_l, \rho_l \geq 0$ be \mathcal{F}_l -measurable random variables such that

$$E[U_{l+1}|\mathcal{F}_l] \leq (1 + \beta_l)U_l + \xi_l - \rho_l, \quad l = 1, 2, \dots$$

If $\sum_l \beta_l < \infty$ and $\sum_l \xi_l < \infty$, then $U_l \rightarrow U < \infty$ a.s. and $\sum_l \rho_l < \infty$ a.s.

3.1 Uniform sampling

Within this subsection we assume that the sampling in step S1 of SLiSeS method is uniform. This implies that

$$E[f_{\mathcal{N}_k}(x_k)|\mathcal{F}_k] = f(x_k) \quad \text{and} \quad E[g_k|\mathcal{F}_k] = \nabla f(x_k), \quad (3)$$

for all k such that $\text{mod}(k, m) = 0$.

We state the following common assumption for subsampling methods.

Assumption 2 (A2). There exists a constant $G > 0$ such that for all k with $\text{mod}(k, m) = 0$ there holds

$$E[\|g_k - \nabla f(x_k)\|^2|\mathcal{F}_k] \leq G.$$

Assumption A2 together with (3) implies

$$E[\|g_k\|^2|\mathcal{F}_k] \leq G + \|\nabla f(x_k)\|^2, \quad (4)$$

at any iteration k such that $\text{mod}(k, m) = 0$.

3.1.1 Subsampling strategies

Given that \mathcal{N}_k is finite let s denote the number of possible samples $\mathcal{N}_k : \mathcal{N}_k^1, \dots, \mathcal{N}_k^s$. Then, for any \mathcal{N}_k at iteration k such that $\text{mod}(k, m) = 0$, we have $\|\nabla f(x_k)\|^2 = E[\nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k)|\mathcal{F}_k]$ and

$$\|\nabla f(x_k)\|^2 = \frac{1}{s} \left(\sum_{i \in I_k} \nabla f(x_k)^T \nabla f_{\mathcal{N}_k^i}(x_k) + \sum_{j \in J_k} \nabla f(x_k)^T \nabla f_{\mathcal{N}_k^j}(x_k) \right), \quad (5)$$

where $I_k = \{i : \nabla f(x_k)^T \nabla f_{\mathcal{N}_k^i}(x_k) \geq 0\}$ and $J_k = \{j : \nabla f(x_k)^T \nabla f_{\mathcal{N}_k^j}(x_k) < 0\}$. Notice that (5) implies that $I_k \neq \emptyset$ for all k . We state the following assumption.

Assumption 3 (A3). There exist constants $C, \theta > 0$ and $0 < \delta < 1$ such that at any iteration k such that $\text{mod}(k, m) = 0$,

$$P(B_k)E[\nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k)|\mathcal{F}_k, B_k] \leq \frac{\theta}{s} \|\nabla f(x_k)\|^2 + \frac{C}{sk^\delta}, \quad (6)$$

where B_k is the event that $\nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k) \geq 0$ and $P(B_k)$ is the probability of this event.

Notice that (6) can also be obtained by assuming the following

$$E[\|\nabla f_{\mathcal{N}_k}(x_k)\|^2|\mathcal{F}_k] \leq \frac{\tilde{C}}{k^\delta} + \tilde{\theta} \|\nabla f(x_k)\|^2, \quad (7)$$

for some $\tilde{\theta}, \tilde{C}, \tilde{\delta} > 0$ (see [9, Lemma 2]). Thus, the Assumption A3 can be replaced with (7), but we decide to keep the weaker assumption that states (6). The assumption is closely related to the so-called orthogonality test in [7], which is used as the sample increase test. Assuming that A3 holds could be interpreted as an assumption that the subsample is large enough, i.e. in our case that S is large enough. Another example of functions that satisfy (7) with $\tilde{C} = 0$ are the ones that satisfy the following strong growth condition (SGC):

$$\frac{1}{N} \sum_{i=1}^N \|\nabla f_i(x_k)\|^2 \leq \tilde{\theta} \|\nabla f(x_k)\|^2,$$

for a given $\tilde{\theta}$, coupled with uniform sampling. Such functions arise in binary classification with a linear classifier employing the squared hinge loss or logistic regression when the data is linearly separable, as shown in [15, Appendix A].

Theorem 3.2. *Suppose that the Assumptions A1, A2 and A3 hold. Then, provided that $s \geq \theta$, with probability 1 there holds $\liminf_{k \rightarrow \infty} \|\nabla f(x_k)\| = 0$.*

Proof. Let us consider a subsequence $K \subset \mathbb{N}$ such that $k \in K$ if $\text{mod}(k, m) = 0$. For the sake of readability we derive the proof for $m = 3$ but the reasoning is exactly the same for larger values of m . So, given $k \in K$ and x_k , we choose a sample \mathcal{N}_k and compute x_{k+1} , x_{k+2} and x_{k+3} taking the directions $-\nabla f_{\mathcal{N}_k}(x_k)$, $-\nabla f_{\mathcal{N}_k}(x_{k+1})$ and $-\nabla f_{\mathcal{N}_k}(x_{k+2})$ and applying the line search algorithm. Therefore we have

$$x_{k+3} = x_k + \sum_{j=0}^2 \alpha_{k+j} d_{k+j}$$

with $d_{k+j} = -\gamma_{k+j} \nabla f_{\mathcal{N}_k}(x_{k+j})$. Letting $\nabla^2 f_{\mathcal{N}_k}(x)$ be the Hessian matrix of $f_{\mathcal{N}_k}$ at x , the Taylor expansion yields

$$\nabla f_{\mathcal{N}_k}(x_{k+1}) = \nabla f_{\mathcal{N}_k}(x_k) - \alpha_k \gamma_k \nabla^2 f_{\mathcal{N}_k}(\theta_k) \nabla f_{\mathcal{N}_k}(x_k), \quad (8)$$

for some $\theta_k \in B(x_k, \alpha_k \|d_k\|)$, and

$$\begin{aligned} \nabla f_{\mathcal{N}_k}(x_{k+2}) &= \nabla f_{\mathcal{N}_k}(x_{k+1}) - \alpha_{k+1} \gamma_{k+1} \nabla^2 f_{\mathcal{N}_k}(\theta_{k+1}) \nabla f_{\mathcal{N}_k}(x_{k+1}) \\ &= \nabla f_{\mathcal{N}_k}(x_k) - \alpha_k \gamma_k \nabla^2 f_{\mathcal{N}_k}(\theta_k) \nabla f_{\mathcal{N}_k}(x_k) - \\ &\quad \alpha_{k+1} \gamma_{k+1} \nabla^2 f_{\mathcal{N}_k}(\theta_{k+1}) (\nabla f_{\mathcal{N}_k}(x_k) - \alpha_k \gamma_k \nabla^2 f_{\mathcal{N}_k}(\theta_k) \nabla f_{\mathcal{N}_k}(x_k)) \\ &= \nabla f_{\mathcal{N}_k}(x_k) - \alpha_k \gamma_k \nabla^2 f_{\mathcal{N}_k}(\theta_k) \nabla f_{\mathcal{N}_k}(x_k) - \alpha_{k+1} \gamma_{k+1} \nabla^2 f_{\mathcal{N}_k}(\theta_{k+1}) \nabla f_{\mathcal{N}_k}(x_k) + \\ &\quad \alpha_k \gamma_k \alpha_{k+1} \gamma_{k+1} \nabla^2 f_{\mathcal{N}_k}(\theta_{k+1}) \nabla^2 f_{\mathcal{N}_k}(\theta_k) \nabla f_{\mathcal{N}_k}(x_k), \end{aligned} \quad (9)$$

for some $\theta_{k+1} \in B(x_{k+1}, \alpha_{k+1} \|d_{k+1}\|)$. Since $\alpha_k \leq 1$ and $\gamma_k \leq \gamma_{\max}/k \leq \gamma_{\max}$ for all k , the assumption A1 further implies the existence of constants C_1, C_2 such that

$$\|\nabla f_{\mathcal{N}_k}(x_{k+1})\| \leq C_1 \|\nabla f_{\mathcal{N}_k}(x_k)\|, \quad \|\nabla f_{\mathcal{N}_k}(x_{k+2})\| \leq C_2 \|\nabla f_{\mathcal{N}_k}(x_k)\|. \quad (10)$$

Furthermore, we can express $\nabla f(x_k)^T(x_{k+3} - x_k)$ as follows,

$$\begin{aligned}
\nabla f(x_k)^T(x_{k+3} - x_k) &= -\left(\sum_{j=0}^2 \alpha_{k+j}\gamma_{k+j}\right)\nabla f(x_k)^T\nabla f_{\mathcal{N}_k}(x_k) + \\
&\quad \alpha_{k+1}\gamma_{k+1}\alpha_k\gamma_k\nabla f(x_k)^T\nabla^2 f_{\mathcal{N}_k}(\theta_k)\nabla f_{\mathcal{N}_k}(x_k) + \\
&\quad \alpha_k\gamma_k\alpha_{k+2}\gamma_{k+2}\nabla f(x_k)^T\nabla^2 f_{\mathcal{N}_k}(\theta_k)\nabla f_{\mathcal{N}_k}(x_k) + \\
&\quad \alpha_{k+1}\gamma_{k+1}\alpha_{k+2}\gamma_{k+2}\nabla f(x_k)^T\nabla^2 f_{\mathcal{N}_k}(\theta_{k+1})\nabla f_{\mathcal{N}_k}(x_k) - \\
&\quad \alpha_k\gamma_k\alpha_{k+1}\gamma_{k+1}\alpha_{k+2}\gamma_{k+2}\nabla f(x_k)^T\nabla^2 f_{\mathcal{N}_k}(\theta_{k+1})\nabla^2 f_{\mathcal{N}_k}(\theta_k)\nabla f_{\mathcal{N}_k}(x_k). \quad (11)
\end{aligned}$$

Notice that for all $k \in \mathbb{N}$ we have

$$\frac{\alpha_{\min}\gamma_{\min}}{k} \leq \alpha_k\gamma_k \leq \frac{\gamma_{\max}}{k}. \quad (12)$$

Moreover, since the Assumption A1 implies bounded Hessians, we obtain

$$\begin{aligned}
\nabla f(x_k)^T(x_{k+3} - x_k) &\leq -\left(\sum_{j=0}^2 \alpha_{k+j}\gamma_{k+j}\right)\nabla f(x_k)^T\nabla f_{\mathcal{N}_k}(x_k) + \\
&\quad \frac{\gamma_{\max}^2}{k^2}L\|\nabla f(x_k)\|\|\nabla f_{\mathcal{N}_k}(x_k)\|(3 + \frac{\gamma_{\max}}{k}L) \\
&= -\left(\sum_{j=0}^2 \alpha_{k+j}\gamma_{k+j}\right)\nabla f(x_k)^T\nabla f_{\mathcal{N}_k}(x_k) + \frac{\tilde{C}_1}{k^2}\|\nabla f(x_k)\|\|\nabla f_{\mathcal{N}_k}(x_k)\|, \quad (13)
\end{aligned}$$

with $\tilde{C}_1 = \gamma_{\max}^2L(3 + \gamma_{\max}L)$. Furthermore, applying (10) we obtain

$$\begin{aligned}
\|x_{k+3} - x_k\|^2 &\leq 4\sum_{j=0}^2 \|\alpha_{k+j}d_{k+j}\|^2 \leq 4\frac{\gamma_{\max}^2}{k^2}\|\nabla f_{\mathcal{N}_k}(x_k)\|^2(1 + C_1^2 + C_2^2) \\
&= \frac{\tilde{C}_2}{k^2}\|\nabla f_{\mathcal{N}_k}(x_k)\|^2, \quad (14)
\end{aligned}$$

with $\tilde{C}_2 = 4\gamma_{\max}^2(1 + C_1^2 + C_2^2)$. Using Assumption A1 and the descent lemma we get

$$f(x_{k+3}) \leq f(x_k) + \nabla f(x_k)^T(x_{k+3} - x_k) + \frac{L}{2}\|x_{k+3} - x_k\|^2.$$

Subtracting f_{low} and applying the conditional expectation with respect to \mathcal{F}_k we obtain

$$E[f(x_{k+3}) - f_{low}|\mathcal{F}_k] \leq f(x_k) - f_{low} + E[\nabla f(x_k)^T(x_{k+3} - x_k)|\mathcal{F}_k] + \frac{L}{2}E[\|x_{k+3} - x_k\|^2|\mathcal{F}_k]. \quad (15)$$

Further, by (4) and (14) we obtain

$$E[\|x_{k+3} - x_k\|^2|\mathcal{F}_k] \leq \frac{\tilde{C}_2}{k^2}E[\|\nabla f_{\mathcal{N}_k}(x_k)\|^2|\mathcal{F}_k] \leq \frac{\tilde{C}_2}{k^2}(G + \|\nabla f(x_k)\|^2). \quad (16)$$

Now, let us upper bound $E[\nabla f(x_k)^T(x_{k+3} - x_k)|\mathcal{F}_k]$. Starting from (13) and using the fact that $\nabla f(x_k)$ is \mathcal{F}_k -measurable we get

$$\begin{aligned}
E[\nabla f(x_k)^T(x_{k+3} - x_k)|\mathcal{F}_k] &= E[-\sum_{j=0}^2 \alpha_{k+j}\gamma_{k+j} \nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k)|\mathcal{F}_k] + \\
&\quad E[\frac{\tilde{C}_1}{k^2} \|\nabla f(x_k)\| \|\nabla f_{\mathcal{N}_k}(x_k)\| |\mathcal{F}_k] \\
&= E[-\sum_{j=0}^2 \alpha_{k+j}\gamma_{k+j} \nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k)|\mathcal{F}_k] + \\
&\quad \frac{\tilde{C}_1}{k^2} \|\nabla f(x_k)\| E[\|\nabla f_{\mathcal{N}_k}(x_k)\| |\mathcal{F}_k]. \tag{17}
\end{aligned}$$

Let us upper bound the second term first. There holds

$$E[\|\nabla f_{\mathcal{N}_k}(x_k)\| |\mathcal{F}_k] \leq \sqrt{E[\|\nabla f_{\mathcal{N}_k}(x_k)\|^2 |\mathcal{F}_k]} \leq \sqrt{G} + \|\nabla f(x_k)\|. \tag{18}$$

Thus, considering both cases, i.e., $\|\nabla f(x_k)\| \leq 1$ and $\|\nabla f(x_k)\| > 1$, we obtain

$$\frac{\tilde{C}_1}{k^2} \|\nabla f(x_k)\| E[\|\nabla f_{\mathcal{N}_k}(x_k)\| |\mathcal{F}_k] \leq \frac{\tilde{C}_3}{k^2} (\sqrt{G} + \|\nabla f(x_k)\|)^2, \tag{19}$$

where $\tilde{C}_3 := \tilde{C}_1(\sqrt{G} + 1)$. Finally, to upper bound $E[-\sum_{j=0}^2 \alpha_{k+j}\gamma_{k+j} \nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k)|\mathcal{F}_k]$ we proceed as follows. For any realizations such that B_k holds, using (12) we have

$$-\sum_{j=0}^2 \alpha_{k+j}\gamma_{k+j} \nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k) \leq -3 \frac{\gamma_{\min} \alpha_{\min}}{k+2} \nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k), \tag{20}$$

while if B_k is not true

$$-\sum_{j=0}^2 \alpha_{k+j}\gamma_{k+j} \nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k) \leq -3 \frac{\gamma_{\max}}{k} \nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k). \tag{21}$$

Moreover, letting \bar{B}_k the complementary event of B_k it holds:

$$\begin{aligned}
E[-\sum_{j=0}^2 \alpha_{k+j}\gamma_{k+j} \nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k)|\mathcal{F}_k] &= P(B_k) E[-\sum_{j=0}^2 \alpha_{k+j}\gamma_{k+j} \nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k)|\mathcal{F}_k, B_k] + \\
&\quad P(\bar{B}_k) E[-\sum_{j=0}^2 \alpha_{k+j}\gamma_{k+j} \nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k)|\mathcal{F}_k, \bar{B}_k] - \\
&\leq -3 \frac{\gamma_{\min} \alpha_{\min}}{k+2} P(B_k) E[\nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k)|\mathcal{F}_k, B_k] \\
&\quad - 3 \frac{\gamma_{\max}}{k} P(\bar{B}_k) E[\nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k)|\mathcal{F}_k, \bar{B}_k]. \tag{22}
\end{aligned}$$

We further notice that

$$\begin{aligned}
\|\nabla f(x_k)\|^2 &= E[\nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k)|\mathcal{F}_k] \\
&= P(B_k) E[\nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k)|\mathcal{F}_k, B_k] + P(\bar{B}_k) E[\nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k)|\mathcal{F}_k, \bar{B}_k].
\end{aligned}$$

Then,

$$-P(\bar{B}_k)E[\nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k)|\mathcal{F}_k, \bar{B}_k] = -\|\nabla f(x_k)\|^2 + P(B_k)E[\nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k)|\mathcal{F}_k, B_k],$$

and by (22) and Assumption A3 we get

$$\begin{aligned} E\left[-\sum_{j=0}^2 \alpha_{k+j} \gamma_{k+j} \nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k)|\mathcal{F}_k\right] &\leq -3 \frac{\gamma_{\min} \alpha_{\min}}{k+2} P(B_k) E[\nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k)|\mathcal{F}_k, B_k] \\ &\quad + 3 \frac{\gamma_{\max}}{k} (-\|\nabla f(x_k)\|^2 + P(B_k) E[\nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k)|\mathcal{F}_k, B_k]) \\ &= 3(P(B_k) E[\nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k)|\mathcal{F}_k, B_k]) \left(\frac{\gamma_{\max}}{k} - \frac{\gamma_{\min} \alpha_{\min}}{k+2}\right) \\ &\quad - \frac{\gamma_{\max}}{k} \|\nabla f(x_k)\|^2 \\ &\leq 3 \left(\left(\frac{\gamma_{\max}}{k} - \frac{\gamma_{\min} \alpha_{\min}}{k+2}\right) \left(\frac{\theta}{s} \|\nabla f(x_k)\|^2 + \frac{C}{s k^{\delta}}\right) - \frac{\gamma_{\max}}{k} \|\nabla f(x_k)\|^2\right) \\ &\leq 3 \left(-\|\nabla f(x_k)\|^2 \left(\frac{\gamma_{\max}}{k} + \frac{\theta \gamma_{\min} \alpha_{\min}}{s(k+2)} - \frac{\theta \gamma_{\max}}{s k}\right) + \frac{\gamma_{\max} C}{s k^{1+\delta}}\right) \\ &\leq -\frac{\tilde{C}_4}{k+2} \|\nabla f(x_k)\|^2 + \frac{\tilde{C}_5}{k^{1+\delta}} \end{aligned} \quad (23)$$

where $\tilde{C}_4 := 3(\theta \gamma_{\min} \alpha_{\min}/s)$, $\tilde{C}_5 := 3\gamma_{\max} C/s$ and the last inequality holds as $s \geq \theta$ by assumption.

Finally, combining (15), (16), (19) and (23) we obtain

$$\begin{aligned} E[f(x_{k+3}) - f_{low}|\mathcal{F}_k] &\leq f(x_k) - f_{low} - \left(\frac{\tilde{C}_4}{k+2} - \frac{\tilde{C}_3 + \tilde{C}_2 L/2}{k^2}\right) \|\nabla f(x_k)\|^2 + \frac{\tilde{C}_6}{k^{1+\delta}} \\ &= f(x_k) - f_{low} - \frac{1}{k} \left(\frac{\tilde{C}_4 k}{k+2} - \frac{\tilde{C}_3 + \tilde{C}_2 L/2}{k}\right) \|\nabla f(x_k)\|^2 + \frac{\tilde{C}_6}{k^{1+\delta}}, \end{aligned} \quad (24)$$

where $\tilde{C}_6 := \tilde{C}_3 \sqrt{G} + \tilde{C}_5 + \tilde{C}_2 G L/2$. Since $\frac{\tilde{C}_4 k}{k+2} \rightarrow \tilde{C}_4$ and $\frac{\tilde{C}_3 + \tilde{C}_2 L/2}{k} \rightarrow 0$ when $k \rightarrow \infty$, there exists $\bar{k} \in K$ such that for all $k \geq \bar{k}, k \in K$ we have

$$E[f(x_{k+3}) - f_{low}|\mathcal{F}_k] \leq f(x_k) - f_{low} - \frac{\bar{C}_4}{2k} \|\nabla f(x_k)\|^2 + \frac{\bar{C}_6}{k^{1+\delta}}. \quad (25)$$

Thus, we have just proved that

$$E[U_{\ell+1}|\mathcal{F}_\ell] \leq (1 + \beta_\ell) U_\ell + \xi_\ell - \rho_\ell, \quad (26)$$

holds for $\ell = 0, 1, 2, \dots$ where

$$U_{\ell+1} := f(x_{\bar{k}+3(\ell+1)}) - f_{low}, \quad \xi_\ell := \frac{\bar{C}_6}{(\bar{k} + 3\ell)^{1+\delta}}, \quad \rho_\ell := \frac{\bar{C}_4}{2(\bar{k} + 3\ell)} \|\nabla f(x_{\bar{k}+3\ell})\|^2, \quad \beta_\ell = 0.$$

Notice that $U_\ell, \xi_\ell, \rho_\ell, \beta_\ell$ are nonnegative and \mathcal{F}_ℓ -measurable. Moreover, there holds $\sum_\ell \xi_\ell < \infty$ and obviously $\sum_\ell \beta_\ell < \infty$. Thus, applying Theorem 3.1 we conclude that

$$\sum_{\ell=1}^{\infty} \rho_\ell = \frac{\bar{C}_4}{2} \sum_{\ell=1}^{\infty} \frac{\|\nabla f(x_{\bar{k}+3\ell})\|^2}{\bar{k} + 3\ell} < \infty \quad \text{a.s.}$$

This further implies that $\liminf_{\ell \rightarrow \infty} \|\nabla f(x_{\bar{k}+3\ell})\|^2 = 0$ a.s., which completes the proof. \square

3.1.2 Algorithmic modifications

We have proved the a.s. convergence result under assumptions A1, A2 and an additional assumption on the sampling strategy (A3). Focusing on the case $m > 1$, we can avoid additional assumptions and retain the a.s convergence by slightly modifying the steps S3-S5 of SLiSeS algorithm as follows. Take $\delta > 0$, arbitrary small and $\tilde{\gamma} > 0$.

S3' If $\text{mod}(k, m) = 0$ set $\gamma_k = \tilde{\gamma}/k$ and go to step S5'. Else, set $s_{k-1} = x_k - x_{k-1}$, $y_{k-1} = g_k - g_{k-1}$, set $c_k = \|s_{k-1}\|^2 / (s_{k-1}^T y_{k-1})$ and go to step S4'.

S4' $\gamma_k = \min\{\gamma_{max}, \max\{\gamma_{min}, c_k\}\} / k^{1+\delta}$.

S5' Set $d_k = -\gamma_k \nabla f_{\mathcal{N}_k}(x_k)$. If $\text{mod}(k, m) = 0$, set $\alpha_k = 1$ and go to step S7. Else, go to step S6.

We note that the choice $\gamma_k = \tilde{\gamma}/k$ when $\text{mod}(k, m) = 0$ yields \mathcal{F}_k -measurable steplengths whenever the subsample is recomputed as γ_k and α_k are not dependent of the subsample \mathcal{N}_k . Notice that taking $\gamma_k = \tilde{\gamma}/k$ instead of $\gamma_k = 1/k$ in step S3 does not change the theoretical analysis and for simplicity we will consider $\tilde{\gamma} = 1$ in the further analysis. The choice of $\tilde{\gamma}$ might influence the numerical performance of the algorithm and in our experiments we used $\tilde{\gamma} = 1/\|g_0\|$ where g_0 is a stochastic estimator of the gradient at x_0 , for further details see Section 4. We further underline that we use stepsizes diminishing as $O(\frac{1}{k^{1+\delta}})$ rather than as $O(\frac{1}{k})$. These two conditions allow to prove convergence without relying on A3 as shown in the following theorem.

Theorem 3.3. *Suppose that the Assumptions A1 and A2 hold. Let $\{x_k\}$ be a sequence generated by modified SLiSeS algorithm with $m > 1$ and steps S3'-S5' in place of of S3-S5. Then, with probability 1, there holds $\liminf_{k \rightarrow \infty} \|\nabla f(x_k)\| = 0$.*

Proof. For simplicity we take $\tilde{\gamma} = 1$ in the proof. Following the same steps as in Theorem 3.2 until (19), we obtain (13), (15) and (16). Then, to upper bound $E[-\sum_{j=0}^2 \alpha_{k+j} \gamma_{k+j} \nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k) | \mathcal{F}_k]$, notice that $\gamma_k = 1/k$ and $\alpha_k = 1$, so they are not dependent on the subsample \mathcal{N}_k and hence they are \mathcal{F}_k -measurable. Thus,

$$\begin{aligned} E[-\sum_{j=0}^2 \alpha_{k+j} \gamma_{k+j} \nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k) | \mathcal{F}_k] &= -\frac{1}{k} \nabla f(x_k)^T E[\nabla f_{\mathcal{N}_k}(x_k) | \mathcal{F}_k] \\ &\quad - E[\sum_{j=1}^2 \alpha_{k+j} \gamma_{k+j} \nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k) | \mathcal{F}_k] \end{aligned} \quad (27)$$

Moreover, there holds

$$-\sum_{j=1}^2 \alpha_{k+j} \gamma_{k+j} \nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k) \leq \sum_{j=1}^2 \frac{\gamma_{max}}{k^{1+\delta}} \|\nabla f(x_k)\| \|\nabla f_{\mathcal{N}_k}(x_k)\|.$$

Thus, we obtain

$$E\left[-\sum_{j=0}^2 \alpha_{k+j} \gamma_{k+j} \nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k) | \mathcal{F}_k\right] = -\frac{1}{k} \nabla f(x_k)^T E[\nabla f_{\mathcal{N}_k}(x_k) | \mathcal{F}_k] \\ + \frac{2\gamma_{max}}{k^{1+\delta}} \|\nabla f(x_k)\| E[\|\nabla f_{\mathcal{N}_k}(x_k)\| | \mathcal{F}_k] \quad (28)$$

$$\leq -\frac{1}{k} \|\nabla f(x_k)\|^2 \\ + \frac{4\gamma_{max}(\sqrt{G}+1)}{k^{1+\delta}} (\sqrt{G} + \|\nabla f(x_k)\|^2), \quad (29)$$

where the last inequality comes from (18), considering both cases $\|\nabla f(x_k)\| \leq 1, \|\nabla f(x_k)\| > 1$. Finally, combining (15) with (16), (19) and (29) we obtain

$$E[f(x_{k+3}) - f_{low} | \mathcal{F}_k] \leq f(x_k) - f_{low} - \frac{1}{k} \left(1 - \frac{4\gamma_{max}(\sqrt{G}+1) + \tilde{C}_3 + L\tilde{C}_2/2}{k^\delta}\right) \|\nabla f(x_k)\|^2 \\ + \frac{\tilde{C}_3\sqrt{G} + L\tilde{C}_2G/2 + 4(G + \sqrt{G})\gamma_{max}}{k^2}$$

and conclude that there exists $\bar{k} \in K$ such that for all $k \geq \bar{k}, k \in K$ we have

$$E[f(x_{k+3}) - f_{low} | \mathcal{F}_k] \leq f(x_k) - f_{low} - \frac{1}{2k} \|\nabla f(x_k)\|^2 + \frac{\tilde{C}_3\sqrt{G} + L\tilde{C}_2G/2 + 4(G + \sqrt{G})\gamma_{max}}{k^2}. \quad (30)$$

Continuing as in Theorem 3.2 we obtain the result. \square

3.2 Nonuniform sampling

In this subsection we discuss the case where the sampling is not uniform. Let us denote the indices that form the subsample \mathcal{N}_k with i_1^k, \dots, i_S^k , i.e., $\mathcal{N}_k = \{i_1^k, \dots, i_S^k\}$, and let us define

$$p_j^k := P(k_\ell = j), \quad j = 1, \dots, N,$$

where k_1, \dots, k_S are i.i.d. random integers and $P(k_\ell = j)$ is the probability of the event $k_\ell = j$. We allow nonuniform sampling in the following sense.

Assumption 4 (A4). *There exist constants $M_3 \geq 0$ and $\varepsilon > 0$ such that*

$$|p_j^k - \frac{1}{N}| \leq \frac{M_3}{k^\varepsilon}, \quad j = 1, \dots, N$$

The above assumption actually means that the sampling we will employ converges to the uniform sampling although we allow it to be nonuniform for arbitrary large number of initial iterations. For example, one can employ modified Adaptive Importance Sampling algorithms [24], for sample sizes of arbitrary size as follows. Starting with set $\pi_i^0, i = 1, \dots, N$, at a generic iteration $k \geq 1$ we define the set of relevant gradient norm values,

$$\pi_{i_j^k}^k = \|\nabla f_{i_j^{k-1}}(x_{k-1})\|, \quad j = 1, \dots, S,$$

and $\pi_i^k = \pi_i^{k-1}$ for any $i \notin \mathcal{N}_{k-1}$. Then, we calculate the probabilities as

$$p_j^k = \frac{1}{k^\varepsilon} \frac{\pi_j^k}{\sum_{i=1}^N \pi_i^k} + \left(1 - \frac{1}{k^\varepsilon}\right) \frac{1}{N}, \quad j = 1, \dots, N. \quad (31)$$

In order to compensate the bias that comes from nonuniform sampling, we state the assumption on similarity of the local cost functions [20]. The assumption is satisfied for a number of important problems, namely the problems coming from machine learning like logistic regression and for quadratic cost function for bounded x .

Assumption 5 (A5). *There exist constants M_1 and M_2 such that for every $x \in \mathbb{R}^n$ and every $i \in \mathcal{N}$ there holds*

$$\|\nabla f_i(x)\| \leq M_1 + M_2 \|\nabla f(x)\|.$$

For sampling defined above we have the following result. Instead of (3) we have sampling as in A4 and instead of assumption A2 we have A5. and we prove the a.s. convergence for the modified version of the algorithm.

Theorem 3.4. *Let A1, A4 and A5 hold. Let $\{x_k\}$ be a sequence generated by modified SLiSeS algorithm with steps S3'-S5' instead of S3-S5. Then, with probability 1, there holds $\liminf_{k \rightarrow \infty} \|\nabla f(x_k)\| = 0$.*

Proof. Following the same steps as in Theorem 3.2, we obtain (13) for k such that $\text{mod}(k, m) = 0$ and

$$E[f(x_{k+3}) - f_{low} | \mathcal{F}_k] \leq f(x_k) - f_{low} + E[\nabla f(x_k)^T (x_{k+3} - x_k) | \mathcal{F}_k] + \frac{\tilde{C}_2 L}{2k^2} E[\|\nabla f_{\mathcal{N}_k}(x_k)\|^2 | \mathcal{F}_k]. \quad (32)$$

Then, to upper bound $E[-\sum_{j=0}^2 \alpha_{k+j} \gamma_{k+j} \nabla f(x_k)^T \nabla f_{\mathcal{N}_k}(x_k) | \mathcal{F}_k]$, notice that $\gamma_k = 1/k$ and $\alpha_k = 1$, so they are \mathcal{F}_k -measurable. Then, we can proceed as in the proof of Theorem 3.3 and we obtain (28). Given that the sampling is nonuniform, let us start estimating the gradient norm using the similarity assumption A4. We have

$$\|\nabla f_{\mathcal{N}_k}(x_k)\| = \left\| \frac{1}{S} \sum_{j=1}^S \nabla f_{i_j^k}(x_k) \right\| \leq M_1 + M_2 \|\nabla f(x_k)\|. \quad (33)$$

Moreover,

$$E[\|\nabla f_{\mathcal{N}_k}(x_k)\|^2 | \mathcal{F}_k] \leq E[2M_1^2 + 2M_2^2 \|\nabla f(x_k)\|^2 | \mathcal{F}_k] = 2M_1^2 + 2M_2^2 \|\nabla f(x_k)\|^2. \quad (34)$$

Furthermore, given that $E[\nabla f_i(x_k) | \mathcal{F}_k] = \sum_{j=1}^N p_j^k \nabla f_j(x_k)$, $i = 1, \dots, N$ and

$$E[\nabla f_{\mathcal{N}_k}(x_k) | \mathcal{F}_k] = \sum_{j=1}^N p_j^k \nabla f_j(x_k) - \nabla f(x_k) + \nabla f(x_k) = \sum_{j=1}^N \left(p_j^k - \frac{1}{N}\right) \nabla f_j(x_k) + \nabla f(x_k),$$

using Assumptions A5 and A4, we get

$$\begin{aligned}
E[-\frac{1}{k}\nabla f(x_k)^T\nabla f_{\mathcal{N}_k}(x_k)|\mathcal{F}_k] &= -\frac{1}{k}\sum_{j=1}^N(p_j^k - \frac{1}{N})\nabla f(x_k)^T\nabla f_j(x_k) - \frac{1}{k}\|\nabla f(x_k)\|^2 \\
&\leq -\frac{1}{k}\|\nabla f(x_k)\|^2 + \frac{1}{k}\sum_{j=1}^N|p_j^k - \frac{1}{N}|\|\nabla f(x_k)\|\|\nabla f_j(x_k)\| \\
&\leq -\frac{1}{k}\|\nabla f(x_k)\|^2 + \frac{1}{k}\sum_{j=1}^N\frac{M_3}{k^\varepsilon}\|\nabla f(x_k)\|\|\nabla f_j(x_k)\| \\
&\leq -\frac{1}{k}\|\nabla f(x_k)\|^2 + \frac{M_3}{k^{1+\varepsilon}}\sum_{j=1}^N\|\nabla f(x_k)\|(M_1 + M_2\|\nabla f(x_k)\|) \\
&\leq -\frac{1}{k}\|\nabla f(x_k)\|^2 + \frac{NM_1M_3}{k^{1+\varepsilon}}\|\nabla f(x_k)\| + \frac{NM_2M_3}{k^{1+\varepsilon}}\|\nabla f(x_k)\|^2.
\end{aligned} \tag{35}$$

Again, considering both cases: $\|\nabla f(x_k)\| \leq 1$ and $\|\nabla f(x_k)\| > 1$ we conclude that

$$E[-\frac{1}{k}\nabla f(x_k)^T\nabla f_{\mathcal{N}_k}(x_k)|\mathcal{F}_k] \leq -\frac{1}{k}\|\nabla f(x_k)\|^2 + \frac{\bar{C}_1}{k^{1+\varepsilon}}\|\nabla f(x_k)\|^2 + \frac{\bar{C}_1}{k^{1+\varepsilon}}, \tag{36}$$

with $\bar{C}_1 = NM_3(M_1 + M_2)$. Now, using (32), (19), (28) and (36) we obtain

$$E[f(x_{k+3}) - f_{low}|\mathcal{F}_k] \leq f(x_k) - f_{low} - \frac{1}{k}\|\nabla f(x_k)\|^2 + \frac{\bar{C}_3}{k^{1+\tau}} + \frac{\bar{C}_4}{k^{1+\tau}}\|\nabla f(x_k)\|^2, \tag{37}$$

where $\tau := \min\{\delta, \varepsilon\}$, $\bar{C}_3 = 2\gamma_{max}M_1 + L\tilde{C}_2M_1^2 + \bar{C}_1$ and $\bar{C}_4 = \bar{C}_1 + L\tilde{C}_2M_2^2 + 2\gamma_{max}M_2$.

Finally, we conclude that there exists \bar{k} such that for $k \in K, k \geq \bar{k}$ we have

$$E[f(x_{k+3}) - f_{low}|\mathcal{F}_k] \leq f(x_k) - f_{low} - \frac{1}{2k}\|\nabla f(x_k)\|^2 + \frac{\bar{C}_3}{k^{1+\tau}},$$

and the statement follows analogously as in Theorem 3.2. \square

3.3 P-L and Strongly convex functions

Notice that the previous result holds for nonconvex problems. However, if we assume that the objective function is strongly convex, we can prove that the sequence of outer iterations converges a.s. towards the unique solution of the original problem.

Theorem 3.5. *Suppose that the assumptions of Theorem 3.2 hold and the function f is strongly convex. Then, the sequence $\{x_{mk-1}\}_{k \in \mathbb{N}}$ converges to the solution of problem (1) x^* almost surely.*

Proof. Let us denote with f^* the optimal value of problem (1). Theorem 3.2 implies that $\liminf_{k \rightarrow \infty} \|\nabla f(x_{mk-1})\|^2 = 0$ a.s. The strong convexity further implies the existence of K such that $\lim_{k \in K} x_{mk-1} = x^*$ a.s. Therefore, $\liminf_{k \rightarrow \infty} f(x_{mk-1}) = f^*$ a.s. Moreover, using the same notation as in the proof of Theorem 3.2 we conclude that Theorem 3.1 implies that U_l converges a.s.

which means that $\{f(x_{mk-1})\}_{k \in \mathbb{N}}$ converges a.s. Combining everything together, we conclude that the whole sequence $\{f(x_{mk-1})\}_{k \in \mathbb{N}}$ converges to f^* a.s. and due to strong convexity we conclude $\lim_{k \rightarrow \infty} x_{mk-1} = x^*$ a.s. \square

Relaxing the strong convexity assumption, we get the following result.

Theorem 3.6. *Suppose that the assumptions of Theorem 3.2 hold and the function f satisfies the Polyak-Lojasiewicz (P-L) condition, i.e., for every $x \in \mathbb{R}^n$ there holds*

$$\|\nabla f(x)\|^2 \geq 2\mu(f(x) - f^*) \quad (38)$$

where μ is a positive constant and f^* is the optimal value of problem (1). Then, the sequence $\{f(x_{mk-1})\}_{k \in \mathbb{N}}$ converges to f^* almost surely.

Proof. Theorem 3.2 implies that $\liminf_{k \rightarrow \infty} \|\nabla f(x_{mk-1})\|^2 = 0$ a.s. and (38) implies that $\liminf_{k \rightarrow \infty} f(x_{mk-1}) = f^*$ a.s.. Proceeding as in the proof of Theorem 3.5 it can be proved that $\lim_{k \rightarrow \infty} f(x_{mk-1}) = f^*$ a.s.. \square

Remark. The two previous statements holds if instead of the assumptions in Theorem 3.2 we impose the assumptions in Theorems 3.3 or 3.4.

4 Numerical results

To give further insight into the proposed subsampling strategy, and to illustrate the practical performance of the SLiSeS algorithm, we present the results of some numerical experiments. We consider two versions of our proposed algorithm: SLiSeS using uniform sampling (SLiSeS-UNI), analyzed in Section 3.1, and SLiSeS using Adaptive Importance Sampling (SLiSeS-AIS), analyzed in Section 3.2 with probabilities defined through (31) with $\varepsilon = 1$ and $\pi_i^0 = 1, i = 1, \dots, N$. Another variant, proposed and analyzed in Section 3.1.2, is to divide c_k by $k^{1+\delta}$, where $\delta > 0$ is a small given number. We will denote this variant as the Modified SLiSeS method. All these versions will be compared with several recently proposed schemes in the literature for solving (1).

Although the SLiSeS algorithm can be applied more generally, in our experiments we focus on strongly convex quadratic functions and also on L2-regularized logistic regression problems, as problems that satisfy our assumption A5. The considered quadratic functions, for $1 \leq i \leq N$, are given by

$$f_i(x) = \frac{1}{2}(x - b_i)^\top A_i(x - b_i),$$

where $b_i \in \mathbb{R}^n$ and the symmetric positive definite matrix $A_i \in \mathbb{R}^{n \times n}$ are obtained as in [22], i.e., vectors b_i are extracted from the Uniform distribution on $[1,31]$, independently from each other. Matrices A_i are of the form $A_i = Q_i D_i Q_i^\top$, where D_i is a diagonal matrix with Uniform distribution on $[1,101]$ and Q_i is the matrix of orthonormal eigenvectors of $\frac{1}{2}(C_i + C_i^\top)$, and the matrix C_i has components drawn independently from the standard Normal distribution. For the logistic regression problems, for $1 \leq i \leq N$, the function f_i is given by

$$f_i(x) = \log(1 + \exp(-b_i(a_i^\top x))) + \frac{\lambda}{2}\|x\|_2^2,$$

where the vectors $a_i \in \mathbb{R}^n$ and the labels $b_i \in \{-1, 1\}$ are given, and the regularization parameter λ is set to 10^{-4} . The vectors a_i and the labels b_i are obtained from three real data sets: Voice (LSVT Voice Rehabilitation) [30] ($n = 309, N = 126$), Park (Parkinson's Disease Classification) [27] ($n = 754,$

$N = 756$), and the Cina0 econometrics dataset ($n = 132$, $N = 16033$) which was downloaded from the CINA website¹.

In all our experiments, we report the computational cost measured by the number of function evaluations versus the observed decrease in the objective function. To calculate the cumulative number of function evaluations we add $S = |\mathcal{N}_k|$ each time we evaluate the function $f_{\mathcal{N}_k}$, and the gradient evaluation comes for free; see, e.g., [2, p. 27]. All runs are terminated when a maximum number of iterations (*maxiter*) is reached. The required input parameters in the SLiSeS algorithms are fixed as follows: $x_0 = 0$, $\eta = 10^{-4}$, $t_k = 1/2^k$, $\gamma_{min} = 10^{-8}$, $\gamma_{max} = 10^8$ and $\tilde{\gamma} = 1/\|g_0\|$ for the Modified SLiSeS. Concerning the sample size, unless otherwise specified, we set $S = |\mathcal{N}_k| = 1$ for all k .




Figure 1: Performance of SLiSeS-UNI on strictly convex quadratics for $m = 1, 3, 5, 10$ (inner iterations), $N = 1000$, *maxiter* = 50, and $n = 10$ (left) or $n = 100$ (right).

We begin by exploring the behavior of Algorithm SLiSeS-UNI and Modified SLiSeS-UNI with $\delta = 0.1$ for different values of the number of internal iterations m , in which we keep the same sample. Let us start by considering strictly convex quadratic functions. In Figure 1 we report the performance of Algorithm SLiSeS for several values of m from 1 to 10, $N = 1000$, *maxiter* = 50, and two different values of n : 10 and 100. We can observe that, for both values of n , the choice $m = 1$ (i.e., the standard strategy of changing the sample at each iteration) produces a much worse behavior than any of the other 3 choices of $m > 1$. We also note, for this particular experiment as well as many others we have run, that the best behavior is observed when $2 \leq m \leq 5$. In Figures 2 and 3 we focus our attention on that interval, for the strictly convex quadratic case and also for some logistic regression problems (Voice and Park datasets). For these latter test problems we report results with the modified version as we observe a general better behaviour of the modified versions with respect to the unmodified counterpart on this class of problems, as it will be shown in Figure 6. We note that for $m = 3$, $m = 4$, $m = 5$ the performance is better than for $m = 2$, and also that none of the choices 3, 4, or 5 shows a clear advantage over the other two. However, in the average, the value $m = 3$ seems to be the one that shows a better performance.

For our second set of experiments, we compare SLiSeS-UNI and Modified SLiSeS-AIS with $\delta = 0.1$ with the full sampled version of the SLiSeS algorithm, i.e., when $S = |\mathcal{N}_k| = N$ for all k . Notice that

¹<http://www.causality.inf.ethz.ch/data/CINA.html>

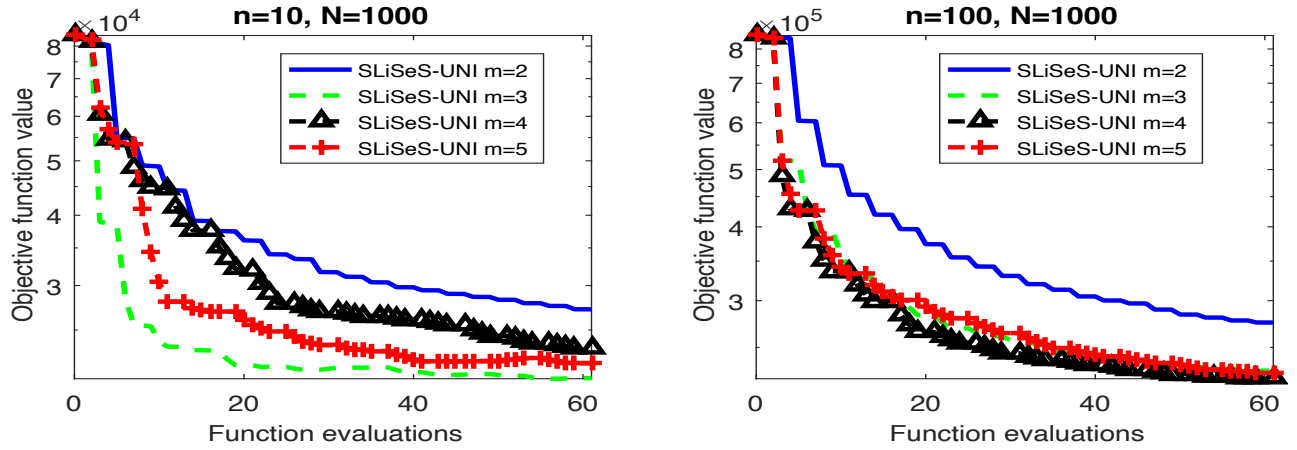


Figure 2: Performance of SLiSeS-UNI on strictly convex quadratics for $m = 2, 3, 4, 5$ (inner iterations), $N = 1000$, $maxiter = 50$, and $n = 10$ (left) or $n = 100$ (right).

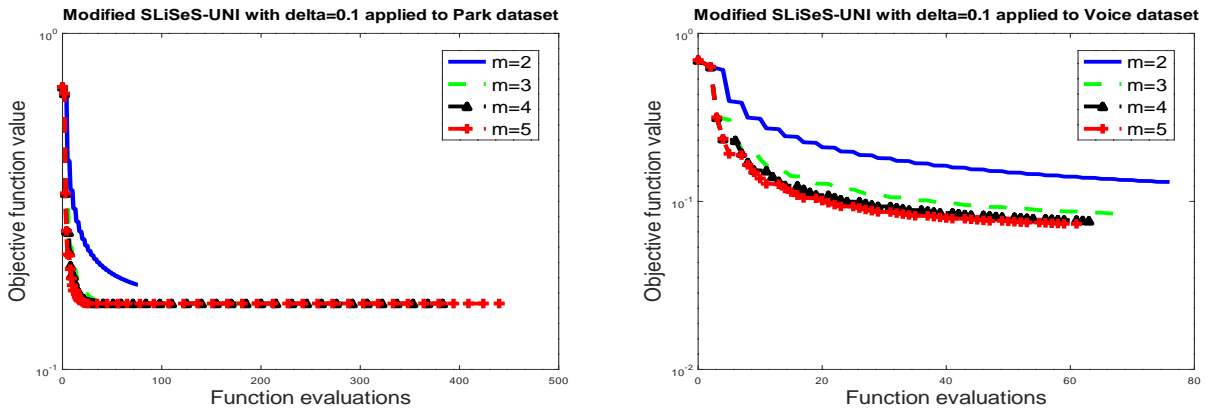


Figure 3: Performance of Modified SLiSeS-UNI with $\delta = 0.1$ on the Park dataset (left) and on the Voice dataset (right) for $m = 2, 3, 4, 5$ (inner iterations), and $maxiter = 50$.

this so-called SpectralLS-Full method matches the spectral gradient method globalized with the line search procedure described in Algorithm 2, for which global convergence is guaranteed; see [6] and references therein. Hence, in the SpectralLS-Full method, there is no need to divide by k at Step S4 of Algorithm 1. To illustrate the benefits of using spectral-BB steplengths in our approach, we also consider the same Algorithm 1 with two additional parameter settings. In the first one we fix $\gamma_k = 1$, so $d_k = -g_k$, while in the second one we fix $\gamma_{min} = \gamma_{max} = 1$, and so $d_k = -(1/k)g_k$. In both of these latter procedures we inhibit the employment of BB steplengths in the SLiSeS framework, and neither of these two settings is competitive with SLiSeS employing the spectral-BB step sizes. In Figure 4 we report the performance of Algorithms SLiSeS-UNI ($m = 3$), SpectralLS-Full, SLiSeS-UNI with $d_k = -g_k$ and Modified SLiSeS-AIS with $\delta = 0.1$ when applied to strictly convex quadratics for $maxiter = 50$ and different combinations of n and N . The case SLiSeS-UNI with $\gamma_{min} = \gamma_{max} = 1$ is not shown in the graphs as it is quite the same as $\gamma_k = 1$; i.e SLiSeS-UNI with $d_k = -g_k$. We also note that for any combination of n and N , the SpectralLS-Full method requires a much higher number of function evaluations, so it is not competitive.

Our convergence analysis in Section 3 was developed using the damping strategy of dividing by k the spectral step length c_k at Step S4 of the SLiSeS algorithm. Nevertheless, in order to assess the benefits of that damping strategy, in our next set of experiments we explore the practical performance of the SLiSeS algorithm when c_k is not divided by k at Step S4. In Figure 5 we show the performance of SLiSeS-UNI ($m = 3$) with the two mentioned variants, when applied to strictly convex quadratics for $maxiter = 100$, $n = 100$, $N = 1000$, and two different values of δ (0.1 and 1).

It can be noted that in the case of strictly convex quadratics the version that does not divide by k produces unstable behavior. Moreover, for many of the experiments we have tried, in the case of logistic regression functions that variant does not converge. We can also observe that the damping strategy using $\delta = 0.1$ shows a better performance than using $\delta = 1$. In Figure 6 we present the advantages of using the damping strategy with $\delta = 0.1$ in the case of two logistic regression functions. These last two figures illustrate that the damping strategy not only guarantees convergence but also adds practical advantages to the SLiSeS algorithm.

Finally, in the last set of experiments we are interested in comparing the SLiSeS method with the natural competitors that exist in the literature to solve problem (1). The so-called Stochastic Gradient Descent (SGD) continues to receive significant attention, and is still a widely used technique for machine learning applications; see, e.g., [8] and references therein. The SGD method can be interpreted as a simplified case of the SLiSeS algorithm where the step length γ_k is obtained either as a sufficiently small constant to guarantee convergence, or in a diminishing way such that $\sum \gamma_k = \infty$ and $\sum \gamma_k^2 < \infty$ to guarantee convergence (the most common choice is $\gamma_k = 1/k$). In either case, the SGD method does not employ a line search strategy. In Figure 7 we report the performance of the modified variants of SLiSeS ($m = 3$ and $\delta = 0.1$) and the SGD method with step length $1/k$, for $maxiter = 100$, when applied to a strictly convex quadratic when $n = 100$ and $N = 1000$, and also to the Cina0 dataset. In Figure 8 we show the performance of the same three methods when applied to the Park dataset and also to the Voice dataset. The results observed in both figures are indicative of the known weaknesses of the SGD method, and also illustrate that the modified SLiSeS Algorithms exhibits better convergence behavior.

More recently, the Barzilai and Borwein spectral step length has been used in connection with stochastic gradient ideas in the so-called SVRG-BB family of methods; see [29]. The SVRG-BB uses a full gradient (i.e., $S = N$) every p iterations (in [29], it is recommended to use $p = 2n$), and it depends on a certain parameter $\eta_0 > 0$ (in [29], it is recommended to set $\eta_0 = 0.01$). Two more members of that family are proposed and discussed in [29], the so-called SGD-BB and the SGD-BB




Figure 4: Performance of SLiSeS-UNI ($m = 3$), SpectralLS-Full, Modified SLiSeS-AIS (for $\delta = 0.1$) and SLiSeS-UNI with $d_k = -g_k$, for strictly convex quadratic functions, different combinations of n and N , and $maxiter = 50$.

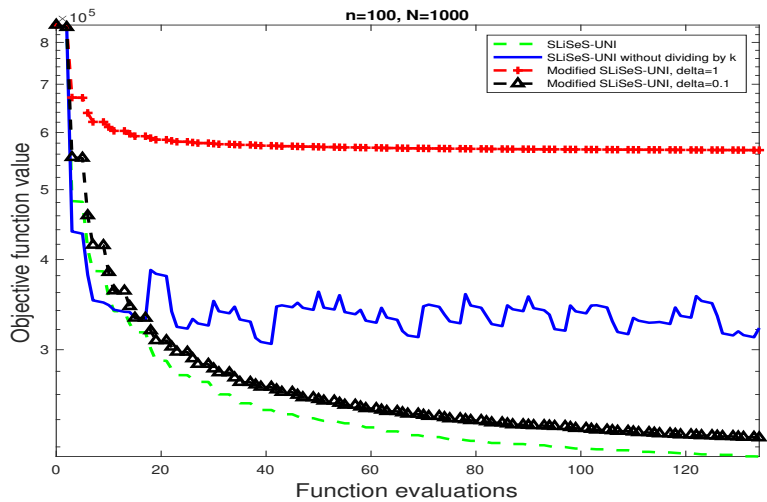


Figure 5: Performance of the SLiSeS-UNI method without dividing by k at Step S4, and the Modified SLiSeS-UNI method (for $\delta = 0.1$ and $\delta = 1$) when compared with the original versions of SLiSeS-UNI for $m = 3$, $maxiter = 100$, $n = 100$, $N = 1000$ and strictly convex quadratic functions.

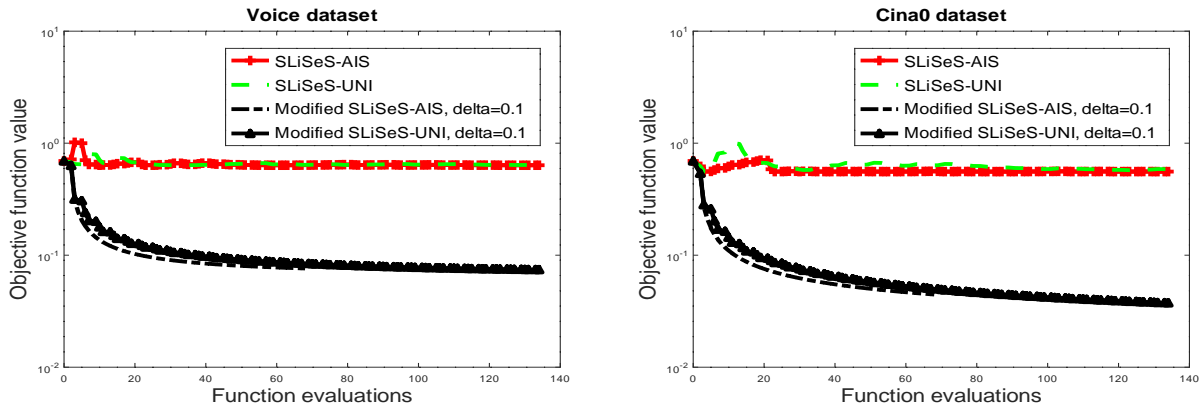


Figure 6: Performance of the SLiSeS methods (AIS and UNI) and their modified versions (with $\delta = 0.1$) for $m = 3$, on the Voice dataset (left), and on the Cina0 dataset (right).




Figure 7: Performance of the modified SLiSeS-UNI method ($m = 3$ and $\delta = 0.1$) and the SGD method with step length $1/k$ ($maxiter = 100$) on a strictly convex quadratic when $n = 100$ and $N = 1000$ (left), and performance of the modified versions of SLiSeS ($\delta = 0.1$) and the SGD method with step length $1/k$ ($maxiter = 100$) on the Cina0 dataset (right).




Figure 8: Performance of the modified SLiSeS methods ($m = 3$ and $\delta = 0.1$) and the SGD method with step length $1/k$, for $maxiter = 100$, on the Park dataset (left) and the Voice dataset (right).

with a specialized smoothing technique (both of them avoid calculating full gradients). These last two options depend on four parameters: $\beta > 0$, $\eta_0 > 0$, $\eta_1 > 0$, and p . (in [29], it is recommended to set $\eta_0 = \eta_1 = 0.01$, $p = n$, and $\beta = 1/p$). In [29], convergence properties of the SVRG-BB family are established for strongly convex functions.




Figure 9: Performance of the SLiSeS-UNI ($m = 3$) on a strictly convex quadratic when compared with SVRG-BB ($p = 2n$ and $\eta_0 > 0.01$), SGD-BB ($p = n$, $\beta = 1/p$, and $\eta_0 = \eta_1 = 0.01$), and SGD-BB with a smoothing technique (same parameters as SGD-BB) for $maxiter = 100$, $N = 1000$, and $n = 10$ (left) and $n = 100$ (right).

In Figure 9 we report, for strictly convex quadratics, the performance of the SLiSeS-UNI ($m = 3$) method and the three described variants of the SVRG-BB family of methods: SVRG-BB, SGD-BB, and SGD-BB with a smoothing technique, when $maxiter = 100$, $N = 1000$, and $n = 10$ (left) and $n = 100$ (right). We note that the SVRG-BB method has a much worse performance than all the other competitors, and this is clearly due to the use of a full gradient every p iterations. We note that the best performance is obtained by the SLiSeS-UNI ($m = 3$) method. We also note that for small values of n the smoothing technique adds some benefit to the SGD-BB method, but once n increases the smoothing effect tends to vanish.

5 Conclusions

Stochastic algorithms based on the negative gradient direction remain one of the most convenient options for solving large-scale finite-sum optimization problems. Among these, the well-known SGD method and its variants present attractive theoretical advantages since they allow establishing worst-case complexity results. However, from a practical point of view, the SGD method has some known weaknesses and can exhibit very slow convergence.

Recently, focusing on practical behavior, some ideas have incorporated the use of non-monotonic step lengths such as the Barzilai and Borwein (BB) spectral one. These ideas have recently been proposed and analyzed in the literature, either avoiding full gradient evaluations altogether or by

allowing full gradient evaluations at some periodic iterations. For all these gradient-type methods, including the SGD method and its variants, a common denominator has been to keep the tradition of changing the sample at each iteration.

In this work, following the line of incorporating the BB spectral step length, and accepting its non-monotonic behavior, we proposed and analyzed a novel strategy: keeping the sample unchanged for a prefixed number of consecutive iterations before sampling again. The motivation behind this strategy is to promote the widely observed sweeping-spectrum behavior associated with spectral BB step lengths. The variants that emerge from this strategy have been analyzed, and convergence results have been obtained for non-convex smooth functions, for both uniform and nonuniform sampling. To illustrate and deepen this sampling strategy, we show a variety of numerical experiments that include comparisons with natural candidates from the literature. Our most important conclusion is that this strategy, based on the experiments shown, is promising and effectively improves the behavior of gradient-type methods. It deserves to be further analyzed with specific applications in mind.

Acknowledgements. We would like to thank an anonymous reviewer for constructive comments and suggestions that helped us to improve the final version of this paper.

Conflicts of interest. No potential conflict of interest was reported by the authors.

Funding. The first author acknowledges financial support received by the INdAM GNCS and by PNRR - Missione 4 Istruzione e Ricerca - Componente C2 Investimento 1.1, Fondo per il Programma Nazionale di Ricerca e Progetti di Rilevante Interesse Nazionale (PRIN) funded by the European Commission under the NextGeneration EU programme, project “Advanced optimization METHods for automated central veIn Sign detection in multiple sclerosis from magneTic resonAnce imaging (AMETISTA)”, code: P2022J9SNP, MUR D.D. financing decree n. 1379 of 1st September 2023 (CUP E53D23017980001), project “Numerical Optimization with Adaptive Accuracy and Applications to Machine Learning”, code: 2022N3ZNAX MUR D.D. financing decree n. 973 of 30th June 2023 (CUP B53D23012670006). The second and third authors were financially supported by the Science Fund of the Republic of Serbia, Grant no. 7359, Project LASCADO. The fourth author was financially supported by Fundação para a Ciência e a Tecnologia (Portuguese Foundation for Science and Technology) under the scope of the projects UIDB/MAT/00297/2020 (doi.org/10.54499/UIDB/00297/2020), and UIDP/MAT/00297/2020 (doi.org/10.54499/UIDP/00297/2020) (Centro de Matemática e Aplicações).

Data availability. The codes and datasets generated during and/or analyzed during the current study are available from the authors on reasonable request.

References

- [1] J. Barzilai and J. Borwein, *Two point step gradient methods*, IMA J. Numer. Anal. 8, (1988), 141–148.
- [2] S. Bellavia, N. Krklec Jerinkić, and G. Malaspina, *Subsampled nonmonotone spectral gradient methods*, Commun. Appl. Ind. Math. 11(1), (2020), 19–34.

- [3] S. Bellavia, T. Bianconcini, N. Krejić and B. Morini, *Subsampled first-order optimization methods with applications in imaging*, Handbook of Mathematical Models and Algorithms in Computer Vision and Imaging, Chen, K., Schonlieb, C., Tai, X., Younces, L. (eds.), Springer, (2021).
- [4] S. Bellavia, N. Krejić, B. Morini and S. Rebegoldi, *A stochastic first-order trust-region method with inexact restoration for finite-sum minimization*, Computational Optimization and Applications, 84, (2023), 53–84.
- [5] D. P. Bertsekas, *Nonlinear programming*, second ed., Athena Scientific Optimization and Computation Series, Athena Scientific, Belmont, MA, 1999.
- [6] E. G. Birgin, J. M. Martínez, and M. Raydan, *Spectral Projected Gradient methods: Review and Perspectives*, Journal of Statistical Software, 60(3), (2014), 1–21.
- [7] R. Bollapragada, R. Byrd and J. Nocedal, *Adaptive Sampling Strategies for Stochastic Optimization*, *SIAM Opt.* 28(4) (2018)
- [8] L. Bottou, F.E. Curtis and J. Nocedal, *Optimization Methods for Large-Scale Machine Learning*, SIAM Review, 60(2), (2018), 223–311.
- [9] F. E. Curtis, K. Scheinberg and R. Shi, *A Stochastic Trust Region Algorithm Based on Careful Step Normalization*, INFORMS Journal on Optimization 1(3), (2019), 200–220.
- [10] F. E. Curtis and K. Scheinberg, *Adaptive Stochastic Optimization: A Framework for Analyzing Stochastic Optimization Algorithms*, IEEE Signal Processing Magazine, 37 (5), (2020) 32–42.
- [11] Y.H. Dai and R. Fletcher, *On the asymptotic behaviour of some new gradient methods*, Math. Programming, 103, (2005), 541–559.
- [12] Y.-H. Dai, W. W. Hager, K. Schittkowski, and H. Zhang, *The cyclic Barzilai-Borwein method for unconstrained optimization*, IMA J. Numer. Anal. 26(3), (2006), 604–627.
- [13] R. De Asmundis, D. Di Serafino, F. Riccio, and G. Toraldo, *On spectral properties of steepest descent methods*, IMA J. Numer. Anal. 33, (2013), 1416–1435.
- [14] D. Di Serafino, V. Ruggiero, G. Toraldo, and L. Zanni, *On the steplength selection in gradient methods for unconstrained optimization*, Appl. Math. Comput. 318, (2018), 176–195.
- [15] C. Fan, C. Thrampoulidis and M. Schmidt, *Fast Convergence of Random Reshuffling Under Over-Parameterization and the Polyak-Lojasiewicz Condition* in Machine Learning and Knowledge Discovery in Databases: Research Track, editors: D. Koutra, C. Plant, M. Gomez Rodriguez, E. Baralis, F. Bonchi, Springer Nature Switzerland, (2023), 301–315.
- [16] G. Ferrandi, M.E. Hochstenbach, and N. Krejić, *A harmonic framework for stepsize selection in gradient methods*, Computational Optimization and Applications 85, (2023), 75–106.
- [17] R. Fletcher, *Low storage methods for unconstrained optimization*, Lectures in Applied Mathematics, 126, (1990), 165–179.
- [18] R. Fletcher, *On the Barzilai-Borwein method*, In: L. Qi, K. Teo, X. Yang, P. M. Pardalos, D. Hearn (eds.) Optimization and Control with Applications, Applied Optimization, Kluwer, Vol. 96, (2005), pp. 235–256.

- [19] G. Frassoldati, L. Zanni, and G. Zanghirati, *New adaptive stepsize selections in gradient methods*, Journal of industrial and management optimization 4(2), (2008), 299–312.
- [20] M.P. Friedlander and M. Schmidt, *Hybrid deterministic-stochastic methods for data fitting*, SIAM J. Sci. Comput. 34(3), (2012), A1380–A1405.
- [21] W. Glunt, T. L. Hayden, and M. Raydan, *Molecular conformations from distance matrices*, Journal of Computational Chemistry, 14, (1993), 114–120.
- [22] D. Jakovetić, N. Krejić, and N. K. Jerinkić, *Exact spectral-like gradient method for distributed optimization*, Comput. Optim. Appl. 74, (2019), 703–728.
- [23] N. Krejić and N. Krklec Jerinkić, *Spectral projected gradient method for stochastic optimization*, Journal of Global Optimization 73, (2018), 59–81.
- [24] H. Liu, X. Wang, J. Li, and A. M.-C. So, *Low-Cost Lipschitz-Independent Adaptive Importance Sampling of Stochastic Gradients*, 2020 25th International Conference on Pattern Recognition (ICPR), Milan, Italy, (2021), pp. 2150–2157.
- [25] M. Raydan, *On the Barzilai and Borwein choice of steplength for the gradient method*, IMA J. Numer. Anal, 13, (1993), 321–326.
- [26] H. Robbins and D. Siegmund, *A convergence theorem for non negative almost supermartingales and some applications*, Optimizing methods in statistics (Proc. Sympos., Ohio State Univ., Columbus, Ohio, 1971), (1971), 233–257.
- [27] C.O. Sakar, G. Serbes, A. Gunduz, H.C. Tunc, H. Nizam, B.E. Sakar, M. Tutuncu, T. Aydin, M.E. Isenkul, and H. Apaydin, *A comparative analysis of speech signal processing algorithms for Parkinson’s disease classification and the use of the tunable Q-factor wavelet transform*, Applied Soft. Computing 74, (2019), 255–263.
- [28] K. Sopyła and P. Drozda, *Stochastic gradient descent with Barzilai-Borwein update step for SVM*, Information Sciences 316, (2015), 218–233.
- [29] C. Tan, S. Ma, Y.-H. Dai, and Y. Qian, *Barzilai-borwein step size for stochastic gradient descent*, Advances in Neural Information Processing Systems 29, (2016), 685–693.
- [30] A. Tsanas, M. A. Little, C. Fox and L. O. Ramig. *Objective Automatic Assessment of Rehabilitative Speech Treatment in Parkinson’s Disease*, IEEE Trans. Neural Syst. Rehabil. Eng. 22(1), (2014), 181–190.
- [31] Z. Yang, C. Wang, Z. Zhang, and J. Li, *Random Barzilai-Borwein step size for mini-batch algorithms*, Engineering Applications of Artificial Intelligence 72, (2018), 124–135.
- [32] Q. Zou and F. Magoulès, *Delayed gradient methods for symmetric and positive definite linear systems*, SIAM Review 34(3), (2022), 517–553.