



# Semi-Stochastic Gradient Descent Methods

Jakub Konečný\* and Peter Richtárik

School of Mathematics, University of Edinburgh, Edinburgh, United Kingdom

## OPEN ACCESS

### Edited by:

Darinka Dentcheva,  
Stevens Institute of Technology,  
United States

### Reviewed by:

Yu Du,  
Rutgers University, United States  
Vladimir Shikhman,  
Technische Universität Chemnitz,  
Germany  
Uday V. Shanbhag,  
Pennsylvania State University,  
United States

### \*Correspondence:

Jakub Konečný  
kubo.konecny@gmail.com

### Specialty section:

This article was submitted to  
Optimization,  
a section of the journal  
Frontiers in Applied Mathematics and  
Statistics

Received: 24 October 2016

Accepted: 08 May 2017

Published: 23 May 2017

### Citation:

Konečný J and Richtárik P (2017)  
Semi-Stochastic Gradient Descent  
Methods. *Front. Appl. Math. Stat.* 3:9.  
doi: 10.3389/fams.2017.00009

In this paper we study the problem of minimizing the average of a large number of smooth convex loss functions. We propose a new method, S2GD (Semi-Stochastic Gradient Descent), which runs for one or several epochs in each of which a single full gradient and a random number of stochastic gradients is computed, following a geometric law. For strongly convex objectives, the method converges linearly. The total work needed for the method to output an epsilon-accurate solution in expectation, measured in the number of passes over data, is proportional to the condition number of the problem and inversely proportional to the number of functions forming the average. This is achieved by running the method with number of stochastic gradient evaluations per epoch proportional to conditioning of the problem. The SVRG method of Johnson and Zhang arises as a special case. To illustrate our theoretical results, S2GD only needs the workload equivalent to about 2.1 full gradient evaluations to find a  $10e-6$  accurate solution for a problem with  $10e9$  functions and a condition number of  $10e3$ .

**Keywords:** stochastic gradient, variance reduction, empirical risk minimization, linear convergence, convex optimization

## 1. INTRODUCTION

Many problems in data science (e.g., machine learning, optimization, and statistics) can be cast as loss minimization problems of the form

$$\min_{x \in \mathbb{R}^d} f(x), \quad (1)$$

where

$$f(x) \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n f_i(x). \quad (2)$$

Here  $d$  typically denotes the number of features / coordinates,  $n$  the number of examples, and  $f_i(x)$  is the loss incurred on example  $i$ . That is, we are seeking to find a predictor  $x \in \mathbb{R}^d$  minimizing the average loss  $f(x)$ . In big data applications,  $n$  is typically very large; in particular,  $n \gg d$ .

Note that this formulation includes more typical formulation of  $L_2$ -regularized objectives— $f(x) = \frac{1}{n} \sum_{i=1}^n \tilde{f}_i(x) + \frac{\lambda}{2} \|x\|^2$ . We hide the regularizer into the function  $f_i(x)$  for the sake of simplicity of resulting analysis.

## 1.1. Motivation

Let us now briefly review two basic approaches to solving problem (1).

1. *Gradient Descent*. Given  $x_k \in \mathbb{R}^d$ , the gradient descent (GD) method sets

$$x_{k+1} = x_k - hf'(x_k),$$

where  $h$  is a stepsize parameter and  $f'(x_k)$  is the gradient of  $f$  at  $x_k$ . We will refer to  $f'(x)$  by the name *full gradient*. In order to compute  $f'(x_k)$ , we need to compute the gradients of  $n$  functions. Since  $n$  is big, it is prohibitive to do this at every iteration.

2. *Stochastic Gradient Descent (SGD)*. Unlike gradient descent, stochastic gradient descent [1, 2] instead picks a random  $i$  (uniformly) and updates

$$x_{k+1} = x_k - hf'_i(x_k).$$

Note that this strategy drastically reduces the amount of work that needs to be done in each iteration (by the factor of  $n$ ). Since

$$\mathbf{E}(f'_i(x_k)) = f'(x_k),$$

we have an unbiased estimator of the full gradient. Hence, the gradients of the component functions  $f_1, \dots, f_n$  will be referred to as *stochastic gradients*. A practical issue with SGD is that consecutive stochastic gradients may vary a lot or even point in opposite directions. This slows down the performance of SGD. On balance, however, SGD is preferable to GD in applications where low accuracy solutions are sufficient. In such cases usually only a small number of passes through the data (i.e., work equivalent to a small number of full gradient evaluations) are needed to find an acceptable  $x$ . For this reason, SGD is extremely popular in fields such as machine learning.

In order to improve upon GD, one needs to reduce the cost of computing a gradient. In order to improve upon SGD, one has to reduce the variance of the stochastic gradients. In this paper we propose and analyze a *Semi-Stochastic Gradient Descent* (S2GD) method. Our method combines GD and SGD steps and reaps the benefits of both algorithms: it inherits the stability and speed of GD and at the same time retains the work-efficiency of SGD.

## 1.2. Brief Literature Review

Several recent papers, e.g., Richtárik and Takáč [3], Roux et al. [4], Schmidt et al. [5], Shalev-Shwartz and Zhang [6], and Johnson and Zhang [7] proposed methods which achieve similar variance-reduction effect, directly or indirectly. These methods enjoy linear convergence rates when applied to minimizing smooth strongly convex loss functions.

The method in Richtárik and Takáč [3] is known as Random Coordinate Descent for Composite functions (RCDC), and can be either applied directly to Equation (1), or to a dual version of Equation (1). Unless specific conditions

on the problem structure are met, application to the primal directly is not as computationally efficient as its dual version<sup>1</sup>. Application of a coordinate descent method to the dual formulation of Equation (1) is generally referred to as Stochastic Dual Coordinate Ascent (SDCA) [9]. The algorithm in Shalev-Shwartz [6] exhibits this duality, and the method in Takáč et al. [10] extends the primal-dual framework to the parallel/mini-batch setting. Parallel and distributed stochastic coordinate descent methods were studied in Richtárik and Takáč [11], Fercoq and Richtárik [12], and Richtárik and Takáč [13].

Stochastic Average Gradient (SAG) by Roux et al. [4], is one of the first SGD-type methods, other than coordinate descent methods, which were shown to exhibit linear convergence. The method of Johnson and Zhang [7], called Stochastic Variance Reduced Gradient (SVRG), arises as a special case in our setting for a suboptimal choice of a single parameter of our method. The Epoch Mixed Gradient Descent (EMGD) method, Zhang et al. [14], is similar in spirit to SVRG, but achieves a quadratic dependence on the condition number instead of a linear dependence, as is the case with SDCA, SAG, SVRG and with our method.

Earlier works of Friedlander and Schmidt [15], Deng and Ferris [16], and Bastin et al. [17] attempt to interpolate between GD and SGD and decrease variance by varying the sample size. These methods however do not realize the kind of improvements as the recent methods above. For partially related classical work on semi-stochastic approximation methods we refer<sup>2</sup> the reader to the papers of Marti and Fuchs [18, 19], which focus on general stochastic optimization.

## 1.3. Outline

We start in Section 2 by describing two algorithms: S2GD, which we analyze, and S2GD+, which we do not analyze, but which exhibits superior performance in practice. We then move to summarizing some of the main contributions of this paper in Section 3. Section 4 is devoted to establishing expectation and high probability complexity results for S2GD in the case of a strongly convex loss. The results are generic in that the parameters of the method are set arbitrarily. Hence, in Section 5 we study the problem of choosing the parameters optimally, with the goal of minimizing the total workload (# of processed examples) sufficient to produce a result of specified accuracy. In Section 6 we establish high probability complexity bounds for S2GD applied to a non-strongly convex loss function. Discussion of efficient implementation for sparse data is in Section 7. Finally, in Section 8 we perform very encouraging numerical experiments on real and artificial problem instances. A brief conclusion can be found in Section 9.

<sup>1</sup>The question of whether or when primal or dual version is better has recently been studied in Csiba and Richtárik [8] to which we refer the reader for further details.

<sup>2</sup>We thank Zaid Harchaoui who pointed us to these papers a few days before we posted our work to arXiv.

## 2. SEMI-STOCHASTIC GRADIENT DESCENT

In this section we describe two novel algorithms: S2GD and S2GD+. We analyze the former only. The latter, however, has superior convergence properties in our experiments.

These following two assumption are regarded as basic setting for smooth convex optimization, under which analysis of methods is typically presented first<sup>3</sup>. We assume throughout the paper that the functions  $f_i$  are convex and  $L$ -smooth.

**Assumption 1.** *The functions  $f_1, \dots, f_n$  have Lipschitz continuous gradients with constant  $L > 0$  (in other words, they are  $L$ -smooth). That is, for all  $x, z \in \mathbb{R}^d$  and all  $i = 1, 2, \dots, n$ ,*

$$f_i(z) \leq f_i(x) + \langle f'_i(x), z - x \rangle + \frac{L}{2} \|z - x\|^2.$$

(This implies that the gradient of  $f$  is Lipschitz with constant  $L$ , and hence  $f$  satisfies the same inequality.)

In one part of the paper (Section 4) we also make the following additional assumption:

**Assumption 2.** *The average loss  $f$  is  $\mu$ -strongly convex,  $\mu > 0$ . That is, for all  $x, z \in \mathbb{R}^d$ ,*

$$f(z) \geq f(x) + \langle f'(x), z - x \rangle + \frac{\mu}{2} \|z - x\|^2. \quad (3)$$

(Note that, necessarily,  $\mu \leq L$ .)

### 2.1. S2GD

Algorithm 1 (S2GD) depends on three parameters: stepsize  $h$ , constant  $m$  limiting the number of stochastic gradients computed in a single epoch, and a  $\nu \in [0, \mu]$ , where  $\mu$  is the strong convexity constant of  $f$ . In practice,  $\nu$  would be a known lower bound on  $\mu$ . Note that the algorithm works also without any knowledge of the strong convexity parameter—the case of  $\nu = 0$ .

---

#### Algorithm 1 Semi-Stochastic Gradient Descent (S2GD)

---

**parameters:**  $m = \max \#$  of stochastic steps per epoch,  $h =$  stepsize,  $\nu =$  lower bound on  $\mu$

```

for  $j = 0, 1, 2, \dots$  do
   $g_j \leftarrow \frac{1}{n} \sum_{i=1}^n f'_i(x_j)$ 
   $y_{j,0} \leftarrow x_j$ 
  Let  $t_j \leftarrow t$  with probability  $(1 - \nu h)^{m-t} / \beta$  for  $t =$ 
   $1, 2, \dots, m$ 
  for  $t = 0$  to  $t_j - 1$  do
    Pick  $i \in \{1, 2, \dots, n\}$ , uniformly at random
     $y_{j,t+1} \leftarrow y_{j,t} - h (g_j + f'_i(y_{j,t}) - f'_i(x_j))$ 
  end for
   $x_{j+1} \leftarrow y_{j,t_j}$ 
end for

```

---

<sup>3</sup>Since the first version of our work, our proposed algorithm has been extended to apply to a more broader class functions in Konečný et al. [20].

The method has an outer loop, indexed by epoch counter  $j$ , and an inner loop, indexed by  $t$ . In each epoch  $j$ , the method first computes  $g_j$ —the full gradient of  $f$  at  $x_j$ . Subsequently, the method produces a random number  $t_j \in [1, m]$  of steps, following a geometric law, where

$$\beta \stackrel{\text{def}}{=} \sum_{t=1}^m (1 - \nu h)^{m-t}, \quad (4)$$

with only two stochastic gradients computed in each step<sup>4</sup>. For each  $t = 0, \dots, t_j - 1$ , the stochastic gradient  $f'_i(x_j)$  is subtracted from  $g_j$ , and  $f'_i(y_{j,t-1})$  is added to  $g_j$ , which ensures that, one has

$$\mathbf{E}(g_j + f'_i(y_{j,t}) - f'_i(x_j)) = f'(y_{j,t}),$$

where the expectation is with respect to the random variable  $i$ .

Hence, the algorithm is stochastic gradient descent—albeit executed in a nonstandard way (compared to the traditional implementation described in the introduction).

Note that for all  $j$ , the expected number of iterations of the inner loop,  $\mathbf{E}(t_j)$ , is equal to

$$\xi = \xi(m, h) \stackrel{\text{def}}{=} \sum_{t=1}^m t \frac{(1 - \nu h)^{m-t}}{\beta}. \quad (5)$$

Also note that  $\xi \in [\frac{m+1}{2}, m)$ , with the lower bound attained for  $\nu = 0$ , and the upper bound for  $\nu h \rightarrow 1$ .

### 2.2. S2GD+

We also implement Algorithm 2, which we call S2GD+. In our experiments, the performance of this method is superior to all methods we tested, including S2GD. However, we do not analyze the complexity of this method and leave this as an open problem.

---

#### Algorithm 2 S2GD+

---

**parameters:**  $\alpha \geq 1$  (e.g.,  $\alpha = 1$ )

- Run SGD for a single pass over the data (i.e.,  $n$  iterations); output  $x$
  - Starting from  $x_0 = x$ , run a version of S2GD in which  $t_j = \alpha n$  for all  $j$
- 

In brief, S2GD+ starts by running SGD for 1 epoch (1 pass over the data) and then switches to a variant of S2GD in which the number of the inner iterations,  $t_j$ , is not random, but fixed to be  $n$  or a small multiple of  $n$ .

The motivation for this method is the following. It is common knowledge that SGD is able to progress much more in one pass over the data than GD (where this would correspond to a single gradient step). However, the very first step of S2GD is the computation of the full gradient of  $f$ . Hence, by starting with a

<sup>4</sup>It is possible to get away with computing only a single stochastic gradient per inner iteration, namely  $f'_i(y_{j,t})$ , at the cost of having to store in memory  $f'_i(x_j)$  for  $i = 1, 2, \dots, n$ . This, however, can be impractical for big  $n$ .

**TABLE 1 | Comparison of performance of selected methods suitable for solving Equation (1).**

Algorithm	Complexity/work
Nesterov's algorithm	$O(\sqrt{\kappa n} \log(1/\varepsilon))$
EMGD	$O((n + \kappa^2) \log(1/\varepsilon))$
SAG	$O(\max\{n, \kappa\} \log(1/\varepsilon))$
SDCA	$O((n + \kappa) \log(1/\varepsilon))$
SVRG	$O((n + \kappa) \log(1/\varepsilon))$
S2GD	$O((n + \kappa) \log(1/\varepsilon))$

The complexity/work is measured in the number of stochastic gradient evaluations needed to find an  $\varepsilon$ -solution.

single pass over data using SGD and *then* switching to S2GD, we obtain a superior method in practice<sup>5</sup>.

### 3. SUMMARY OF RESULTS

In this section we summarize some of the main results and contributions of this work.

- Complexity for strongly convex  $f$ .** If  $f$  is strongly convex, S2GD needs

$$\mathcal{W} = O((n + \kappa) \log(1/\varepsilon)) \quad (6)$$

work (measured as the total number of evaluations of the stochastic gradient, accounting for the full gradient evaluations as well) to output an  $\varepsilon$ -approximate solution (in expectation or in high probability), where  $\kappa = L/\mu$  is the condition number. This is achieved by running S2GD with stepsize  $h = O(1/L)$ ,  $j = O(\log(1/\varepsilon))$  epochs (this is also equal to the number of full gradient evaluations) and  $m = O(\kappa)$  (this is also roughly equal to the number of stochastic gradient evaluations in a single epoch). The complexity results are stated in detail in Sections 4 and 5 (see Theorems 4, 5 and 6; see also Equations 26 and 27).

- Comparison with existing results.** This complexity result (Equation 6) matches the best-known results obtained for strongly convex losses in recent work such as Roux et al. [4], Johnson and Zhang [7], and Zhang and Mahdavi [14]. Our treatment is most closely related to Johnson and Zhang [7], and contains their method (SVRG) as a special case. In **Table 1** we summarize our results in the strongly convex case with other existing results for different algorithms.

We should note that the rate of convergence of Nesterov's algorithm [21] is a deterministic result. EMGD and S2GD results hold with high probability (see Theorem 5 for precise statement). Complexity results for stochastic coordinate descent methods are also typically analyzed in the high probability regime [3]. The remaining results hold in

<sup>5</sup>Using a single pass of SGD as an initialization strategy was already considered in Roux et al. [4]. However, the authors claim that their implementation of vanilla SAG did not benefit from it. S2GD does benefit from such an initialization due to it starting, in theory, with a (heavy) full gradient computation.

expectation. Notion of  $\kappa$  is slightly different for SDCA, which requires explicit knowledge of the strong convexity parameter  $\mu$  to run the algorithm. In contrast, other methods do not algorithmically depend on this, and thus their convergence rate can adapt to any additional strong convexity locally.

- Complexity for convex  $f$ .** If  $f$  is *not* strongly convex, then we propose that S2GD be applied to a perturbed version of the problem, with strong convexity constant  $\mu = O(L/\varepsilon)$ . An  $\varepsilon$ -accurate solution of the original problem is recovered with arbitrarily high probability (see Theorem 8 in Section 6). The total work in this case is

$$\mathcal{W} = O((n + L/\varepsilon) \log(1/\varepsilon)),$$

that is,  $\tilde{O}(1/\varepsilon)$ , which is better than the standard rate of SGD.

- Optimal parameters.** We derive formulas for optimal parameters of the method which (approximately) minimize the total workload, measured in the number of stochastic gradients computed (counting a single full gradient evaluation as  $n$  evaluations of the stochastic gradient). In particular, we show that the method should be run for  $O(\log(1/\varepsilon))$  epochs, with stepsize  $h = O(1/L)$  and  $m = O(\kappa)$ . No such results were derived for SVRG in Johnson and Zhang [7].
- One epoch.** Consider the case when S2GD is run for 1 epoch only, effectively limiting the number of full gradient evaluations to 1, while choosing a target accuracy  $\varepsilon$ . We show that S2GD with  $\nu = \mu$  needs

$$O(n + (\kappa/\varepsilon) \log(1/\varepsilon))$$

work only (see **Table 2**). This compares favorably with the optimal complexity in the  $\nu = 0$  case (which reduces to SVRG), where the work needed is

$$O(n + \kappa/\varepsilon^2).$$

For two epochs one could just say that we need  $\sqrt{\varepsilon}$  decrease in each epoch, thus having complexity of  $O(n + (\kappa/\sqrt{\varepsilon}) \log(1/\sqrt{\varepsilon}))$ . This is already better than general rate of SGD ( $O(1/\varepsilon)$ ).

- Special cases.** GD and SVRG arise as special cases of S2GD, for  $m = 1$  and  $\nu = 0$ , respectively<sup>6</sup>.
- Low memory requirements.** Note that SDCA and SAG, unlike SVRG and S2GD, need to store all gradients  $f'_i$  (or dual variables) throughout the iterative process. While this may not be a problem for a modest sized optimization task, this requirement makes such methods less suitable for problems with very large  $n$ .
- S2GD+.** We propose a “boosted” version of S2GD, called S2GD+, which we do not analyze. In our experiments, however, it performs vastly superior to all other methods we tested, including GD, SGD, SAG and S2GD. S2GD alone is better than both GD and SGD if a highly accurate solution

<sup>6</sup>While S2GD reduces to GD for  $m = 1$ , our *analysis* does not say anything meaningful in the  $m = 1$  case—it is too coarse to cover this case. This is also the reason behind the empty space in the “Complexity” box column for GD in **Table 2**.

**TABLE 2 | Summary of complexity results and special cases.**

Parameters	Method	Complexity
$\nu = \mu, j = O(\log(\frac{1}{\epsilon}))$ & $m = O(\kappa)$	Optimal S2GD	$O((n + \kappa) \log(\frac{1}{\epsilon}))$
$m = 1$	GD	—
$\nu = 0$	SVRG [7]	$O((n + \kappa) \log(\frac{1}{\epsilon}))$
$\nu = 0, j = 1, m = O(\frac{\kappa}{\epsilon^2})$	Optimal SVRG with 1 epoch	$O(n + \frac{\kappa}{\epsilon^2})$
$\nu = \mu, j = 1, m = O(\frac{\kappa}{\epsilon} \log(\frac{1}{\epsilon}))$	Optimal S2GD with 1 epoch	$O(n + \frac{\kappa}{\epsilon} \log(\frac{1}{\epsilon}))$

Condition number:  $\kappa = L/\mu$  if  $f$  is  $\mu$ -strongly convex and  $\kappa = 2L/\epsilon$  if  $f$  is not strongly convex and  $\epsilon \leq L$ .

is required. The performance of S2GD and SAG is roughly comparable, even though in our experiments S2GD turned to have an edge.

### 4. COMPLEXITY ANALYSIS: STRONGLY CONVEX LOSS

For the purpose of the analysis, let

$$\mathcal{F}_{j,t} \stackrel{\text{def}}{=} \sigma(x_1, x_2, \dots, x_j; y_{j,1}, y_{j,2}, \dots, y_{j,t}) \tag{7}$$

be the  $\sigma$ -algebra generated by the relevant history of S2GD. We first isolate an auxiliary result.

**Lemma 3.** Consider the S2GD algorithm. For any fixed epoch number  $j$ , the following identity holds:

$$\mathbf{E}(f(x_{j+1})) = \frac{1}{\beta} \sum_{t=1}^m (1 - \nu h)^{m-t} \mathbf{E}(f(y_{j,t-1})). \tag{8}$$

*Proof.* By the tower law of conditional expectations and the definition of  $x_{j+1}$  in the algorithm, we obtain

$$\begin{aligned} \mathbf{E}(f(x_{j+1})) &= \mathbf{E}(\mathbf{E}(f(x_{j+1}) \mid \mathcal{F}_{j,m})) \\ &= \mathbf{E}\left(\sum_{t=1}^m \frac{(1 - \nu h)^{m-t}}{\beta} f(y_{j,t-1})\right) \\ &= \frac{1}{\beta} \sum_{t=1}^m (1 - \nu h)^{m-t} \mathbf{E}(f(y_{j,t-1})). \end{aligned}$$

□

We now state and prove the main result of this section.

**Theorem 4.** Let Assumptions 1 and 2 be satisfied. Consider the S2GD algorithm applied to solving problem (1). Choose  $0 \leq \nu \leq \mu$ ,  $0 < h < \frac{1}{2L}$ , and let  $m$  be sufficiently large so that

$$c \stackrel{\text{def}}{=} \frac{(1 - \nu h)^m}{\beta \mu h (1 - 2Lh)} + \frac{2(L - \mu)h}{1 - 2Lh} < 1. \tag{9}$$

Then we have the following convergence in expectation:

$$\mathbf{E}(f(x_j) - f(x_*)) \leq c^j (f(x_0) - f(x_*)). \tag{10}$$

Before we proceed to proving the theorem, note that in the special case with  $\nu = 0$ , we recover the result of Johnson and Zhang [7] (with a minor improvement in the second term of  $c$  where  $L$  is replaced by  $L - \mu$ ), namely

$$c = \frac{1}{\mu h (1 - 2Lh)m} + \frac{2(L - \mu)h}{1 - 2Lh}. \tag{11}$$

If we set  $\nu = \mu$ , then  $c$  can be written in the form (see Equation 4)

$$c = \frac{(1 - \mu h)^m}{(1 - (1 - \mu h)^m)(1 - 2Lh)} + \frac{2(L - \mu)h}{1 - 2Lh}. \tag{12}$$

Clearly, the latter  $c$  is a major improvement on the former one. We shall elaborate on this further later.

*Proof.* It is well-known [21, Theorem 2.1.5] that since the functions  $f_i$  are  $L$ -smooth, they necessarily satisfy the following inequality:

$$\|f'_i(x) - f'_i(x_*)\|^2 \leq 2L [f_i(x) - f_i(x_*) - \langle f'_i(x_*), x - x_* \rangle].$$

By summing these inequalities for  $i = 1, \dots, n$ , and using  $f'(x_*) = 0$ , we get

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \|f'_i(x) - f'_i(x_*)\|^2 &\leq 2L [f(x) - f(x_*) - \langle f'(x_*), x - x_* \rangle] \\ &= 2L(f(x) - f(x_*)). \end{aligned} \tag{13}$$

Let  $G_{j,t} \stackrel{\text{def}}{=} g_j + f'_i(y_{j,t-1}) - f'_i(x_j)$  be the direction of update at  $j^{\text{th}}$  iteration in the outer loop and  $t^{\text{th}}$  iteration in the inner loop. Taking expectation with respect to  $i$ , conditioned on the  $\sigma$ -algebra  $\mathcal{F}_{j,t-1}$  Equation (7), we obtain<sup>7</sup>

$$\begin{aligned} \mathbf{E}(\|G_{j,t}\|^2) &= \mathbf{E}(\|f'_i(y_{j,t-1}) - f'_i(x_*) - f'_i(x_j) + f'_i(x_*) + g_j\|^2) \\ &\leq 2\mathbf{E}(\|f'_i(y_{j,t-1}) - f'_i(x_*)\|^2) \\ &\quad + 2\mathbf{E}(\| [f'_i(x_j) - f'_i(x_*)] - f'(x_j) \|^2) \\ &= 2\mathbf{E}(\|f'_i(y_{j,t-1}) - f'_i(x_*)\|^2) \\ &\quad + 2\mathbf{E}(\|f'_i(x_j) - f'_i(x_*)\|^2) \\ &\quad - 4\mathbf{E}(\langle f'(x_j), f'_i(x_j) - f'_i(x_*) \rangle) + 2\|f'(x_j)\|^2 \\ &\stackrel{(13)}{\leq} 4L [f(y_{j,t-1}) - f(x_*) + f(x_j) - f(x_*)] \\ &\quad - 2\|f'(x_j)\|^2 - 4\langle f'(x_j), f'(x_*) \rangle \\ &\stackrel{(3)}{\leq} 4L [f(y_{j,t-1}) - f(x_*)] \\ &\quad + 4(L - \mu) [f(x_j) - f(x_*)]. \end{aligned} \tag{14}$$

Above we have used the bound  $\|x' + x''\|^2 \leq 2\|x'\|^2 + 2\|x''\|^2$  and the fact that

$$\mathbf{E}(G_{j,t} \mid \mathcal{F}_{j,t-1}) = f'(y_{j,t-1}). \tag{15}$$

<sup>7</sup>For simplicity, we suppress the  $\mathbf{E}(\cdot \mid \mathcal{F}_{j,t-1})$  notation here.

We now study the expected distance to the optimal solution (a standard approach in the analysis of gradient methods):

**Theorem 5.** Consider the setting of Theorem 4. Then, for any  $0 < \rho < 1, 0 < \varepsilon < 1$  and

$$\begin{aligned}
 \mathbf{E}(\|y_{j,t} - x_*\|^2 \mid \mathcal{F}_{j,t-1}) &= \|y_{j,t-1} - x_*\|^2 - 2h\mathbf{E}(G_{j,t} \mid \mathcal{F}_{j,t-1}, y_{j,t-1} - x_*) \\
 &\quad + h^2\mathbf{E}(\|G_{j,t}\|^2 \mid \mathcal{F}_{j,t-1}) \\
 &\stackrel{(14)+(15)}{\leq} \|y_{j,t-1} - x_*\|^2 - 2h\langle f'(y_{j,t-1}), y_{j,t-1} - x_* \rangle \\
 &\quad + 4Lh^2 [f(y_{j,t-1}) - f(x_*)] + 4(L - \mu)h^2 [f(x_j) - f(x_*)] \\
 &\stackrel{(3)}{\leq} \|y_{j,t-1} - x_*\|^2 - 2h [f(y_{j,t-1}) - f(x_*)] - \nu h \|y_{j,t-1} - x_*\|^2 \\
 &\quad + 4Lh^2 [f(y_{j,t-1}) - f(x_*)] + 4(L - \mu)h^2 [f(x_j) - f(x_*)] \\
 &= (1 - \nu h)\|y_{j,t-1} - x_*\|^2 - 2h(1 - 2Lh)[f(y_{j,t-1}) - f(x_*)] \\
 &\quad + 4(L - \mu)h^2 [f(x_j) - f(x_*)]. \tag{16}
 \end{aligned}$$

By rearranging the terms in Equation (16) and taking expectation over the  $\sigma$ -algebra  $\mathcal{F}_{j,t-1}$ , we get the following inequality:

$$j \geq \frac{\log\left(\frac{1}{\varepsilon\rho}\right)}{\log\left(\frac{1}{c}\right)}, \tag{18}$$

$$\begin{aligned}
 \mathbf{E}(\|y_{j,t} - x_*\|^2) + 2h(1 - 2Lh)\mathbf{E}(f(y_{j,t-1}) - f(x_*)) \\
 \leq (1 - \nu h)\mathbf{E}(\|y_{j,t-1} - x_*\|^2) \\
 + 4(L - \mu)h^2\mathbf{E}(f(x_j) - f(x_*)). \tag{17}
 \end{aligned}$$

we have

$$\mathbf{P}\left(\frac{f(x_j) - f(x_*)}{f(x_0) - f(x_*)} \leq \varepsilon\right) \geq 1 - \rho. \tag{19}$$

Finally, we can analyze what happens after one iteration of the outer loop of S2GD, i.e., between two computations of the full gradient. By summing up inequalities Equation (17) for  $t = 1, \dots, m$ , with inequality  $t$  multiplied by  $(1 - \nu h)^{m-t}$ , we get the left-hand side

*Proof.* This follows directly from Markov inequality and Theorem 4:

$$\mathbf{P}(f(x_j) - f(x_*) > \varepsilon(f(x_0) - f(x_*))) \stackrel{(10)}{\leq} \frac{\mathbf{E}(f(x_j) - f(x_*))}{\varepsilon(f(x_0) - f(x_*))} \leq \frac{c^j}{\varepsilon} \stackrel{(18)}{\leq} \rho \quad \square$$

$$\begin{aligned}
 LHS &= \mathbf{E}(\|y_{j,m} - x_*\|^2) + 2h(1 - 2Lh) \\
 &\quad \sum_{t=1}^m (1 - \nu h)^{m-t} \mathbf{E}(f(y_{j,t-1}) - f(x_*)) \\
 &\stackrel{(8)}{=} \mathbf{E}(\|y_{j,m} - x_*\|^2) + 2\beta h(1 - 2Lh)\mathbf{E}(f(x_{j+1}) - f(x_*)),
 \end{aligned}$$

This result will be also useful when treating the non-strongly convex case.

and the right-hand side

$$\begin{aligned}
 RHS &= (1 - \nu h)^m \mathbf{E}(\|x_j - x_*\|^2) \\
 &\quad + 4\beta(L - \mu)h^2 \mathbf{E}(f(x_j) - f(x_*)) \\
 &\stackrel{(3)}{\leq} \frac{2(1 - \nu h)^m}{\mu} \mathbf{E}(f(x_j) - f(x_*)) \\
 &\quad + 4\beta(L - \mu)h^2 \mathbf{E}(f(x_j) - f(x_*)) \\
 &= 2 \left( \frac{(1 - \nu h)^m}{\mu} + 2\beta(L - \mu)h^2 \right) \mathbf{E}(f(x_j) - f(x_*)).
 \end{aligned}$$

Since  $LHS \leq RHS$ , we finally conclude with

$$\begin{aligned}
 \mathbf{E}(f(x_{j+1}) - f(x_*)) &\leq c\mathbf{E}(f(x_j) - f(x_*)) \\
 - \frac{\mathbf{E}(\|y_{j,m} - x_*\|^2)}{2\beta h(1 - 2Lh)} &\leq c\mathbf{E}(f(x_j) - f(x_*)).
 \end{aligned}$$

□

Since we have established linear convergence of expected values, a high probability result can be obtained in a straightforward way using Markov inequality.

## 5. OPTIMAL CHOICE OF PARAMETERS

The goal of this section is to provide insight into the choice of parameters of S2GD; that is, the number of epochs (equivalently, full gradient evaluations)  $j$ , the maximal number of steps in each epoch  $m$ , and the stepsize  $h$ . The remaining parameters ( $L, \mu, n$ ) are inherent in the problem and we will hence treat them in this section as given.

In particular, ideally we wish to find parameters  $j, m$  and  $h$  solving the following optimization problem:

$$\min_{j,m,h} \tilde{\mathcal{W}}(j, m, h) \stackrel{\text{def}}{=} j(n + 2\xi(m, h)), \tag{20}$$

subject to

$$\mathbf{E}(f(x_j) - f(x_*)) \leq \varepsilon(f(x_0) - f(x_*)). \tag{21}$$

Note that  $\tilde{\mathcal{W}}(j, m, h)$  is the *expected work*, measured by the number number of stochastic gradient evaluations, performed by S2GD when running for  $j$  epochs. Indeed, the evaluation of  $g_j$  is

equivalent to  $n$  stochastic gradient evaluations, and each epoch further computes on average  $2\xi(m, h)$  stochastic gradients (see Equation 5). Since  $\frac{m+1}{2} \leq \xi(m, h) < m$ , we can simplify and solve the problem with  $\xi$  set to the conservative upper estimate  $\xi = m$ .

In view of Equation (10), accuracy constraint Equation (21) is satisfied if  $c$  (which depends on  $h$  and  $m$ ) and  $j$  satisfy

$$c^j \leq \varepsilon. \tag{22}$$

We therefore instead consider the parameter fine-tuning problem:

$$\min_{j,m,h} \mathcal{W}(j, m, h) \stackrel{\text{def}}{=} j(n + 2m) \quad \text{subject to} \quad c \leq \varepsilon^{1/j}. \tag{23}$$

In the following we (approximately) solve this problem in two steps. First, we fix  $j$  and find (nearly) optimal  $h = h(j)$  and  $m = m(j)$ . The problem reduces to minimizing  $m$  subject to  $c \leq \varepsilon^{1/j}$  by fine-tuning  $h$ . While in the  $\nu = 0$  case it is possible to obtain closed form solution, this is not possible for  $\nu > \mu$ .

However, it is still possible to obtain a good formula for  $h(j)$  leading to expression for good  $m(j)$  which depends on  $\varepsilon$  in the correct way. We then plug the formula for  $m(j)$  obtained this way back into Equation (23), and study the quantity  $\mathcal{W}(j, m(j), h(j)) = j(n + 2m(j))$  as a function of  $j$ , over which we optimize at the end.

**Theorem 6 (Choice of parameters).** Fix the number of epochs  $j \geq 1$ , error tolerance  $0 < \varepsilon < 1$ , and let  $\Delta = \varepsilon^{1/j}$ . If we run S2GD with the stepsize

$$h = h(j) \stackrel{\text{def}}{=} \frac{1}{\frac{4}{\Delta}(L - \mu) + 2L} \tag{24}$$

and

$$m \geq m(j) \stackrel{\text{def}}{=} \begin{cases} \left( \frac{4(\kappa-1)}{\Delta} + 2\kappa \right) \log \left( \frac{2}{\Delta} + \frac{2\kappa-1}{\kappa-1} \right), & \text{if } \nu = \mu, \\ \frac{8(\kappa-1)}{\Delta^2} + \frac{8\kappa}{\Delta} + \frac{2\kappa^2}{\kappa-1}, & \text{if } \nu = 0, \end{cases} \tag{25}$$

then  $\mathbf{E}(f(x_j) - f(x_*)) \leq \varepsilon(f(x_0) - f(x_*))$ .

In particular, if we choose  $j^* = \lceil \log(1/\varepsilon) \rceil$ , then  $\frac{1}{\Delta} \leq \exp(1)$ , and hence  $m(j^*) = O(\kappa)$ , leading to the workload

$$\mathcal{W}(j^*, m(j^*), h(j^*)) = \lceil \log \left( \frac{1}{\varepsilon} \right) \rceil (n + O(\kappa)) = O \left( (n + \kappa) \log \left( \frac{1}{\varepsilon} \right) \right). \tag{26}$$

*Proof.* We only need to show that  $c \leq \Delta$ , where  $c$  is given by Equation (12) for  $\nu = \mu$  and by Equation (11) for  $\nu = 0$ . We denote the two summands in expressions for  $c$  as  $c_1$  and  $c_2$ . We choose the  $h$  and  $m$  so that both  $c_1$  and  $c_2$  are smaller than  $\Delta/2$ , resulting in  $c_1 + c_2 = c \leq \Delta$ .

The stepsize  $h$  is chosen so that

$$c_2 \stackrel{\text{def}}{=} \frac{2(L - \mu)h}{1 - 2Lh} = \frac{\Delta}{2},$$

and hence it only remains to verify that  $c_1 = c - c_2 \leq \frac{\Delta}{2}$ . In the  $\nu = 0$  case,  $m(j)$  is chosen so that  $c - c_2 = \frac{\Delta}{2}$ . In the  $\nu = \mu$  case,

$c - c_2 = \frac{\Delta}{2}$  holds for  $m = \log \left( \frac{2}{\Delta} + \frac{2\kappa-1}{\kappa-1} \right) / \log \left( \frac{1}{1-H} \right)$ , where  $H = \left( \frac{4(\kappa-1)}{\Delta} + 2\kappa \right)^{-1}$ . We only need to observe that  $c$  decreases as  $m$  increases, and apply the inequality  $\log \left( \frac{1}{1-H} \right) \geq H$ . □

We now comment on the above result:

1. **Workload.** Notice that for the choice of parameters  $j^*$ ,  $h = h(j^*)$ ,  $m = m(j^*)$  and any  $\nu \in [0, \mu]$ , the method needs  $\log(1/\varepsilon)$  computations of the full gradient (note this is independent of  $\kappa$ ), and  $O(\kappa \log(1/\varepsilon))$  computations of the stochastic gradient. This result, and special cases thereof, are summarized in **Table 2**.
2. **Simpler formulas for  $m$ .** If  $\kappa \geq 2$ , we can instead of Equation (25) use the following (slightly worse but) simpler expressions for  $m(j)$ , obtained from Equation (25) by using the bounds  $1 \leq \kappa - 1, \kappa - 1 \leq \kappa$  and  $\Delta < 1$  in appropriate places (e.g.,  $\frac{8\kappa}{\Delta} < \frac{8\kappa}{\Delta^2}, \frac{\kappa}{\kappa-1} \leq 2 < \frac{2}{\Delta^2}$ ):
 
$$m \geq \tilde{m}(j) \stackrel{\text{def}}{=} \begin{cases} \frac{6\kappa}{\Delta} \log \left( \frac{5}{\Delta} \right), & \text{if } \nu = \mu, \\ \frac{20\kappa}{\Delta^2}, & \text{if } \nu = 0. \end{cases} \tag{27}$$
3. **Optimal stepsize in the  $\nu = 0$  case.** Theorem 6 does not claim to have solved problem (23); the problem in general does not have a closed form solution. However, in the  $\nu = 0$  case a closed-form formula can easily be obtained:

$$h(j) = \frac{1}{\frac{4}{\Delta}(L - \mu) + 4L}, \quad m \geq m(j) \stackrel{\text{def}}{=} \frac{8(\kappa - 1)}{\Delta^2} + \frac{8\kappa}{\Delta}. \tag{28}$$

Indeed, for fixed  $j$ , Equation (23) is equivalent to finding  $h$  that minimizes  $m$  subject to the constraint  $c \leq \Delta$ . In view of Equation (11), this is equivalent to searching for  $h > 0$  maximizing the quadratic  $h \rightarrow h(\Delta - 2(\Delta L + L - \mu)h)$ , which leads to Equation (28).

Note that both the stepsize  $h(j)$  and the resulting  $m(j)$  are slightly larger in Theorem 6 than in Equation (28). This is because in the theorem the stepsize was for simplicity chosen to satisfy  $c_2 = \frac{\Delta}{2}$ , and hence is (slightly) suboptimal. Nevertheless, the dependence of  $m(j)$  on  $\Delta$  is of the correct (optimal) order in both cases. That is,  $m(j) = O \left( \frac{\kappa}{\Delta} \log \left( \frac{1}{\Delta} \right) \right)$  for  $\nu = \mu$  and  $m(j) = O \left( \frac{\kappa}{\Delta^2} \right)$  for  $\nu = 0$ .

4. **Stepsize choice.** In cases when one does not have a good estimate of the strong convexity constant  $\mu$  to determine the stepsize via Equation (24), one may choose suboptimal stepsize that does not depend on  $\mu$  and derive similar results to those above. For instance, one may choose  $h = \frac{\Delta}{6L}$ .

In **Table 3** we provide comparison of work needed for small values of  $j$ , and different values of  $\kappa$  and  $\varepsilon$ . Note, for instance, that for any problem with  $n = 10^9$  and  $\kappa = 10^3$ , S2GD outputs a highly accurate solution ( $\varepsilon = 10^{-6}$ ) in the amount of work equivalent to 2.12 evaluations of the full gradient of  $f$ !

**TABLE 3 | Comparison of work sufficient to solve a problem with  $n = 10^9$ , and various values of  $\kappa$  and  $\epsilon$ .**

$\epsilon = 10^{-3}, \kappa = 10^3$			$\epsilon = 10^{-6}, \kappa = 10^3$			$\epsilon = 10^{-9}, \kappa = 10^3$		
$j$	$\mathcal{W}_\mu(\hat{j})$	$\mathcal{W}_0(\hat{j})$	$j$	$\mathcal{W}_\mu(\hat{j})$	$\mathcal{W}_0(\hat{j})$	$j$	$\mathcal{W}_\mu(\hat{j})$	$\mathcal{W}_0(\hat{j})$
1	<b>1.06n</b>	17.0n	1	116n	$10^7n$	2	7.58n	$10^4n$
2	2.00n	<b>2.03n</b>	2	<b>2.12n</b>	34.0n	3	<b>3.18n</b>	51.0n
3	3.00n	3.00n	3	3.01n	<b>3.48n</b>	4	4.03n	6.03n
4	4.00n	4.00n	4	4.00n	4.06n	5	5.01n	<b>5.32n</b>
5	5.00n	5.00n	5	5.00n	5.02n	6	6.00n	6.09n
$\epsilon = 10^{-3}, \kappa = 10^6$			$\epsilon = 10^{-6}, \kappa = 10^6$			$\epsilon = 10^{-9}, \kappa = 10^6$		
$j$	$\mathcal{W}_\mu(\hat{j})$	$\mathcal{W}_0(\hat{j})$	$j$	$\mathcal{W}_\mu(\hat{j})$	$\mathcal{W}_0(\hat{j})$	$j$	$\mathcal{W}_\mu(\hat{j})$	$\mathcal{W}_0(\hat{j})$
2	4.14n	35.0n	4	8.29n	70.0n	5	17.3n	328n
3	<b>3.77n</b>	8.29n	5	<b>7.30n</b>	26.3n	8	<b>10.9n</b>	32.5n
4	4.50n	<b>6.39n</b>	6	7.55n	16.5n	10	11.9n	21.4n
5	5.41n	6.60n	8	9.01n	<b>12.7n</b>	13	14.3n	<b>19.1n</b>
6	6.37n	7.28n	10	10.8n	13.2n	20	21.0n	23.5n
$\epsilon = 10^{-3}, \kappa = 10^9$			$\epsilon = 10^{-6}, \kappa = 10^9$			$\epsilon = 10^{-9}, \kappa = 10^9$		
$j$	$\mathcal{W}_\mu(\hat{j})$	$\mathcal{W}_0(\hat{j})$	$j$	$\mathcal{W}_\mu(\hat{j})$	$\mathcal{W}_0(\hat{j})$	$j$	$\mathcal{W}_\mu(\hat{j})$	$\mathcal{W}_0(\hat{j})$
6	378n	1,293n	13	737n	2,409n	15	1,251n	4,834n
8	<b>358n</b>	1,063n	16	<b>717n</b>	2,126n	24	<b>1,076n</b>	3,189n
11	376n	<b>1,002n</b>	19	727n	2,025n	30	1,102n	3,018n
15	426n	1,058n	22	752n	<b>2,005n</b>	32	1,119n	<b>3,008n</b>
20	501n	1,190n	30	852n	2,116n	40	1,210n	3,078n

The work was computed using formula (23), with  $m(\hat{j})$  as in Equation (27). The notation  $\mathcal{W}_v(\hat{j})$  indicates what parameter  $v$  was used, and optimal values are highlighted in bold.

### 6. COMPLEXITY ANALYSIS: CONVEX LOSS

If  $f$  is convex but not strongly convex, we define  $\hat{f}_i(x) \stackrel{\text{def}}{=} f_i(x) + \frac{\mu}{2} \|x - x_0\|^2$ , for small enough  $\mu > 0$  (we shall see below how the choice of  $\mu$  affects the results), and consider the perturbed problem

$$\min_{x \in \mathbb{R}^d} \hat{f}(x), \tag{29}$$

where

$$\hat{f}(x) \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \hat{f}_i(x) = f(x) + \frac{\mu}{2} \|x - x_0\|^2. \tag{30}$$

Note that  $\hat{f}$  is  $\mu$ -strongly convex and  $(L + \mu)$ -smooth. In particular, the theory developed in the previous section applies. We propose that S2GD be instead applied to the perturbed problem, and show that an approximate solution of Equation (29) is also an approximate solution of Equation (1) (we will assume that this problem has a minimizer).

Let  $\hat{x}_*$  be the (necessarily unique) solution of the perturbed problem (29). The following result describes an important connection between the original problem and the perturbed problem.

**Lemma 7.** *If  $\hat{x} \in \mathbb{R}^d$  satisfies  $\hat{f}(\hat{x}) \leq \hat{f}(\hat{x}_*) + \delta$ , where  $\delta > 0$ , then*

$$f(\hat{x}) \leq f(x_*) + \frac{\mu}{2} \|x_0 - x_*\|^2 + \delta.$$

*Proof.* The statement is almost identical to Lemma 9 in Richtárik and Takáč [3]; its proof follows the same steps with only minor adjustments.  $\square$

We are now ready to establish a complexity result for non-strongly convex losses.

**Theorem 8.** *Let Assumption 1 be satisfied. Choose  $\mu > 0$ ,  $0 \leq v \leq \mu$ , stepsize  $0 < h < \frac{1}{2(L+\mu)}$ , and let  $m$  be sufficiently large so that*

$$\hat{c} \stackrel{\text{def}}{=} \frac{(1 - vh)^m}{\beta\mu h(1 - 2(L + \mu)h)} + \frac{2Lh}{1 - 2(L + \mu)h} < 1. \tag{31}$$

*Pick  $x_0 \in \mathbb{R}^d$  and let  $\hat{x}_0 = x_0, \hat{x}_1, \dots, \hat{x}_j$  be the sequence of iterates produced by S2GD as applied to problem (29). Then, for any  $0 < \rho < 1$ ,  $0 < \epsilon < 1$  and*

$$j \geq \frac{\log(1/(\epsilon\rho))}{\log(1/\hat{c})}, \tag{32}$$

*we have*

$$\mathbf{P} \left( f(\hat{x}_j) - f(x_*) \leq \epsilon(f(x_0) - f(x_*)) + \frac{\mu}{2} \|x_0 - x_*\|^2 \right) \geq 1 - \rho. \tag{33}$$

*In particular, if we choose  $\mu = \epsilon < L$  and parameters  $j^*, h(j^*), m(j^*)$  as in Theorem 6, the amount of work performed by S2GD to guarantee Equation (33) is*

$$\mathcal{W}(j^*, h(j^*), m(j^*)) = O \left( (n + \frac{L}{\epsilon}) \log(\frac{1}{\epsilon}) \right),$$

which consists of  $O(\frac{1}{\epsilon})$  full gradient evaluations and  $O(\frac{L}{\epsilon} \log(\frac{1}{\epsilon}))$  stochastic gradient evaluations.

*Proof.* We first note that

$$\hat{f}(\hat{x}_0) - \hat{f}(\hat{x}_*) \stackrel{(30)}{=} f(\hat{x}_0) - \hat{f}(\hat{x}_*) \leq f(\hat{x}_0) - f(\hat{x}_*) \leq f(x_0) - f(x_*), \tag{34}$$

where the first inequality follows from  $f \leq \hat{f}$ , and the second one from optimality of  $x_*$ . Hence, by first applying Lemma 7 with  $\hat{x} = \hat{x}_j$  and  $\delta = \epsilon(f(x_0) - f(x_*))$ , and then Theorem 5, with  $c \leftarrow \hat{c}$ ,  $f \leftarrow \hat{f}$ ,  $x_0 \leftarrow \hat{x}_0$ ,  $x_* \leftarrow \hat{x}_*$ , we obtain

$$\begin{aligned} \mathbf{P}\left(f(\hat{x}_j) - f(x_*) \leq \delta + \frac{\mu}{2} \|x_0 - x_*\|^2\right) \\ \stackrel{(\text{Lemma 7})}{\geq} \mathbf{P}\left(\hat{f}(\hat{x}_j) - \hat{f}(\hat{x}_*) \leq \delta\right) \\ \stackrel{(34)}{\geq} \mathbf{P}\left(\frac{\hat{f}(\hat{x}_j) - \hat{f}(\hat{x}_*)}{\hat{f}(\hat{x}_0) - \hat{f}(\hat{x}_*)} \leq \epsilon\right) \stackrel{(19)}{\geq} 1 - \rho. \end{aligned}$$

The second statement follows directly from the second part of Theorem 6 and the fact that the condition number of the perturbed problem is  $\kappa = \frac{L+\epsilon}{\epsilon} \leq \frac{2L}{\epsilon}$ .  $\square$

## 7. IMPLEMENTATION FOR SPARSE DATA

In our sparse implementation of Algorithm 1, described in this section and formally stated as Algorithm 3, we make the following structural assumption:

*Assumption 9.* The loss functions arise as the composition of a univariate smooth loss function  $\phi_i$ , and an inner product with a data point/example  $a_i \in \mathbb{R}^d$ :

$$f_i(x) = \phi_i(a_i^T x), \quad i = 1, 2, \dots, n.$$

In this case,  $f'_i(x) = \phi'_i(a_i^T x) a_i$ .

This is the structure in many cases of interest, including linear or logistic regression.

A natural question one might want to ask is whether S2GD can be implemented efficiently for sparse data.

Let us first take a brief detour and look at SGD, which performs iterations of the type:

$$x_{j+1} \leftarrow x_j - h \phi'_i(a_i^T x) a_i. \tag{35}$$

Let  $\omega_i$  be the number of nonzero features in example  $a_i$ , i.e.,  $\omega_i \stackrel{\text{def}}{=} \|a_i\|_0 \leq d$ . Assuming that the computation of the derivative of the univariate function  $\phi_i$  takes  $O(1)$  amount of work, the computation of  $\nabla f_i(x)$  will take  $O(\omega_i)$  work. Hence, the update step Equation (35) will cost  $O(\omega_i)$ , too, which means the method can naturally speed up its iterations on sparse data.

The situation is not as simple with S2GD, which for loss functions of the type described in Assumption 9 performs inner iterations as follows:

$$y_{j,t+1} \leftarrow y_{j,t} - h \left( g_j + \phi'_i(a_i^T y_{j,t}) a_i - \phi'_i(a_i^T x_j) a_i \right). \tag{36}$$

Indeed, note that  $g_j = f'(x_j)$  is in general be fully dense even for sparse data  $\{a_i\}$ . As a consequence, the update in Equation (36) might be as costly as  $d$  operations, irrespective of the sparsity level  $\omega_i$  of the active example  $a_i$ . However, we can use the following “lazy/delayed” update trick. We split the update to the  $y$  vector into two parts: immediate, and delayed. Assume index  $i = i_t$  was chosen at inner iteration  $t$ . We immediately perform the update

$$\tilde{y}_{j,t+1} \leftarrow y_{j,t} - h \left( \phi'_{i_t}(a_{i_t}^T y_{j,t}) - \phi'_{i_t}(a_{i_t}^T x_j) \right) a_{i_t},$$

which costs  $O(a_{i_t})$ . Note that we have not computed the  $y_{j,t+1}$ . However, we “know” that

$$y_{j,t+1} = \tilde{y}_{j,t+1} - h g_j,$$

without having to actually compute the difference. At the next iteration, we are supposed to perform update Equation (36) for  $i = i_{t+1}$ :

$$y_{j,t+2} \leftarrow y_{j,t+1} - h g_j - h \left( \phi'_{i_{t+1}}(a_{i_{t+1}}^T y_{j,t+1}) - \phi'_{i_{t+1}}(a_{i_{t+1}}^T x_j) \right) a_{i_{t+1}}. \tag{37}$$

---

**Algorithm 3** Semi-Stochastic Gradient Descent (S2GD) for sparse data; “lazy” updates

---

**parameters:**  $m = \max$  # of stochastic steps per epoch,  $h =$  stepsize,  $\nu =$  lower bound on  $\mu$

**for**  $j = 0, 1, 2, \dots$  **do**

$g_j \leftarrow \frac{1}{n} \sum_{i=1}^n f'_i(x_j)$

$y_{j,0} \leftarrow x_j$

$\chi^{(s)} \leftarrow 0$  for  $s = 1, 2, \dots, d$   $\triangleright$  Store when a coordinate was updated last time

Let  $t_j \leftarrow t$  with probability  $(1 - \nu h)^{m-t} / \beta$  for  $t = 1, 2, \dots, m$

**for**  $t = 0$  to  $t_j - 1$  **do**

Pick  $i \in \{1, 2, \dots, n\}$ , uniformly at random

**for**  $s \in \text{nnz}(a_i)$  **do**

$y_{j,t}^{(s)} \leftarrow y_{j,t}^{(s)} - (t - \chi^{(s)}) h g_j^{(s)}$   $\triangleright$  Update what will be needed

$\chi^{(s)} = t$

**end for**

$y_{j,t+1} \leftarrow y_{j,t} - h \left( \phi'_i(a_i^T y_{j,t}) - \phi'_i(a_i^T x_j) \right) a_i$   $\triangleright$  A sparse update

**end for**

**for**  $s = 1$  to  $d$  **do**  $\triangleright$  Finish all the “lazy” updates

$y_{j,t_j}^{(s)} \leftarrow y_{j,t_j}^{(s)} - (t_j - \chi^{(s)}) h g_j^{(s)}$

**end for**

$x_{j+1} \leftarrow y_{j,t_j}$

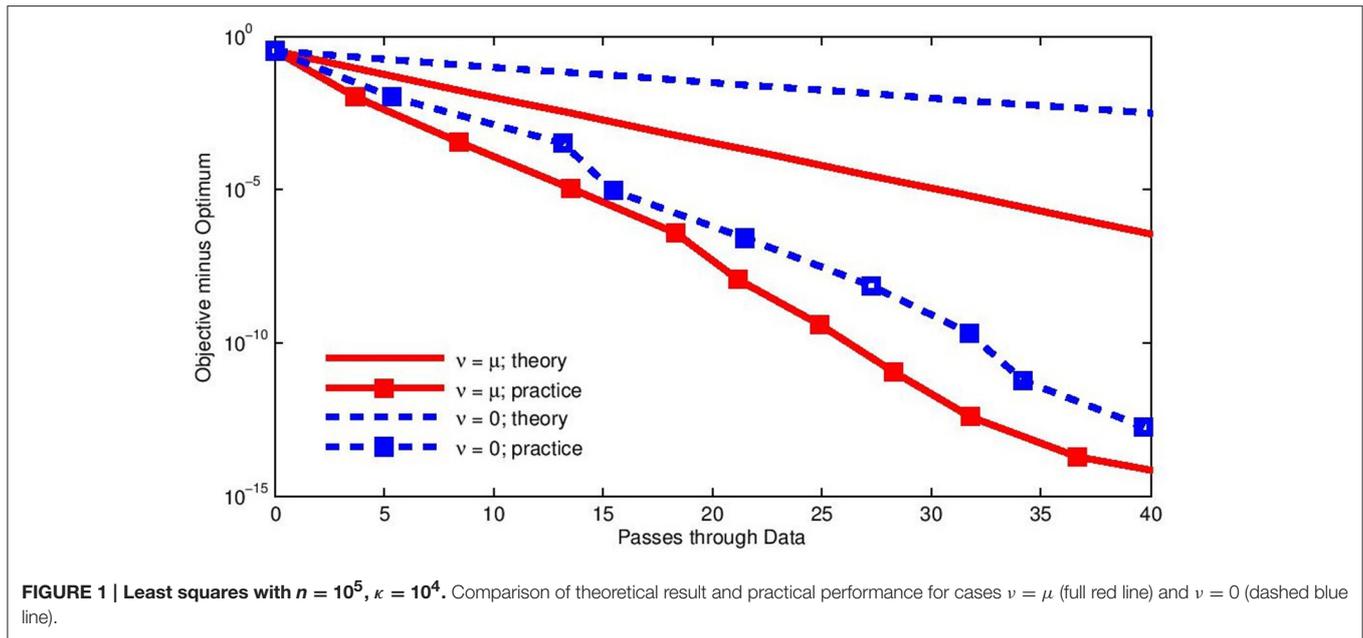
**end for**

---

However, notice that we can't compute

$$\phi'_{i_{t+1}}(a_{i_{t+1}}^T y_{j,t+1}) \tag{38}$$

as we never computed  $y_{j,t+1}$ . However, here lies the trick: as  $a_{i_{t+1}}$  is sparse, we only need to know those coordinates  $s$  of  $y_{j,t+1}$  for



which  $a_{i_{t+1}}^{(s)}$  is nonzero. So, just before we compute the (sparse part of) of the update Equation (37), we perform the update

$$y_{j,t+1}^{(s)} \leftarrow \tilde{y}_{j,t+1}^{(s)} - hg_j^{(s)}$$

for coordinates  $s$  for which  $a_{i_{t+1}}^{(s)}$  is nonzero. This way we know that the inner product appearing in Equation (38) is computed correctly (despite the fact that  $y_{j,t+1}$  potentially is not!). In turn, this means that we can compute the sparse part of the update in Equation (37).

We now continue as before, again only computing  $\tilde{y}_{j,t+3}$ . However, this time we have to be more careful as it is no longer true that

$$y_{j,t+2} = \tilde{y}_{j,t+2} - hg_j.$$

We need to remember, for each coordinate  $s$ , the last iteration counter  $t$  for which  $a_{i_t}^{(s)} \neq 0$ . This way we will know how many times did we “forget” to apply the dense update  $-hg_j^{(s)}$ . We do it in a just-in-time fashion, just before it is needed.

Algorithm 3 (sparse S2GD) performs these lazy updates as described above. It produces exactly the same result as Algorithm 1 (S2GD), but is much more efficient for sparse data as iteration picking example  $i$  only costs  $O(\omega_i)$ . This is done with a memory overhead of only  $O(d)$  (as represented by vector  $x \in \mathbb{R}^d$ ).

## 8. NUMERICAL EXPERIMENTS

In this section we conduct computational experiments to illustrate some aspects of the performance of our algorithm. In Section 8.1 we consider the least squares problem with synthetic data to compare the practical performance and the theoretical bound on convergence in expectations. We demonstrate that for

**TABLE 4 | Datasets used in the experiments.**

Dataset	Training examples ( $n$ )	Variables ( $d$ )	$L$	$\mu$	$\kappa$
<i>ijcnn</i>	49,990	23	1.23	$1/n$	61,696
<i>rcv1</i>	20,242	47,237	0.50	$1/n$	10,122
<i>real-sim</i>	72,309	20,959	0.50	$1/n$	36,155
<i>url</i>	2,396,130	3,231,962	128.70	$100/n$	3,084,052

both SVRG and S2GD, the practical rate is substantially better than the theoretical one. In Section 8.2 we compare the S2GD algorithm on several real datasets with other algorithms suitable for this task. We also provide efficient implementation of the algorithm, as described in Section 7, for the case of logistic regression in the MLOSS repository<sup>8</sup>.

### 8.1. Comparison with Theory

Figure 1 presents a comparison of the theoretical rate and practical performance on a larger problem with artificial data, with a condition number we can control (and choose it to be poor). In particular, we consider the L2-regularized least squares with

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

for some  $a_i \in \mathbb{R}^d$ ,  $b_i \in \mathbb{R}$  and  $\lambda > 0$  is the regularization parameter.

We consider an instance with  $n = 100,000$ ,  $d = 1,000$  and  $\kappa = 10,000$ . We run the algorithm with both parameters  $v = \lambda$  (our best estimate of  $\mu$ ) and  $v = 0$ . Recall that the latter choice

<sup>8</sup><http://mloss.org/software/view/556/>

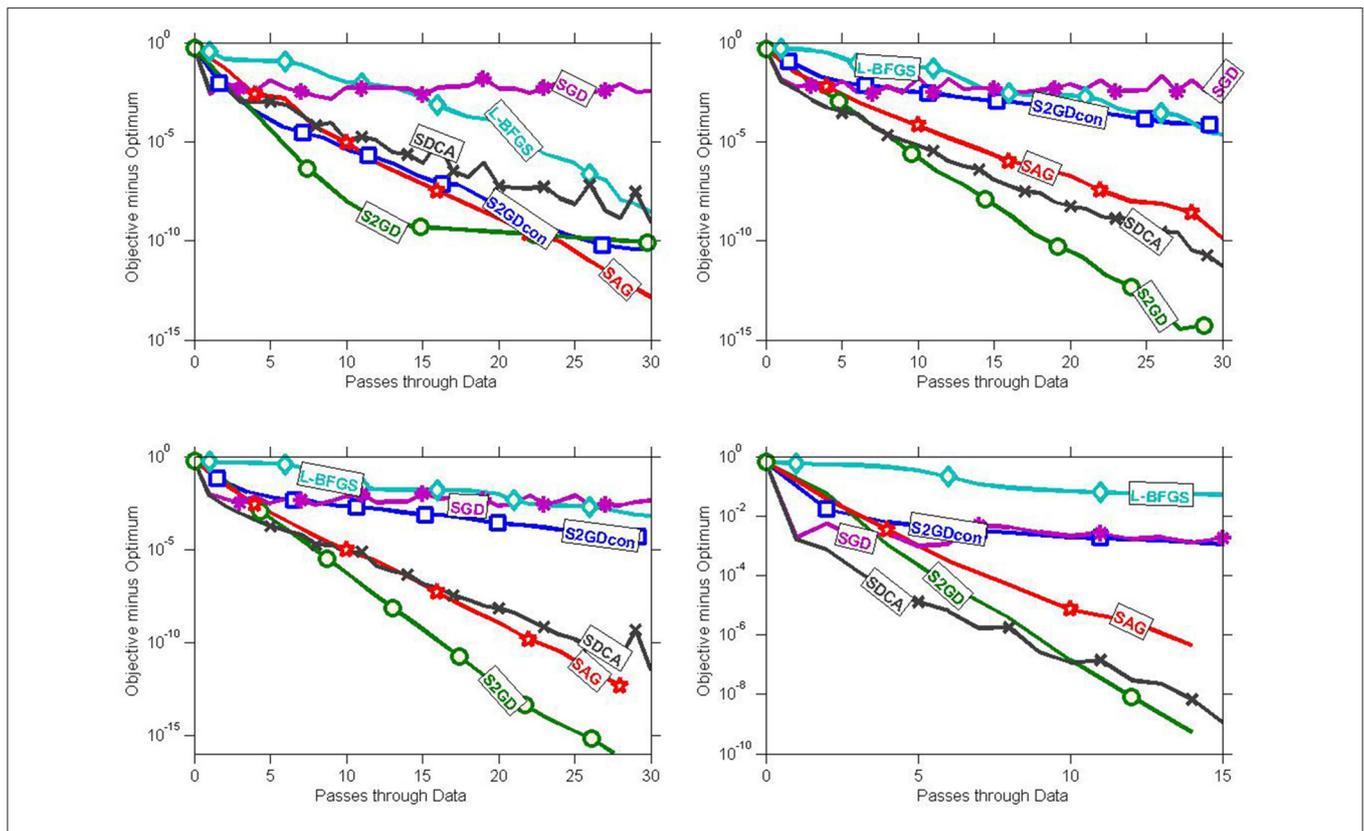


FIGURE 2 | Practical performance for logistic regression and the following datasets: *ijcnn*, *rcv* (first row), *realsim*, *url* (second row).

leads to the SVRG method of [7]. We chose parameters  $m$  and  $h$  as a (numerical) solution of the work-minimization problem (20), obtaining  $m = 261,063$  and  $h = 1/11.4L$  for  $\nu = \lambda$  and  $m = 426,660$  and  $h = 1/12.7L$  for  $\nu = 0$ . The practical performance is obtained after a single run of the S2GD algorithm.

The figure demonstrates linear convergence of S2GD in practice, with the convergence rate being significantly better than the already strong theoretical result. Recall that the bound is on the expected function values. We can observe a rather strong convergence to machine precision in work equivalent to evaluating the full gradient only 40 times. Needless to say, neither SGD nor GD have such speed. Our method is also an improvement over [7], both in theory and practice.

### 8.2. Comparison with other Methods

The S2GD algorithm can be applied to several classes of problems. We perform experiments on an important and in many applications used L2-regularized logistic regression for binary classification on several datasets. The functions  $f_i$  in this case are:

$$f_i(x) = \log \left( 1 + \exp \left( l_i a_i^T x \right) \right) + \frac{\lambda}{2} \|x\|^2,$$

where  $l_i$  is the label of  $i^{\text{th}}$  training example  $a_i$ . In our experiments we set the regularization parameter  $\lambda = O(1/n)$  so that the

TABLE 5 | Time required to produce plots in Figure 2.

Algorithm	Time in seconds			
	<i>ijcnn</i>	<i>rcv1</i>	<i>real-sim</i>	<i>url</i>
S2GDcon	0.25	0.43	1.01	125.53
S2GD	0.29	0.49	1.02	54.04
SAG	0.41	0.73	1.87	71.74
L-BFGS	0.15	0.67	0.76	309.14
SGD	0.39	0.57	1.54	62.73
SDCA	0.33	0.38	1.10	126.32

condition number  $\kappa = O(n)$ , which is about the most ill-conditioned problem used in practice. We added a (regularized) bias term to all datasets.

All the datasets we used, listed in Table 4, are freely available<sup>9</sup> benchmark binary classification datasets.

In the experiment, we compared the following algorithms:

- **SGD:** Stochastic Gradient Descent. After various experiments, we decided to use a variant with constant step-size that gave the best practical performance in hindsight.

<sup>9</sup> Available at <http://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/>.

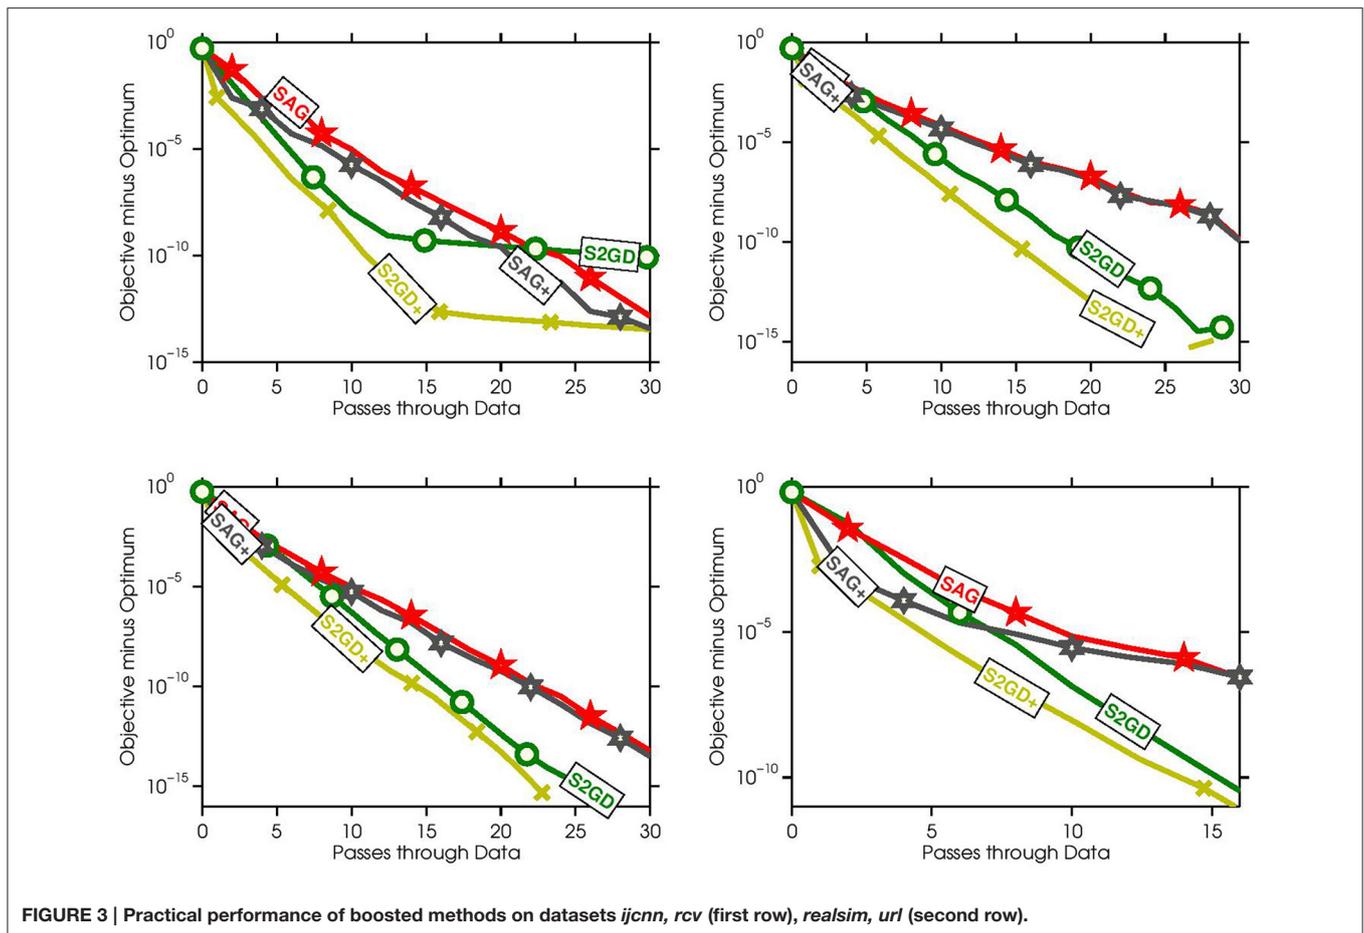


FIGURE 3 | Practical performance of boosted methods on datasets *ijcnn*, *rcv* (first row), *realsim*, *url* (second row).

- **L-BFGS**: A publicly-available limited-memory quasi-Newton method that is suitable for broader classes of problems. We used a popular implementation by Mark Schmidt<sup>10</sup>.
- **SAG**: Stochastic Average Gradient, Schmidt et al. [5]. This is the most important method to compare to, as it also achieves linear convergence using only stochastic gradient evaluations. Although the method has been analyzed for stepsize  $h = 1/16L$ , we experimented with various stepsizes and chose the one that gave the best performance for each problem individually.
- **SDCA**: Stochastic Dual Coordinate Ascent, where we used approximate solution to the one-dimensional dual step, as in Section 6.2 of Shalev-Shwartz and Zhang [6].
- **S2GDcon**: The S2GD algorithm with conservative stepsize choice, i.e., following the theory. We set  $m = O(\kappa)$  and  $h = 1/10L$ , which is approximately the value you would get from Equation (24).
- **S2GD**: The S2GD algorithm, with stepsize that gave the best performance in hindsight. The best value of  $m$  was between  $n$  and  $2n$  in all cases, but optimal  $h$  varied from  $1/2L$  to  $1/10L$ .

Note that SAG needs to store  $n$  gradients in memory in order to run. In case of relatively simple functions, one can store only  $n$  scalars, as the gradient of  $f_i$  is always a multiple of  $a_i$ . If we are comparing with SAG, we are implicitly assuming that our memory limitations allow us to do so. Although not included in Algorithm (1), we could also store these gradients we used to compute the full gradient, which would mean we would only have to compute a single stochastic gradient per inner iteration (instead of two).

We plot the results of these methods, as applied to various different, in the **Figure 2** for first 15–30 passes through the data (i.e., amount of work equivalent to 15–30 full gradient evaluations).

There are several remarks we would like to make. First, our experiments confirm the insight from Schmidt et al. [5] that for this types of problems, reduced-variance methods consistently exhibit substantially better performance than the popular L-BFGS algorithm.

The performance gap between S2GDcon and S2GD differs from dataset to dataset. A possible explanation for this can be found in an extension of SVRG to proximal setting Xiao and Zhang [22], released after the first version of this paper was put onto arXiv (i.e., after December 2013). Instead Assumption 1, where all loss functions are assumed to be associated with the

<sup>10</sup><http://www.di.ens.fr/~mschmidt/Software/minFunc.html>

same constant  $L$ , the authors of Xiao and Zhang [22] instead assume that each loss function  $f_i$  has its own constant  $L_i$ . Subsequently, they sample proportionally to these quantities as opposed to the uniform sampling. In our case,  $L = \max_i L_i$ . This weighted sampling has an impact on the convergence: one gets dependence on the average of the quantities  $L_i$  and not in their maximum.

The number of passes through data seems a reasonable way to compare performance, but some algorithms could need more time to do the same amount of passes through data than others. In this sense, S2GD should be in fact faster than SAG due to the following property. While SAG updates the test point after each evaluation of a stochastic gradient, S2GD does not always make the update—during the evaluation of the full gradient. This claim is supported by computational evidence: SAG needed about 20–40% more time than S2GD to do the same amount of passes through data.

Finally, in Table 5 we provide the time it took the algorithm to produce these plots on a desktop computer with Intel Core i7 3610QM processor, with  $2 \times 4$  GB DDR3 1,600 MHz memory. The numbers for the *url* dataset is are not representative, as the algorithm needed extra memory, which slightly exceeded the memory limit of our computer.

### 8.3. Boosted variants of S2GD and SAG

In this section we study the practical performance of boosted methods, namely S2GD+ (Algorithm 2) and variant of SAG suggested by its authors [5, Section 4.2].

SAG+ is a simple modification of SAG, where one does not divide the sum of the stochastic gradients by  $n$ , but by the number of training examples seen during the run of the algorithm, which has the effect of producing larger steps at the beginning. The authors claim that this method performed better in practice than a hybrid SG/SAG algorithm.

We have observed that, in practice, starting SAG from a point close to the optimum, leads to an initial “away jump.” Eventually, the method exhibits linear convergence. In contrast, S2GD converges linearly from the start, regardless of the starting position.

Figure 3 shows that S2GD+ consistently improves over S2GD, while SAG+ does not improve always: sometimes it

performs essentially the same as SAG. Although S2GD+ is overall a superior algorithm, one should note that this comes at the cost of having to choose stepsize parameter for SGD initialization. If one chooses these parameters poorly, then S2GD+ could perform worse than S2GD. The other three algorithms can work well without any parameter tuning.

## 9. CONCLUSION

We have developed a new semi-stochastic gradient descent method (S2GD) and analyzed its complexity for smooth convex and strongly convex loss functions. Our methods need  $O((\kappa/n) \log(1/\varepsilon))$  work only, measured in units equivalent to the evaluation of the full gradient of the loss function, where  $\kappa = L/\mu$  if the loss is  $L$ -smooth and  $\mu$ -strongly convex, and  $\kappa \leq 2L/\varepsilon$  if the loss is merely  $L$ -smooth.

Our results in the strongly convex case match or improve on a few very recent results, while at the same time generalizing and simplifying the analysis. Additionally, we proposed S2GD+—a method which equips S2GD with an SGD pre-processing step—which in our experiments exhibits superior performance to all methods we tested. We leave the analysis of this method as an open problem.

## AUTHOR CONTRIBUTIONS

All authors listed, have made substantial, direct and intellectual contribution to the work, and approved it for publication.

## FUNDING

The work of both authors was supported by the Centre for Numerical Algorithms and Intelligent Software (funded by EPSRC grant EP/G036136/1 and the Scottish Funding Council). Both authors also thank the Simons Institute for the Theory of Computing, UC Berkeley, where this work was conceived and finalized. The work of PR was also supported by the EPSRC grant EP/I017127/1 (Mathematics for Vast Digital Resources) and EPSRC grant EP/K02325X/1 (Accelerated Coordinate Descent Methods for Big Data Problems).

## REFERENCES

- Nemirovski A, Juditsky A, Lan G, Shapiro A. Robust stochastic approximation approach to stochastic programming. *SIAM J Optim.* (2009) **19**:1574–609. doi: 10.1137/070704277
- Zhang T. Solving large scale linear prediction problems using stochastic gradient descent algorithms. In: *The International Conference on Machine Learning*. Banff, AB (2004).
- Richtárik P, Takáč M. Iteration complexity of randomized block-coordinate descent methods for minimizing a composite function. *Math Progr.* (2014) **144**:1–38. doi: 10.1007/s10107-012-0614-z
- Roux NL, Schmidt M, Bach FR. A stochastic gradient method with an exponential convergence rate for finite training sets. In: *Advances in Neural Information Processing Systems*. Lake Tahoe (2012) p. 2663–71.
- Schmidt M, Le Roux N, Bach F. Minimizing finite sums with the stochastic average gradient. *Math Progr.* (2017) **162**:83–112. doi: 10.1007/s10107-016-1030-6
- Shalev-Shwartz S, Zhang T. Stochastic dual coordinate ascent methods for regularized loss minimization. *J Mach Learn Res.* (2013) **14**:567–99. Available online at: <http://www.jmlr.org/papers/v14/shalev-shwartz13a.html>
- Johnson R, Zhang T. Accelerating stochastic gradient descent using predictive variance reduction. In: *Advances in Neural Information Processing Systems*. Lake Tahoe (2013).
- Csiba D, Richtárik P. Coordinate descent face-off: primal or dual? arXiv preprint arXiv:160508982 (2016).
- Hsieh CJ, Chang KW, Lin CJ, Keerthi SS, Sundarajan S. A dual coordinate descent method for large-scale linear SVM. In: *International Conference on Machine Learning*. Helsinki (2008). doi: 10.1145/1390156.1390208

10. Takáč M, Bijral A, Richtárik P, Srebro N. Mini-batch primal and dual methods for SVMs. In: *International Conference on Machine Learning* (2013).
11. Richtárik P, Takáč M. Parallel coordinate descent methods for big data optimization. *Math Progr.* (2016) **156**:433–84. doi: 10.1007/s10107-015-0901-6
12. Fercoq O, Richtárik P. Smooth minimization of nonsmooth functions with parallel coordinate descent methods. arXiv:13095885 (2013).
13. Richtárik P, Takáč M. Distributed coordinate descent method for learning with big data. *J Mach Learn Res.* (2016) **17**:1–25. Available online at: <http://www.jmlr.org/papers/v17/15-001.html>
14. Zhang L, Mahdavi M, Jin R. Linear convergence with condition number independent access of full gradients. In: *Advances in Neural Information Processing Systems*. Lake Tahoe (2013).
15. Friedlander MP, Schmidt M. Hybrid deterministic-stochastic methods for data fitting. *SIAM J Sci Comput.* (2012) **34**:A1380–405. doi: 10.1137/110830629
16. Deng G, Ferris MC. Variable-number sample-path optimization. *Math Progr.* (2009) **117**:81–109. doi: 10.1007/s10107-007-0164-y
17. Bastin F, Cirillo C, Toint PL. Convergence theory for nonconvex stochastic programming with an application to mixed logit. *Math Progr.* (2006) **108**:207–34. doi: 10.1007/s10107-006-0708-6
18. Marti K, Fuchs E. On solutions of stochastic programming problems by descent procedures with stochastic and deterministic directions. *Methods Operat Res.* (1979) **33**:281–93.
19. Marti K, Fuchs E. Rates of convergence of semi-stochastic approximation procedures for solving stochastic optimization problems. *Optimization* (1986) **17**:243–65. doi: 10.1080/02331938608843124
20. Konečný J, Liu J, Richtárik P, Takáč M. Mini-batch semi-stochastic gradient descent in the proximal setting. *IEEE J Select Top Signal Process.* (2016) **10**:242–55. doi: 10.1109/JSTSP.2015.2505682
21. Nesterov Y. *Introductory Lectures on Convex Optimization: A Basic Course*. Vol. 87. Springer (2004).
22. Xiao L, Zhang T. A proximal stochastic gradient method with progressive variance reduction. *SIAM J Optim.* (2014) **24**:2057–75. doi: 10.1137/140961791

**Conflict of Interest Statement:** The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2017 Konečný and Richtárik. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.