

Utilizing Second Order Information in Minibatch Stochastic Variance Reduced Proximal Iterations

Jialei Wang

*Department of Computer Science
University of Chicago
Chicago, IL, 60637*

JIALEI@CS.UCHICAGO.EDU

Tong Zhang

*Computer Science & Mathematics
Hong Kong University of Science and Technology
Hong Kong*

TONGZHANG@TONGZHANG-ML.ORG

Editor: Francis Bach

Abstract

We present a novel minibatch stochastic optimization method for empirical risk minimization of linear predictors. The method efficiently leverages both sub-sampled first-order and higher-order information, by incorporating variance-reduction and acceleration techniques. We prove improved iteration complexity over state-of-the-art methods under suitable conditions. In particular, the approach enjoys global fast convergence for quadratic convex objectives and local fast convergence for general convex objectives. Experiments are provided to demonstrate the empirical advantage of the proposed method over existing approaches in the literature.

1. Introduction

We consider the following optimization problem of finite-sums:

$$\min_w f(w) := \frac{1}{n} \sum_{i=1}^n f_i(w), \quad (1)$$

which arises in many machine learning applications. In this paper, we focus on regularized loss minimization of linear predictors, where $f_i(w) = \ell(w^\top x_i; b_i) + \frac{\lambda}{2} \|w\|^2$. Let $x_1, x_2, \dots, x_n \in \mathbb{R}^d$ be feature vectors of n data samples, and $b_1, \dots, b_n \in \mathbb{R}$ or $\{-1, 1\}$ be the corresponding target variables of interest. We note that (1) covers many popular models used in machine learning: for example when $\ell(w^\top x_i; b_i) = \frac{1}{2}(b_i - w^\top x_i)^2$ we obtain ridge regression; when $\ell(w^\top x_i; b_i) = \log(1 + \exp(-b_i w^\top x_i))$ we obtain ℓ_2 regularized logistic regression. Other examples such as SVMs can also be obtained by setting different loss functions $\ell(w^\top x_i; b_i)$. The ubiquitousness of such finite-sum optimization problems and the massive scale of modern datasets motivate significant research effort on efficient optimization algorithms to solve (1).

We denote $w^* = \arg \min_w f(w)$ to be the optimum of (1). For any approximate solution w , we say that it achieves ϵ -objective suboptimality if $f(w) - f(w^*) \leq \epsilon$. For any $\epsilon > 0$, the runtime cost of optimization algorithms to find an ϵ -suboptimal minimizer typically

depends on the target accuracy ϵ , as well as the conditions of the problem (1). Throughout the paper, we use the following notion of strong convexity and smoothness to discuss the iteration complexity of various optimization algorithms.

Definition 1 *A function $f_i(w)$ is L -smooth with respect to w if $f_i(w)$ is differentiable and its gradient is L -Lipschitz continuous, i.e. we have*

$$\|\nabla f_i(w) - \nabla f_i(w')\| \leq L \|w - w'\|, \quad \forall w, w' \in \mathbb{R}^d.$$

A function $f_i(w)$ is λ -strongly convex with respect to w , if we have

$$f_i(w) \geq f_i(w') + \langle \nabla f_i(w'), w - w' \rangle + \frac{\lambda}{2} \|w - w'\|^2, \quad \forall w, w' \in \mathbb{R}^d.$$

It is well-known that a consequence of L -smoothness is the following quadratic upper bound for $f_i(w)$:

$$f_i(w) \leq f_i(w') + \langle \nabla f_i(w'), w - w' \rangle + \frac{L}{2} \|w - w'\|^2, \quad \forall w, w' \in \mathbb{R}^d.$$

In recent years there have been significant advances in developing fast optimization algorithms for (1), and we refer the readers to (Bottou et al., 2016) for a comprehensive survey of these developments. For large-scale problems in the form of (1), randomized methods are particularly efficient because of their low per iteration computation. Below we briefly review two lines of research: i) randomized variance reduced first-order methods; ii) randomized methods leveraging second-order information. Through out this section we focus on problems where each function $f_i(w), \forall i \in [n]$ is λ -strongly convex and L -smooth. For non-strongly convex or non-smooth objectives, we can use regularization or smoothing technique to reduce it to strongly convex and smooth objectives (Allen-Zhu and Hazan, 2016).

1.1 Variance Reduced First-order Methods

The key technique for developing fast stochastic first-order methods is variance reduction, which makes the variance of the randomized update direction approaching zero when the iterate gets closer to the optimum. Representative methods of this category include SAG (Roux et al., 2012), SVRG (Johnson and Zhang, 2013), SDCA (Shalev-Shwartz and Zhang, 2013) and SAGA (Defazio et al., 2014), etc. In order to find a solution that reaches ϵ -suboptimality, these methods requires

$$\mathcal{O}\left(\left(n + \frac{L}{\lambda}\right) \log\left(\frac{1}{\epsilon}\right)\right) \quad (2)$$

calls to the first-order oracle to compute the gradient of an individual function. In the large condition number regime (e.g. $\frac{L}{\lambda} > n$), by using an acceleration technique (such as Catalyst (Shalev-Shwartz and Zhang, 2016; Frostig et al., 2015; Lin et al., 2015a), APCG (Lin et al., 2015b), SPDC (Zhang and Xiao, 2015), Katyusha (Allen-Zhu, 2016), etc), one can further improve the iteration complexity to

$$\mathcal{O}\left(\left(n + \sqrt{n \cdot \frac{L}{\lambda}}\right) \log\left(\frac{1}{\epsilon}\right)\right). \quad (3)$$

Recently, Woodworth and Srebro (2016) provided a first-order oracle lower bound for (1) which has the same iteration complexity as (3). This means that (3) is not improvable for general (high dimensional) problems, if the optimization procedure only uses the first order and proximal oracles of individual functions.

1.2 Leveraging Second-order Information

Second-order information is often useful in improving the convergence of optimization algorithms. However for large-scale problems, obtaining and inverting the exact Hessian matrix is often computational expensive, making vanilla Newton methods not well suited in solving (1). Therefore, there have been emerging studies in designing randomized algorithms that effectively utilize the *approximated* Hessian information. One line of such research is the sub-sampled Newton method which approximates the Hessian matrix based on a sub-sampled minibatch of data. Several work, such as (Byrd et al., 2011; Erdogdu and Montanari, 2015; Roosta-Khorasani and Mahoney, 2016a,b; Bollapragada et al., 2016), established local linear convergence rates for different variants of sub-sampled Newton methods, when the size of the sub-sampled data is large enough. (Xu et al., 2016) considered non-uniform sampling in constructing the stochastic Hessian matrix, and showed that weighted sampling according to individual smoothness or leverage score can help constructing better approximation of Hessian, which improves the convergence. (Pilanci and Wainwright, 2017) discussed how to use sketching instead of sub-sampling to approximate the Hessian matrix. All of the aforementioned methods employ the full first-order gradient, and require calling the second-order oracle to compute the Hessian and solving a resulting linear system at every iteration. Therefore the computational complexity (as compared in Table 2 of (Xu et al., 2016)) are often worse than stochastic algorithms such as SVRG.

Another line of research is to consider randomness in both first and second-order information to design lower per-iteration cost algorithms. In particular, (Byrd et al., 2016) considered combining minibatch SGD with Limited-memory BFGS (L-BFGS) (Liu and Nocedal, 1989) type update. Inspired by this, (Moritz et al., 2016; Gower et al., 2016; Wang et al., 2017b) proposed to combine variance reduced stochastic gradient with L-BFGS, and proved linear convergence for strongly convex and smooth objectives. Moreover, experiments in (Moritz et al., 2016; Gower et al., 2016; Wang et al., 2017b) also demonstrated superior empirical performance of this type of methods. However, theoretically it is hard to guarantee the quality of approximated Hessian using L-BFGS update, and thus the iteration complexity obtained by (Moritz et al., 2016; Gower et al., 2016; Wang et al., 2017b) can be pessimistic, and can be much worse than vanilla SVRG. Moreover, (Qu et al., 2015) proposed to incorporate curvature information in minibatch SDCA methods, and showed improved convergence over SDCA, but the method proposed in (Qu et al., 2015) involves solving a much more expensive subproblem than minibatch SDCA. Therefore the overall runtime benefit is still unclear. (Gonen et al., 2016) considered ridge regression problems specifically, and proposed to use sketching to compute the rank- k approximation of the Hessian matrix, based on which SVRG is ran with preconditioning. (Yang et al., 2016) suggested to use preconditioning as a preprocessing step for general stochastic optimization problems. (Agarwal et al., 2016) proposed to approximate the inverse Hessian matrix directly by sampling from its Taylor expansion. In terms of lower bound, (Arjevani and Shamir, 2016) showed under

some mild algorithmic assumptions, second-order oracle generally cannot improve the oracle complexity over first-order methods, for very high-dimensional problems.

Though lower bound have been established showing that the iteration complexity of accelerated SVRG methods cannot be improved in general, the establishments of these lower bounds are based on constructing very high-dimensional hard problems where dimension d can be much larger than sample size n . Such pessimistic situations are not consistent with practice, where second-order information is observed to be very helpful in improving the convergence speed. Therefore it is still interesting to analyze theoretically how second-order information can be helpful under suitable conditions, and more importantly, how to design more efficient methods that make use of possibly noisy second-order information. In this paper, we make effort in solving empirical risk minimization problems with linear predictors more efficiently, and make the following contributions:

- We propose a novel approach that combines the advantages of variance-reduced first-order methods and sub-sampled Newton methods, in a efficient way that does not require expensive Hessian matrix computation and inversion. The method can be naturally extended to solve composite optimization problems with non-smooth regularization.
- We theoretically show under certain conditions the proposed approach can improve state-of-the-art iteration complexity, and empirically demonstrate it can substantially improve the convergence of existing methods. Note that though we only establish theoretical guarantees for empirical risk minimization of linear predictors, the methods it self can be applied to general finite-sum problems.

1.3 Notations

We use $[n]$ to denote the set $1, \dots, n$. For a vector $w \in \mathbb{R}^d$, we use $\|w\|$ to denote its ℓ_2 norm and $\|w\|_1$ to denote its ℓ_1 norm. For a matrix X , we use $\|X\|_2$ to denote its spectral norm and $\|X\|_F$ to denote its Frobenius norm, and $\lambda_{\min}(X)$ to denote its minimum singular value, and $\text{tr}(X)$ to denote the trace of X . We use I to denote an identity matrix. For two symmetric matrices A and B , we denote $A \succeq B$ if $A - B$ is positive semi-definite. For two sequences of numbers $\{a_n\}$ and $\{b_n\}$, we say $a_n = \mathcal{O}(b_n)$ if $a_n \leq Cb_n$ for n large enough, with some positive constant C . We use the notation $\tilde{\mathcal{O}}(\cdot)$ to hide poly-log factors. We also use $a_n \lesssim b_n$ to denote $a_n = \mathcal{O}(b_n)$, $a_n \gtrsim b_n$ to denote $b_n = \mathcal{O}(a_n)$, $a_n \asymp b_n$ to denote $a_n = \mathcal{O}(b_n)$ and $b_n = \mathcal{O}(a_n)$.

1.4 Organization

The rest of the paper is organized as follows: we introduce the proposed method in Section 2, and present the main theoretical results in Section 3. In Section 4 we present the convergence analysis of inexact minibatch accelerated SVRG update as a key step in proving the main results, and this analysis might be of independent interest. We provide some empirical comparisons over existing approaches in Section 5, and conclusions in Section 6. Some detailed proofs are deferred to the Appendices.

2. Minibatch Stochastic Variance Reduced Proximal Iterations

In this section we present the proposed approach for minimizing (1), which naturally integrate both noisy first-order and higher-order information in an efficient way.

2.1 Algorithm Description

The high-level description is given in Algorithm 1, which consists of two main innovations:

- simultaneous incorporation of first-order and higher-order information via sub-sampling;
- allowing larger minibatch sizes through acceleration with inexact minimization.

We explain these features in the remaining of this section.

2.1.1 SAMPLING BOTH FIRST-ORDER AND HIGHER-ORDER INFORMATION

We first discuss the main building block of the algorithm. Suppose at iteration t , given the previous iterate w_{t-1} , we consider the following update rule:

$$\begin{aligned}
 w_t = \arg \min_w & \frac{1}{b} \sum_{i \in \bar{B}} f_i(w) - \left\langle \frac{1}{b} \sum_{i \in \bar{B}} \nabla f_i(w_{t-1}), w \right\rangle \\
 & + \left\langle \frac{\eta}{b} \sum_{i \in B_t} \nabla f_i(w_{t-1}) + \eta \nabla f(\tilde{w}) - \frac{\eta}{b} \sum_{i \in B_t} \nabla f_i(\tilde{w}), w \right\rangle + \frac{\bar{\lambda}}{2} \|w - w_{t-1}\|^2, \quad (6)
 \end{aligned}$$

where \bar{B}, B_t are some randomly sampled minibatch from $1, \dots, n$, both with minibatch size b ; η and $\bar{\lambda}$ are stepsize parameters, and \tilde{w} is a ‘‘reference’’ predictor used for reducing variance. The main feature of updating rule (6) is it considered both noisy first-order and higher-order information via minibatch sampling. In particular, the term $\left\langle \frac{\eta}{b} \sum_{i \in B_t} \nabla f_i(w_{t-1}) + \eta \nabla f(\tilde{w}) - \frac{\eta}{b} \sum_{i \in B_t} \nabla f_i(\tilde{w}), w \right\rangle$ and $\frac{1}{b} \sum_{i \in \bar{B}} f_i(w) - \left\langle \frac{1}{b} \sum_{i \in \bar{B}} \nabla f_i(w_{t-1}), w \right\rangle$ in (6) incorporates variance-reduced first-order and noisy higher-order information of the objective function, respectively. To see this, second-order methods (e.g. Newton methods) often use the following second-order Taylor approximation of the function:

$$\begin{aligned}
 \frac{1}{b} \sum_{i \in \bar{B}} f_i(w) & \approx \frac{1}{b} \sum_{i \in \bar{B}} f_i(w_{t-1}) + \left\langle \frac{1}{b} \sum_{i \in \bar{B}} \nabla f_i(w_{t-1}), w - w_{t-1} \right\rangle \\
 & + \frac{1}{2} (w - w_{t-1})^\top \left(\frac{1}{b} \sum_{i \in \bar{B}} \nabla^2 f_i(w_{t-1}) \right) (w - w_{t-1}). \quad (7)
 \end{aligned}$$

Algorithm 1 MB-SVRP: Minibatch Stochastic Variance Reduced Proximal Iterations.

Parameters $\eta, b, \bar{\lambda}, \nu, \varepsilon$.

Initialize $\tilde{w}_0 = 0$.

Sampling Sampling b items from $[n]$ to form a minibatch \bar{B} .

for $s = 1, 2, \dots, S$ **do**

Calculate $\tilde{v} = \frac{1}{n} \sum_{i=1}^n \nabla f_i(\tilde{w}_{s-1})$.

Initialize $y_0 = w_0 = \tilde{w}_s$.

for $t = 1, 2, \dots, T$ **do**

Sampling b items from $[n]$ to form a minibatch B_t .

Find w_t **that approximately solve** (4) (**Option-I**) or (5) (**Option-II**) such that $\tilde{f}_t(w_t) - \min_w \tilde{f}_t(w) \leq \varepsilon$.

$$\begin{aligned} \text{Option-I: } \tilde{f}_t(w) &:= \frac{1}{b} \sum_{i \in \bar{B}} f_i(w) - \left\langle \frac{1}{b} \sum_{i \in \bar{B}} \nabla f_i(y_{t-1}), w \right\rangle \\ &\quad + \left\langle \frac{\eta}{b} \sum_{i \in B_t} \nabla f_i(y_{t-1}) + \eta \tilde{v} - \frac{\eta}{b} \sum_{i \in B_t} \nabla f_i(\tilde{w}_{s-1}), w \right\rangle + \frac{\bar{\lambda}}{2} \|w - y_{t-1}\|^2, \end{aligned} \tag{4}$$

$$\begin{aligned} \text{Option-II: } \tilde{f}_t(w) &:= \frac{1}{2} (w - y_{t-1})^\top \left(\frac{1}{b} \sum_{i \in \bar{B}} \nabla^2 f_i(y_{t-1}) \right) (w - y_{t-1}) \\ &\quad + \left\langle \frac{\eta}{b} \sum_{i \in B_t} \nabla f_i(y_{t-1}) + \eta \tilde{v} - \frac{\eta}{b} \sum_{i \in B_t} \nabla f_i(\tilde{w}_{s-1}), w \right\rangle + \frac{\bar{\lambda}}{2} \|w - y_{t-1}\|^2. \end{aligned} \tag{5}$$

Update:

$$y_t = w_t + \nu(w_t - w_{t-1}).$$

end for

Update $\tilde{w}_s = w_m$.

end for

Return \tilde{w}_s

If we plug the second-order approximation (7) into (6), we get the update rule as

$$\begin{aligned} w_t &\approx \arg \min_w \frac{1}{2} (w - w_{t-1})^\top \left(\frac{1}{b} \sum_{i \in \bar{B}} \nabla^2 f_i(w_{t-1}) \right) (w - w_{t-1}) + \frac{\bar{\lambda}}{2} \|w - w_{t-1}\|^2 \\ &\quad + \left\langle \frac{\eta}{b} \sum_{i \in B_t} \nabla f_i(w_{t-1}) + \eta \nabla f(\tilde{w}) - \frac{\eta}{b} \sum_{i \in B_t} \nabla f_i(\tilde{w}), w \right\rangle \end{aligned} \tag{8}$$

$$= w_{t-1} - \eta \left(\frac{1}{b} \sum_{i \in \bar{B}} \nabla^2 f_i(w_{t-1}) + \bar{\lambda} I \right)^{-1} \left(\frac{1}{b} \sum_{i \in B_t} \nabla f_i(w_{t-1}) + \nabla f(\tilde{w}) - \frac{1}{b} \sum_{i \in B_t} \nabla f_i(\tilde{w}) \right), \tag{9}$$

which can be treated as a variant of sub-sampled Newton method, combined with the minibatch stochastic gradient with variance reduction. Moreover, when $f_i(w)$ is a quadratic function of w , the approximation (7) is exact (thus Option-I and Option-II in Algorithm 1 are coincident). Therefore the update rule (6) can be treated as a preconditioned minibatch SVRG update rule, with $(\frac{1}{b} \sum_{i \in \bar{B}} \nabla^2 f_i(w_{t-1}) + \bar{\lambda}I)^{-1}$ as a precondition matrix. This is formalized in the proposition below.

Proposition 2 *When $f_i(w), \forall i$ is a quadratic function of w , then the update rule of (4) and (5) is equivalent to the following preconditioned minibatch SVRG update rule:*

$$w_t \leftarrow w_{t-1} - \eta (\bar{H} + \bar{\lambda}I)^{-1} \left(\frac{1}{b} \sum_{i \in B_t} \nabla f_i(w_{t-1}) + \nabla f(\tilde{w}) - \frac{1}{b} \sum_{i \in B_t} \nabla f_i(\tilde{w}) \right),$$

where

$$\bar{H} = \frac{1}{b} \sum_{i \in \bar{B}} \nabla^2 f_i(w)$$

is the sub-sampled Hessian matrix.

2.1.2 CONNECTION TO EXISTING METHODS

Depending on the choice of these parameters, we observe that the update rule of (6) can be viewed as a generalization of several update rules proposed recently:

- When $b = 1, \bar{\lambda} \rightarrow \infty$ and $\eta = \frac{\bar{\lambda}}{L}$, the term $\frac{1}{b} \sum_{i \in \bar{B}} f_i(w) - \langle \frac{1}{b} \sum_{i \in \bar{B}} \nabla f_i(w_{t-1}), w \rangle$ is negligible, thus (6) reduced to standard SVRG update (Johnson and Zhang, 2013):

$$w_t \leftarrow w_{t-1} - \frac{1}{L} \left(\nabla f_i(w_{t-1}) + \frac{1}{n} \sum_{i=1}^n \nabla f_i(\tilde{w}) - \nabla f_i(\tilde{w}) \right).$$

- When $b > 1, \bar{\lambda}, \eta \rightarrow \infty$, then the term $\frac{1}{b} \sum_{i \in \bar{B}} f_i(w) - \langle \frac{1}{b} \sum_{i \in \bar{B}} \nabla f_i(w_{t-1}), w \rangle$ in (6) is negligible as well, (6) reduces to the following update:

$$w_t \leftarrow w_{t-1} - \frac{\eta}{\bar{\lambda}} \left(\frac{1}{b} \sum_{i \in B_t} \nabla f_i(w_{t-1}) + \frac{1}{n} \sum_{i=1}^n \nabla f_i(\tilde{w}) - \frac{1}{b} \sum_{i \in B_t} \nabla f_i(\tilde{w}) \right),$$

which recovers the update rule of minibatch semi-stochastic gradient methods (a.k.a minibatch SVRG) (Konečný et al., 2016).

- When $\bar{B} = B_t, b = 1, \eta = 1$, (6) reduced to stochastic variance reduced proximal iterations. More specifically, (6) will be reduced to:

$$w_t = \arg \min_w f_i(w) + \left\langle \frac{1}{n} \sum_{i=1}^n \nabla f_i(\tilde{w}) - \nabla f_i(\tilde{w}), w \right\rangle + \frac{\bar{\lambda}}{2} \|w - w_{t-1}\|^2, \quad (10)$$

by checking the first order optimality condition it is easy to verify that (10) is performing the following update:

$$w_t \leftarrow w_{t-1} - \frac{1}{\bar{\lambda}} \left(\nabla f_i(w_t) + \frac{1}{n} \sum_{i=1}^n \nabla f_i(\tilde{w}) - \nabla f_i(\tilde{w}) \right),$$

compared with standard SVRG update with stepsize $\frac{1}{\lambda}$, we see the only difference is the gradient evaluation on f_i is based on “future” iterate w_t rather than the current iterate w_{t-1} in SVRG. Stochastic proximal iterations based on SAGA method has been proposed and analyzed in (Defazio, 2016) recently.

- When $B_t = \bar{B}$, $b > 1$, $\eta = 1$, and $\tilde{w} = w_{t-1}$, (6) will be reduced to:

$$w_t = \arg \min_w \frac{1}{b} \sum_{i \in \bar{B}} f_i(w) + \left\langle \frac{1}{n} \sum_{i=1}^n \nabla f_i(w_{t-1}) - \frac{1}{b} \sum_{i \in \bar{B}} \nabla f_i(w_{t-1}), w \right\rangle + \frac{\bar{\lambda}}{2} \|w - w_{t-1}\|^2, \quad (11)$$

(11) covers the update rule of DANE algorithm (Shamir et al., 2014), which is a communication-efficient distributed optimization algorithm. DANE uses the data on local machine to form the minibatch \bar{B} and every round machines communicate the gradient vector based on their local data. As shown in (Shamir et al., 2014), DANE is provably communication more efficient than the first-order methods in certain scenarios.

- When $\bar{B} = B_t$, and we ignore the linear term $-\langle \frac{1}{b} \sum_{i \in \bar{B}} \nabla f_i(w_{t-1}), w \rangle + \langle \frac{\eta}{b} \sum_{i \in B_t} \nabla f_i(w_{t-1}) + \eta \nabla f(\tilde{w}) - \frac{\eta}{b} \sum_{i \in B_t} \nabla f_i(\tilde{w}), w \rangle$, then (6) reduces to the minibatch proximal iterations (Li et al., 2014; Wang et al., 2017a):

$$w_t = \arg \min_w \frac{1}{b} \sum_{i \in B_t} f_i(w) + \frac{\bar{\lambda}}{2} \|w - w_{t-1}\|^2.$$

Such an update allows larger minibatch size than standard minibatch SGD, but without the linear correction term as we considered in (6). For such methods only sublinear convergence can be established for finite-sum problems.

2.1.3 LARGE MINIBATCH SIZE VIA ACCELERATION WITH INEXACT MINIMIZATION

Using (6) as the building block, we propose the MB-SVRP (minibatch stochastic variance reduced proximal iterations) method, which is detailed in Algorithm 1. At the beginning of the algorithm, we form a minibatch \bar{B} by sampling from $1, \dots, n$ and fix it for the whole optimization process¹. Then following the SVRG method (Johnson and Zhang, 2013), the algorithm is divided to multiple stages, indexed by s . At each stage, we iteratively solve a minimization problem of the form (4) (Option I) or (5) (Option II) based on the randomly sampled minibatch B_t .

Compared with the simple update rule in (6), the major difference in Algorithm 1 is that we consider a momentum scheme by maintaining two sequences $\{w_t, y_t\}$, which is inspired by Nesterov’s acceleration technique (Nesterov, 2004) and its recent SVRG variant (Nitanda, 2014). The main theoretical advantage over minibatch SVRG without momentum (Konečný et al., 2016) is that such an acceleration allows us to use a much larger minibatch

1. We can also consider the varying \bar{B} option by simply setting $\bar{B} = B_t$, which makes the algorithm simpler to implement. In practice, we observe no significant difference between this option and the pre-fixed \bar{B} . Here we consider fixed \bar{B} mainly for the sake of simplicity in our theoretical analysis.

size (up to a size of $\mathcal{O}(\sqrt{n})$) without slowing down the convergence. In comparison the iteration complexity of standard minibatch SVRG will become worse when the minibatch size (Konečný et al., 2016) increases. The advantage of allowing larger minibatch size using acceleration in SVRG type algorithms is in analogy to the situation of stochastic gradient descent (without variance reduction) type algorithms (Dekel et al., 2012; Cotter et al., 2011; Lan, 2012).

Since it is often expensive to find the exact minimizer of (4), we consider an approximate minimizer with objective suboptimality ε . When we choose the appropriate $\bar{\lambda}$ for $\tilde{f}_t(w)$ to obtain enough strong convexity, the subproblems defined in Option-I and Option-II can be both be solved to high accuracy efficiently. For example, using SVRG-type algorithms (Johnson and Zhang, 2013) to solve (4) or (5) allows us to find an approximate solution with a small suboptimality ε with a constant number of passes over the data in the minibatch $\bar{B} \cup B_t$, if $\bar{\lambda}$ is set appropriately (discussed in details in our theoretical analysis). This approach avoids Hessian matrix construction and inversion operations, which are often computationally expensive for large-scale problems. Allowing error in gradient oracle has been analyzed in several first-order methods (Schmidt et al., 2011; Devolder et al., 2014) in the batch setting, but has been largely unexplored in stochastic gradient methods. The recent work of (Wang et al., 2017a) analyzed inexact minibatch proximal updates, which can be treated as an implicit minibatch stochastic gradient with errors. In Section 4, we establish convergence rate of *inexact, minibatched, accelerated* SVRG method, which might be of independent interest.

2.2 Extension to Composite Minimization

For many methods that try to incorporate second-order information such as sub-sampled Newton and L-BFGS type algorithms, it is not clear how to extend them to solving non-smooth composite problems. In contrast, the proposed approach can be easily extended to handle non-smooth composite minimization problems. Consider the minimization of:

$$F(w) := \frac{1}{n} \sum_{i=1}^n f_i(w) + g(w), \quad (12)$$

where the component functions $f_i(w), \forall i \in [n]$ are smooth and strongly convex, and $g(w)$ is a non-smooth regularizer. For example when $g(w) = \mu \|w\|_1$ and $f_i(w) = \frac{1}{2}(b_i - w^\top x_i)^2$, we get the Lasso objective (Tibshirani, 1996); when $g(w) = \mu \|w\|_1 + \frac{\lambda}{2} \|w\|^2$ and $f_i(w) = \log(1 + \exp(-b_i w^\top x_i))$, we get the elastic net regularized logistic regression (Zou and Hastie, 2005).

We can easily modify Algorithm 1 to solve (12). The idea is rather straightforward: at each inner iteration, rather than (4), we simply solve the following minibatch composite minimization problem:

$$\begin{aligned} w_t \approx \arg \min_w \tilde{F}_t(w) &:= \frac{1}{\eta b} \sum_{i \in \bar{B}} f_i(w) - \left\langle \frac{1}{\eta b} \sum_{i \in \bar{B}} \nabla f_i(y_{t-1}), w \right\rangle \\ &+ \left\langle \frac{1}{b} \sum_{i \in B_t} \nabla f_i(y_{t-1}) + \tilde{v} - \frac{1}{b} \sum_{i \in B_t} \nabla f_i(\tilde{w}_{s-1}), w \right\rangle + \frac{\bar{\lambda}}{2\eta} \|w - y_{t-1}\|^2 + g(w). \end{aligned} \quad (13)$$

We slightly re-scale the approximated loss term to make sure that the relative weight between the approximated loss and the regularization term is correct. Since (13) is a standard finite-sum composite minimization problem, we could apply prox-SVRG or prox-SAGA to solve (13) efficiently when the term $\frac{\lambda}{2\eta} \|w - y_{t-1}\|^2$ has sufficient strong convexity.

3. Convergence Analysis

In this section we present theoretical results for the proposed approach. As declared in the introduction, we focus on establishing convergence analysis on linear model empirical risk minimization problems because the Hessian structure is easier to interpret. However, the algorithms itself can be implemented for solving general finite-sum problems, and the theory can also be extended to handle such case. For non-quadratic convex objectives, our convergence analysis is local, i.e. we assume the initial and subsequent solutions are within a bounded region near the optimum, which is common in the analysis of second-order methods in optimization (Nesterov, 2004; Boyd and Vandenberghe, 2004). As will be discussed in the sequel, the size of the region depends on the condition number of the problem, as well as the Lipschitz parameter of the Hessian. For quadratic objectives, the radius of the region can be infinite, which implies global fast convergence.

In this section, for simplicity, we focus on the analysis of Option-II in Algorithm 1 which uses second-order approximation to construct each sub-problem (5). Similar results hold for Option-I because locally around the optimal solution, the two methods are equivalent. Note that for linear predictors $f_i(w)$ admits the form $f_i(w) = \ell(w^\top x_i; b_i) + \frac{\lambda}{2} \|w\|^2$, where the Hessian can be computed as $\nabla^2 f_i(y_{t-1}) = \ell''(y_{t-1}^\top x_i; b_i) x_i x_i^\top + \lambda I$. Moreover, we do not need to explicitly maintain the Hessian matrix: when calculating the stochastic gradient of equation (5), especially the gradient of $(1/2)(w - y_{t-1})^\top \nabla^2 f_i(y_{t-1})(w - y_{t-1})$, we get $\nabla^2 f_i(y_{t-1})(w - y_{t-1}) = \ell''(y_{t-1}^\top x_i; b_i)(x_i^\top (w - y_{t-1}))x_i + \lambda(w - y_{t-1})$, which can be computed in time $\mathcal{O}(d)$. Thus approximately solving objective equation (5) would only requires a few stochastic gradients calculation (or equivalent operations in $\mathcal{O}(d)$), we count total number of such gradient computation operations as a measure of Algorithms 1's efficiency in optimization.

Besides the strong convexity and smoothness, we also need the following Hessian Lipschitz condition.

Condition 1 For each loss function $\ell(w^\top x_i; b_i), \forall i \in [n]$, its second derivative is \bar{M} -Lipschitz, i.e.

$$|\ell''(w_1^\top x_i; b_i) - \ell''(w_2^\top x_i; b_i)| \leq \bar{M} \|w_1 - w_2\|, \forall w_1, w_2 \in \mathbb{R}^d, \forall i \in [n].$$

Conditions of Lipschitz Hessians are typically not required for first-order methods such as SVRG and SAGA, but are typically required for second-order methods, such as Newton methods and sub-sampled Newton methods. Moreover, Lipschitz Hessians condition is often satisfied for popular loss functions used in machine learning such as logistic loss and square loss. We consider the component function in form of convex loss function of linear predictor: $f_i(w) = \ell(w^\top x_i; b_i) + \frac{\lambda}{2} \|w\|_2^2$. By Condition 1, combined with the condition that

maximum norm of data points are bounded: $\max_i \|x_i\|^2 \leq D$, we know Hessian matrix of each component loss function $f_i(w)$ is $M = \bar{M}D + \lambda$ -Lipschitz, i.e.

$$\|\nabla^2 f_i(w_1) - \nabla^2 f_i(w_2)\| \leq M \|w_1 - w_2\|, \forall w_1, w_2 \in \mathbb{R}^d, \forall i \in [n].$$

For $f(w) = (1/n) \sum_i f_i(w)$, we have the following formula of Hessian matrix:

$$H_\lambda(w) = \frac{1}{n} \sum_{i=1}^n \nabla^2 f_i(w) = \frac{1}{n} \sum_{i=1}^n \ell''(w^\top x_i; b_i) x_i x_i^\top + \lambda I.$$

Let $H_\lambda(w_{t-1})$ be the Hessian matrix at w_{t-1} , and suppose we have an approximated Hessian

$$\bar{H}_\lambda(w_{t-1}) = \frac{1}{b} \sum_{i \in \bar{B}} \ell''(w_{t-1}^\top x_i; b_i) x_i x_i^\top + \tilde{\lambda} I,$$

the following proposition states that performing the preconditioned minibatch SVRG update is equivalent to the standard minibatch SVRG in a linear transformed (preconditioned) space:

Proposition 3 *Considering the following minibatch SVRG type update when solving $\min_w f(w)$:*

$$w_t \leftarrow w_{t-1} - \eta \bar{H}_\lambda(w_{t-1})^{-1} \left(\frac{1}{b} \sum_{i \in B_t} \nabla f_i(w_{t-1}) + \frac{1}{n} \sum_{i=1}^n \nabla f_i(\tilde{w}) - \frac{1}{b} \sum_{i \in B_t} \nabla f_i(\tilde{w}) \right), \quad (14)$$

which is equivalent to solving a minimization problem with respect to z :

$$\min_z f(\bar{H}_\lambda^{-1/2}(w^*)z)$$

via the variable transform $z = \bar{H}_\lambda^{1/2}(w^*)w$. Moreover, the update rule (14) is equivalent to the following update on z :

$$\begin{aligned} z_t \leftarrow z_{t-1} - \eta \left(\frac{1}{b} \sum_{i \in B_t} \nabla_z f_i(\bar{H}_\lambda^{-1/2}(w^*)z_{t-1}) + \frac{1}{n} \sum_{i=1}^n \nabla_z f_i(\bar{H}_\lambda^{-1/2}(w^*)\tilde{z}) - \frac{1}{b} \sum_{i \in B_t} \nabla_z f_i(\bar{H}_\lambda^{-1/2}(w^*)\tilde{z}) \right) \\ + \eta \bar{H}_\lambda^{1/2}(w^*) (\bar{H}_\lambda(w^*)^{-1} - \bar{H}_\lambda(w_{t-1})^{-1}) \left(\frac{1}{b} \sum_{i \in B_t} \nabla f_i(w_{t-1}) + \frac{1}{n} \sum_{i=1}^n \nabla f_i(\tilde{w}) - \frac{1}{b} \sum_{i \in B_t} \nabla f_i(\tilde{w}) \right). \end{aligned} \quad (15)$$

With the above proposition, and Proposition 2, we can see that the proposed updates (4) and (5) are implicitly performing inexact minibatch SVRG update on the following transformed problem:

$$\min_z f(\bar{H}_\lambda^{-1/2}(w^*)z) := \frac{1}{n} \sum_{i=1}^n \ell(z^\top \bar{H}_\lambda^{-1/2}(w^*)x_i; b_i) + \lambda \left\| \bar{H}_\lambda^{-1/2}(w^*)z \right\|^2, \quad (16)$$

by choosing $\bar{\lambda} = \tilde{\lambda} - \lambda$. In the following, we focus on discussing the choice of $\tilde{\lambda}$, and $\bar{\lambda}$ can be settled accordingly. As a straightforward extension, it is also easy to see that using the acceleration technique as Algorithm 1 is equivalent to performing inexact *accelerated* minibatch SVRG update, similar to the exact update analyzed in (Nitanda, 2014). The analysis of inexact algorithms are presented in Section 4.

The benefit of the transformed problem is that when $\bar{H}_{\tilde{\lambda}}(w^*) \approx H_{\lambda}(w^*)$, the condition number of the new problem (16) becomes smaller than that of the original problem. As a consequence, when the computational cost of update (14) is cheap, we expect improved convergence and runtime guarantees for solving the original problem. Before analyzing the condition number, we state the following lemma, which connects the inexactness in w space when performing the update (8) (and (14)) to the inexactness in z space when performing the update (15).

Lemma 4 *Let $\tilde{f}_t(w)$ be the function to be minimized in (8), and $w_t^* = \arg \min_w \tilde{f}_t(w)$ to be its exact minimizer. Let*

$$\bar{z}_t = z_{t-1} - \eta \left(\frac{1}{b} \sum_{i \in B_t} \nabla_z f_i(\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*)z_{t-1}) + \frac{1}{n} \sum_{i=1}^n \nabla_z f_i(\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*)\tilde{z}) - \frac{1}{b} \sum_{i \in B_t} \nabla_z f_i(\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*)\tilde{z}) \right),$$

be the exact minibatch SVRG update solution in the preconditioned space, and $z_t = \bar{H}_{\tilde{\lambda}}^{-1/2}(w^)w_t$ be the actual update in precondition space by approximately solving (8). Let $\vartheta_t = \frac{1}{b} \sum_{i \in B_t} \nabla f_i(w_{t-1}) + \frac{1}{n} \sum_{i=1}^n \nabla f_i(\tilde{w}) - \frac{1}{b} \sum_{i \in B_t} \nabla f_i(\tilde{w})$, We have the following upper bound on $\|z_t - \bar{z}_t\|$:*

$$\|z_t - \bar{z}_t\| \leq \sqrt{\frac{2(\tilde{f}_t(w_t) - \tilde{f}_t(w_t^*))(L + \tilde{\lambda})}{\lambda + \tilde{\lambda}}} + \frac{\eta M \|w_{t-1} - w^*\| \|\vartheta_t\|}{\tilde{\lambda}^{3/2}}.$$

Remark 5 *Lemma 4 states that the inexactness in the minibatch SVRG update in the z space can be decomposed to two sources, the first source is the fact that we solve the minimization problem of objective $\tilde{f}_t(w)$ inexactly, and the second source is the fact that the Hessian matrix is changing at different locations. Nevertheless, we can fix the z space that use Hessian matrix at optimum $\bar{H}_{\tilde{\lambda}}(w^*)$ as the preconditioner. For quadratic objectives where $M = 0$, the second error term becomes zero.*

3.1 Improved Condition Number

In this section we analyze the condition number in the “preconditioned” space (16), where \bar{B} is sampled uniformly from $1, \dots, n$. For the analysis, we introduce two notions which describes the global and local properties of data. The following definitions of effective dimension and bounded statistical leverage extend the ones defined in (Hsu et al., 2014) which were used to analyzing the generalization performance of ridge regression.

Definition 6 (Effective dimension) *Let the $\lambda_1, \dots, \lambda_d$ be the top- d eigenvalues of $H_0(w^*) = (1/n) \sum_{i=1}^n \ell''(w^{*\top} x_i; b_i) x_i x_i^\top$, define the effective dimension $d_{\tilde{\lambda}}$ (for some $\tilde{\lambda} \geq 0$) of $H_0(w^*)$ being*

$$d_{\tilde{\lambda}} = \sum_{j=1}^d \frac{\lambda_j}{\lambda_j + \tilde{\lambda}}.$$




Figure 1: Effective dimension and maximum statistical leverage with different regularization parameters $\tilde{\lambda}$ on some empirical datasets for ridge regression problems.

We see $d_{\tilde{\lambda}}$ is a decreasing function of $\tilde{\lambda}$, and $d_0 = d$ when $\tilde{\lambda} = 0$. When the spectrum of $\frac{1}{n} \sum_{i=1}^n \ell''(w^{*\top} x_i; b_i) x_i x_i^\top$ is decaying very fast, the effective dimension $d_{\tilde{\lambda}}$ can be significantly smaller than d for moderate $\tilde{\lambda}$. The following lemma demonstrated this for Hessian matrices with fast decaying spectrum, which can be well approximated by low-rank matrices.

Lemma 7 *We have the following upper bound for $d_{\tilde{\lambda}}$:*

$$d_{\tilde{\lambda}} \leq \min_k \left\{ k + \frac{\sum_{j>k} \lambda_j}{\tilde{\lambda}} \right\}.$$

As an example of data with fast decaying spectrum, consider generalized linear predictors in a Hilbert space induced by a Gaussian kernel $K(x_i, x'_i) = \exp(-\|x - x'_i\|^2 / (2\sigma^2))$ where σ is a bandwidth parameter. The covariance matrix of the induced feature space has exponentially decaying eigenvalues with $\lambda_j \leq c_1 \exp(-c_2 j^2)$, where c_1 and c_2 are constants (Zhang et al., 2015). With exponentially decaying eigenvalues, it is shown in Corollary 5 (Zhang et al., 2015) that $d_{\tilde{\lambda}} \leq \mathcal{O}(\sqrt{\log(1/\tilde{\lambda})})$ even though the Gaussian kernel maps the original data to an infinite dimensional space.

The following notion of statistical leverage have been used in regression analysis (Chatterjee and Hadi, 2009; Hsu et al., 2014) and matrix approximation (Mahoney et al., 2011).

Definition 8 (Statistical leverage at $\tilde{\lambda}$) Let $H_{\tilde{\lambda}}(w^*) = \frac{1}{n} \sum_{i=1}^n \ell''(w^{*\top} x_i; b_i) x_i x_i^\top + \tilde{\lambda} I$, we say the statistical leverage of data matrix X is bounded by $\rho_{\tilde{\lambda}}$ at $\tilde{\lambda}$ if

$$\max_{i \in [n]} \frac{\left\| H_{\tilde{\lambda}}^{-1/2}(w^*) \ell''(w^{*\top} x_i; b_i)^{1/2} x_i \right\|}{\sqrt{(1/n) \sum_{j=1}^n \left\| H_{\tilde{\lambda}}^{-1/2}(w^*) \ell''(w^{*\top} x_j; b_j)^{1/2} x_j \right\|^2}} \leq \rho_{\tilde{\lambda}}. \quad (17)$$

For ridge regression problems, the above definition is slightly different from the one used in (Hsu et al., 2014) in the sense that the empirical Hessian matrix $H_{\tilde{\lambda}}$ in (17) is replaced by the population Hessian matrix $\mathbb{E}[x_i^\top x_i] + \tilde{\lambda}I$ in (Hsu et al., 2014). When the sample size n is large, the differences in these two definitions are minor. As argued in (Hsu et al., 2014), when x_i is i.i.d. randomly drawn from sub-Gaussian distributions, $(\mathbb{E}[x_i^\top x_i])^{-1/2} x_i$ is isotropic. In this case the statistical leverage only grows logarithmically with dimension. For the empirical statistical leverage considered in (17), if we consider the un-regularized form (when $\lambda = 0$), then Definition 8 is closely related to the leverage score used in (Mahoney et al., 2011; Drineas and Mahoney, 2017), where the leverage score for x_i is defined as $x_i^\top \left((1/n) \sum_{j=1}^n x_j x_j^\top \right)^{-1} x_i$, and thus statistical leverage is just the ratio between “maximum leverage score” and “average leverage score”. Though the in worse case $\rho_{\tilde{\lambda}}$ can be as large as \sqrt{n} , in many situations it is small. For instance, if the data is generated by random Gaussian distributions, or when we apply certain random projections on the data, the (projected) data matrix is known to have almost uniform leverage scores (i.e. the statistical leverage parameter in (17) can be upper bounded by $\mathcal{O}(\log d)$, see Lemma 70 in (Drineas and Mahoney, 2017) for example).

Our theoretical analysis relies on the condition that the two quantities defined above are not too large. This condition often holds in practice, as shown in Figure 1, which plots the effective dimension and the statistical leverage for several real world datasets with varying regularization parameters.

The convergence of stochastic gradient based algorithms for minimizing the objective (16) depends on two important quantities: strong convexity and smoothness. The strong convexity for problem (16) (with respect to z) is

$$\min_z \lambda_{\min} \left(\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*) \left(\frac{1}{n} \sum_{i=1}^n \ell''(z^\top \bar{H}_{\tilde{\lambda}}^{-1/2}(w^*) x_i; b_i) x_i x_i^\top + \lambda I \right) \bar{H}_{\tilde{\lambda}}^{-1/2}(w^*) \right), \quad (18)$$

which is a property of the overall objective function rather than individual component. However, we must consider the smoothness parameter for individual function $f_i(\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*) z)$ (with respect to z), which in our context, is

$$\max_i \max_z \lambda_{\max} \left(\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*) \left(\ell''(z^\top \bar{H}_{\tilde{\lambda}}^{-1/2}(w^*) x_i; b_i) x_i x_i^\top + \lambda I \right) \bar{H}_{\tilde{\lambda}}^{-1/2}(w^*) \right). \quad (19)$$

By the variable transformation $w = \bar{H}_{\tilde{\lambda}}^{-1/2}(w^*) z$, we get $\ell''(z^\top \bar{H}_{\tilde{\lambda}}^{-1/2}(w^*) x_i; b_i) = \ell''(w^\top x_i; b_i)$. Moreover, we consider a local region where w is close to w^* measured by diameter R : $\|w - w^*\| \leq R$, and study the local strong convexity and smoothness quantities((18) and (19)).

Both of these quantities are related to how close is the constructed Hessian approximation $\bar{H}_{\tilde{\lambda}}(w^*)$ to the true Hessian $H_{\tilde{\lambda}}(w^*)$ in spectral norm. The following lemma bound this quantity using matrix concentration.

Lemma 9 *If \bar{B} is formed by uniform sampling with replacement from $[n]$, then we have the following concentration bound, with probability at least $1 - \delta$,*

$$\left\| H_{\tilde{\lambda}}^{-1}(w^*) \left(\frac{1}{b} \sum_{i \in \bar{B}} \ell''(w^{*\top} x_i; b_i) x_i x_i^\top - \frac{1}{n} \sum_{i=1}^n \ell''(w^{*\top} x_i; b_i) x_i x_i^\top \right) \right\|_2 \leq 2 \log \left(\frac{d}{\delta} \right) \cdot \sqrt{\frac{\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}}{b}}.$$

3.1.1 STRONG CONVEXITY

Based on above lemma, we have the following lower bound of the strong convexity for (16), specified in the following lemma.

Lemma 10 *If we set $\tilde{\lambda}$ such that $\tilde{\lambda} \geq \lambda$, then with probability at least $1 - \delta$ over the random choice of \bar{B} to form $\bar{H}_{\tilde{\lambda}}$, we have*

$$\min_{\|w-w^*\| \leq R} \lambda_{\min} \left(\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*) H_{\lambda}(w) \bar{H}_{\tilde{\lambda}}^{-1/2}(w^*) \right) \geq \frac{\lambda}{\tilde{\lambda}} \frac{1}{1 + 2 \log(d/\delta) \sqrt{(\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}})/b(1 + MR/\tilde{\lambda}) + MR/\tilde{\lambda}}}.$$

3.1.2 SMOOTHNESS

Next we explore the local smoothness parameter of (16). Using the variable transform $w = \bar{H}_{\tilde{\lambda}}^{-1/2}(w^*)z$, if we consider the individual smoothness parameter (with respect to z) around the neighborhood of w^* with diameter R , it can be upper bounded by

$$\max_{i \in [n]} \max_{\|w-w^*\| \leq R} \lambda_{\max} \left(\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*) \left(\ell''(w^{\top} x_i; b_i) x_i x_i^{\top} + \lambda I \right) \bar{H}_{\tilde{\lambda}}^{-1/2}(w^*) \right).$$

Since

$$\begin{aligned} & \lambda_{\max} \left(\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*) \left(\ell''(w^{\top} x_i; b_i) x_i x_i^{\top} + \lambda I \right) \bar{H}_{\tilde{\lambda}}^{-1/2}(w^*) \right) \\ & \leq \lambda_{\max} \left(\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*) \left(\ell''(w^{\top} x_i; b_i) x_i x_i^{\top} \right) \bar{H}_{\tilde{\lambda}}^{-1/2}(w^*) \right) + \lambda_{\max} \left(\lambda \bar{H}_{\tilde{\lambda}}^{-1}(w^*) \right) \\ & \leq \ell''(w^{\top} x_i; b_i) x_i^{\top} \bar{H}_{\tilde{\lambda}}^{-1}(w^*) x_i + \frac{\lambda}{\tilde{\lambda}} \end{aligned} \quad (20)$$

The most straightforward way to upper bound $\ell''(w^{\top} x_i; b_i) x_i^{\top} \bar{H}_{\tilde{\lambda}}^{-1}(w^*) x_i$ is

$$\begin{aligned} & \max_{i \in [n]} \max_{\|w-w^*\| \leq R} \ell''(w^{\top} x_i; b_i) x_i^{\top} \bar{H}_{\tilde{\lambda}}^{-1}(w^*) x_i \\ & \leq \max_{i \in [n]} \max_{\|w-w^*\| \leq R} \ell''(w^{\top} x_i; b_i) \|x_i\|^2 \lambda_{\min}(\bar{H}_{\tilde{\lambda}}^{-1}(w^*)) \\ & = \max_{i \in [n]} \max_{\|w-w^*\| \leq R} \frac{\ell''(w^{\top} x_i; b_i) \|x_i\|^2}{\tilde{\lambda}} = \frac{L}{\tilde{\lambda}}. \end{aligned}$$

In this way, the condition number after preconditioning becomes

$$\mathcal{O} \left(\frac{4\tilde{\lambda}}{\lambda} \cdot \frac{L}{\tilde{\lambda}} \right) = \mathcal{O} \left(\frac{L}{\lambda} \right),$$

which didn't show any advantage of using preconditioning. This is not surprising because it is known that for high dimensional problems, the worse case behavior of second order methods is no better than that of the first order methods.

In the lemma below, we provide an improved analysis of smoothness which is based on the notion of effective dimension (Definition 6) and statistical leverage (Definition 8).

Lemma 11 *If we choose $\tilde{\lambda} \geq \lambda$ and $b \geq 16\rho_\lambda^2 d_\lambda \log^2(d/\delta)$, then with probability at least $1 - \delta$ over the random choice of \bar{B} to form $\bar{H}_{\tilde{\lambda}}(w^*)$, we have*

$$\max_{i \in [n]} \max_{\|w-w^*\| \leq R} \ell''(w^\top x_i; b_i) x_i^\top \bar{H}_{\tilde{\lambda}}^{-1}(w^*) x_i \leq 2\rho_\lambda^2 d_\lambda + \frac{MR}{\tilde{\lambda}}.$$

Based on above analysis, we have the following corollary states the local individual smoothness parameter of (16).

Corollary 12 *If we choose $\tilde{\lambda} \geq \lambda$ and $b \geq 16\rho_\lambda^2 d_\lambda \log^2(d/\delta)$, and $R \leq \frac{\tilde{\lambda}}{M}$. Then with probability at least $1 - \delta$ over the random choice of \bar{B} to form $\bar{H}_{\tilde{\lambda}}(w^*)$, we have the following upper bound hold for the local smoothness parameter*

$$\max_{i \in [n]} \max_{\|w-w^*\| \leq R} \lambda_{\max} \left(\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*) \left(\ell''(w^\top x_i; b_i) x_i x_i^\top + \lambda I \right) \bar{H}_{\tilde{\lambda}}^{-1/2}(w^*) \right) \leq 2\rho_\lambda^2 d_\lambda + 2.$$

Proof Combining (20) and Lemma 11, as well as the conditions of $\tilde{\lambda}$ and R stated in the corollary we get the desired results. \blacksquare

Combining Lemma 10 and Corollary 12 we get the following result about the condition number for (16). If we choose $\tilde{\lambda}$ and b that satisfy

$$\tilde{\lambda} \geq \max \left\{ \lambda, \frac{L}{b} \right\}, \quad b \geq 16\rho_\lambda^2 d_\lambda \log^2 \left(\frac{d}{\delta} \right), \quad R \leq \min \left\{ \frac{1}{4ML}, \frac{\tilde{\lambda}}{4M} \right\},$$

then with probability at least $1 - \delta$, the condition number for stochastic gradient algorithms after preconditioning scales as $\frac{4\rho_\lambda^2 d_\lambda \tilde{\lambda}}{\lambda}$, more specifically when $\tilde{\lambda} = \max \left\{ \lambda, \frac{L}{b} \right\}$, then the condition number of (16) can be upper bounded by:

$$\rho_\lambda^2 d_\lambda \max \left\{ 4, \frac{4L}{\lambda b} \right\},$$

which improves the original condition number L/λ by a factor of at least $\tilde{\mathcal{O}}(\rho_\lambda^2 d_\lambda / b)$.

3.2 Improved Runtime Guarantee

Based on the above analysis, we have the following main results stating the iteration complexity of MB-SVRP algorithms on general convex objectives (1).

Theorem 13 *Consider Algorithm 1 on problems with L -smooth and λ -strongly convex functions that satisfied Condition 1 with Hessian Lipschitz parameter M . Suppose we sample minibatch \bar{B} uniformly from $1, \dots, n$, and set the tuning parameters as*

$$\tilde{\lambda} = \max \left\{ \lambda, \frac{L}{b} \right\}, \quad b \asymp \min \left\{ n, \rho_\lambda^2 d_\lambda \left(\frac{L}{\lambda} \right)^{1/3} \right\}, \quad \eta \asymp \min \left\{ \frac{b^3 \lambda}{(\rho_\lambda^2 d_\lambda)^2 L}, \frac{1}{\rho_\lambda^2 d_\lambda} \right\},$$

and suppose we start from the initialization point that is sufficiently close to the optimum:

$$\|w_0 - w^*\| \leq R = \mathcal{O}\left(\frac{\lambda^4}{L^2 M}\right). \quad (21)$$

Given $\epsilon > 0$, and let

$$\epsilon \leq \frac{1}{10^5} \cdot \left(\frac{\lambda}{L}\right)^7 \epsilon. \quad (22)$$

Then for the MB-SVRP algorithm to find the approximate solution \tilde{w}_s of (1) that reaches the expected ϵ -objective suboptimality

$$\mathbb{E}f(\tilde{w}_s) - \min_w f(w) \leq \epsilon,$$

we can set

$$T \leq \mathcal{O}\left(\max\left\{\left(\frac{L}{\lambda}\right)^{1/3}, \frac{\rho_\lambda^2 d_\lambda L}{\lambda n^2}\right\}\right),$$

and uses $S \leq \mathcal{O}(\log(1/\epsilon))$ total stages in Algorithm 1. Thus the total number of oracle calls to solve (8) is no more than

$$\mathcal{O}\left(\max\left\{\left(\frac{L}{\lambda}\right)^{1/3}, \frac{\rho_\lambda^2 d_\lambda L}{\lambda n^2}\right\} \cdot \log\left(\frac{1}{\epsilon}\right)\right).$$

Remark 14 There are two source of randomness in Algorithm 1: *i)* sampling \bar{B} to form the precondition minibatch before the iterations begins; *ii)* sampling B_t during the iterations. In particular, sampling \bar{B} affects the new condition number as Section 3.1 analyzed, while sampling B_t affects the precondition SVRG convergence. We state the results about the randomness in \bar{B} in high-probability form as Lemma 10 and Corollary 12; while state the randomness in B_t in expectation results, it is also possible to obtain high-probability results for randomness in B_t by using Markov's inequality, for example as Corollary 1 in (Xiao and Zhang, 2014).

Note that in Algorithm 1, we can use SVRG to solve each subproblem in (8). This leads to the following result.

Corollary 15 Assume that the conditions of Theorem 13 hold. Moreover, if we use SVRG to solve each subproblem in (8) up to the suboptimality ϵ , then the total number of gradient evaluations used in the whole MB-SVRP algorithm can be upper bounded by

$$\mathcal{O}\left(\rho_\lambda^2 d_\lambda \max\left\{\left(\frac{L}{\lambda}\right)^{2/3}, \frac{L}{\lambda n}\right\} \cdot \log\left(\frac{L}{\lambda}\right) \cdot \log^2\left(\frac{1}{\epsilon}\right) + n \cdot \log\left(\frac{1}{\epsilon}\right)\right).$$

We make the following remarks about Theorem 13 and Corollary 15.

Remark 16 *The solution accuracy requirement in (22) (in practice we can monitoring ε via $\|\nabla \tilde{f}_t(w)\|$ since $\tilde{f}_t(w)$ is $\tilde{\lambda}$ -strongly convex) is conservative, but does not affect the overall iteration complexity because solving (4) and (5) to high precision only introduces an additional logarithmic factor using SVRG. In practice, we observe that using only one pass SVRG warm started from previous iterate performs well (please refer to the experiments).*

Remark 17 *For quadratic objectives, where $M = 0$, the local fast region defined in (21) is the whole space, implying global fast convergence for quadratic objectives. Moreover, in experiments, we observe that even for the general non-quadratic objectives, the proposed MB-SVRP method also converges fast globally. Therefore the size of the local fast convergence region (21) established here might be too conservative. We leave the investigation on improving the size of the local fast convergence region to future work.*

Remark 18 *Corollary 15 states that the iteration complexity of Algorithm 1 is*

$$\tilde{O} \left(n + \rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}} \max \left\{ \left(\frac{L}{\tilde{\lambda}} \right)^{2/3}, \frac{L}{\tilde{\lambda}n} \right\} \right),$$

where $\tilde{O}(\cdot)$ hides the minor poly logarithmic factors. When $\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}$ is small so that it can be treated as a constant, then the iteration complexity of MB-SVRP improves over standard SVRG by a factor of $\min \left\{ n, \left(\frac{L}{\tilde{\lambda}} \right)^{1/3} \right\}$ when the condition number $\frac{L}{\tilde{\lambda}}$ is much larger than n . Even compared with accelerated methods (such as SVRG equipped with catalyst acceleration (Shalev-Shwartz and Zhang, 2016; Frostig et al., 2015; Lin et al., 2015a)), it can sometimes be better (when $n < L/\tilde{\lambda} < n^2$, see Table 1 for details).

Remark 19 *Though the optimal parameter choice stated in Theorem 13 might depend on the property of the Hessian matrix at w^* , the algorithm actually works for a wide range of parameters without the knowledge of w^* because i) for quadratic objectives the Hessian matrix does not change with w so the optimal parameter choice will not depend on w^* ; ii) we only requires some constants $\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}$ to bound certain quantities of the Hessian matrix at w^* , but do not need exactly where w^* is; iii) if we are uncertain about the data we can choose some conservative parameters (as discussed in Section 3.1) to ensure the method converges at least as fast as vanilla SVRG.*

Remark 20 *When $\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}$ is small such that can be treated as constant, then the optimal choice of b according to our theory is often smaller than n , except the case where the condition number is extremely large: $L/\tilde{\lambda} \gtrsim (n/(\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}))^3$. Moreover, when combine the proposed approach with Catalyst acceleration as in Section 3.3, the condition number in each proximal point problem can always be controlled to be not too large, in which case the optimal batch size b is smaller than n .*

Remark 21 *MB-SVRP method is able to improve the iteration complexity over SVRG when $\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}$ is small, thus it is also possible to use MB-SVRP itself to solve the minibatch subproblem (4). Such a nested approach allows us to choose a even smaller $\tilde{\lambda}$, thus can further reduce the dependency on condition number in the iteration complexity. However, this is at the cost of increasing the dependency on $\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}$ and more complicated implementation. We will leave such investigation to future research.*

3.3 With Catalyst Acceleration

Algorithm 2 Acc-MB-SVRP: Accelerated MB-SVRP method.

Initialize $w_0 = z_0 = 0$.

for $r = 1, 2, \dots$ **do**

 Call **MB-SVRP** algorithm 1 to approximately solve

$$w_r \approx \arg \min_w f(w) + \frac{\gamma}{2} \|w - z_{r-1}\|^2. \quad (23)$$

Update

$$z_r = w_r + \nu_r(w_r - w_{r-1}).$$

end for

Return w_r .

Table 1: Comparison of iteration complexity of various finite-sum quadratic optimization algorithms when effective dimension and statistical leverage are bounded, where we compare the different relative scale of condition number $\kappa = L/\lambda$ and sample size n , ignoring logarithmic factors.

	$\kappa \leq n$	$\kappa = n^{4/3}$	$\kappa = n^{3/2}$	$\kappa = n^2$	$\kappa = n^3$
SVRG	n	$n^{4/3}$	$n^{3/2}$	n^2	n^3
MB-SVRP	n	$\rho_\lambda^2 d_\lambda \cdot n$	$\rho_\lambda^2 d_\lambda \cdot n$	$\rho_\lambda^2 d_\lambda \cdot n^{3/2}$	$\rho_\lambda^2 d_\lambda \cdot n^{5/2}$
Acc-SVRG	n	$n^{7/6}$	$n^{5/4}$	$n^{3/2}$	n^2
Acc-MB-SVRP	n	$\left(\rho_\lambda^2 d_\lambda\right)^{1/2} \cdot n$	$\left(\rho_\lambda^2 d_\lambda\right)^{1/2} \cdot n$	$\left(\rho_\lambda^2 d_\lambda\right)^{1/2} \cdot n^{5/4}$	$\left(\rho_\lambda^2 d_\lambda\right)^{1/2} \cdot n^{7/4}$

Corollary 15 stated that if the condition number is not too large: $\frac{L}{\lambda} \leq n^{5/4}$, then MB-SVRP only requires logarithmic passes over data to find a solution with high accuracy. When the condition number is large, MB-SVRP can still be slow. By using the accelerated proximal point framework proposed in (Shalev-Shwartz and Zhang, 2016; Frostig et al., 2015; Lin et al., 2015a), it is possible to obtain an accelerated convergence rate which has milder dependence on condition number, the algorithm is outlined in Algorithm 2, which iteratively, approximately call the original algorithm MB-SVRP to solve an augmented proximal point problem. The main iteration complexity is stated in the theorem below.

Theorem 22 For ill-condition problems where $L/\lambda \geq n^{3/2}$, if we set the parameter $\gamma \asymp \frac{\rho_\lambda^2 d_\lambda L}{n^{3/2}} - \lambda$, Algorithm 2 has iteration complexity of

$$\tilde{O}\left(\left(\rho_\lambda^2 d_\lambda\right)^{1/2} \cdot n^{1/4} \left(\frac{L}{\lambda}\right)^{1/2}\right).$$

Remark 23 From theorem 22 we see the iteration complexity of accelerated MB-SVRP grows at a square root rate with respect to the condition number, which is much better than that of the non-accelerated algorithms, and this is similar to the behavior of accelerated SVRG algorithms. Moreover, accelerated MB-SVRP improves the iteration complexity of accelerated SVRG by a factor of $n^{1/4}$ when $\sqrt{\rho_\lambda^2 d_\lambda}$ is small.

4. Convergence Analysis of Inexact Accelerated Minibatch SVRG

Algorithm 3 IMBA-SVRG: Inexact Minibatch Accelerated SVRG Method.

Parameters $\alpha = \sqrt{\eta\lambda/2}$.

Initialize $\tilde{w}_0 = 0$.

for $s = 1, 2, \dots, S$ **do**

Calculate $\tilde{v} = \frac{1}{n} \sum_{i=1}^n \nabla f_i(\tilde{w}_{s-1})$.

Initialize $y_0 = w_0 = \tilde{w}_s$.

for $t = 1, 2, \dots, T$ **do**

Sampling b items from $[n]$ to form a minibatch B_t .

Inexact update w_t such that

$$w_t = \bar{w}_t + \xi_t, \tag{24}$$

where:

$$\bar{w}_t = y_{t-1} - \eta \left(\frac{1}{b} \sum_{i \in B_t} \nabla f_i(y_{t-1}) - \frac{1}{b} \sum_{i \in B_t} \nabla f_i(\tilde{w}_{s-1}) + \tilde{v} \right).$$

Update

$$y_t = w_t + \left(\frac{1 - \alpha}{1 + \alpha} \right) (w_t - w_{t-1}).$$

end for

Update $\tilde{w}_s = w_m$.

end for

Return \tilde{w}_s

As discussed in Propositions 2 and 3, the proposed MB-SVRP method (Algorithm 1) essentially perform the inexact minibatch accelerated SVRG updates in the preconditioned space using $\bar{H}_\lambda^{-1/2}(w^*)$ as the preconditioner. Thus, the main results established in Section 3 rely on the analysis for inexact, minibatch, accelerated SVRG update (IMBA-SVRG Algorithm 3), where in each stochastic step, we allow a small error ξ_t in the updating. In this section, we show that as long as the inexactness at each iteration is small enough, IMBA-SVRG can use a large minibatch size without slowing down the convergence.

Theorem 24 Consider the IMBA-SVRG algorithm. If we choose the stepsize as

$$\eta = \min \left\{ \frac{b^2 \lambda}{6400L^2}, \frac{1}{8L} \right\},$$

and when the deviation ξ_t satisfies $\forall 1 \leq t \leq T$

$$\|\xi_t\|^2 \leq \frac{2\lambda^2\eta^3\sqrt{\lambda\eta}}{15}(f(\tilde{w}_{s-1}) - f(w^*)), \quad (25)$$

and the number of iterations T satisfies

$$T \geq \frac{10}{9\sqrt{\lambda\eta}} \log(36),$$

then

$$\mathbb{E}[f(\tilde{w}_s) - f(w^*)] \leq \frac{1}{2}[f(\tilde{w}_{s-1}) - f(w^*)].$$

Remark 25 Depending on the relative magnitude of minibatch size b and condition number $\frac{L}{\lambda}$, the overall iteration complexity can be divided into the following two situations.

- $b \leq 20\sqrt{\frac{2L}{\lambda}}$. If we choose the stepsize as $\eta = \frac{b^2\lambda}{6400L^2}$, and the deviation ξ_t satisfies $\forall 1 \leq t \leq T$: $\|\xi_t\|^2 \leq \frac{2\lambda^2\eta^3\sqrt{\lambda\eta}}{15}(f(\tilde{w}_{s-1}) - f(w^*))$, then $\mathbb{E}[f(\tilde{w}_s) - f(w^*)] \leq \frac{1}{2}[f(\tilde{w}_{s-1}) - f(w^*)]$ when $T \geq \frac{90L}{\lambda b}$.
- $20\sqrt{\frac{2L}{\lambda}} \leq b \leq n$. If we choose stepsize as $\eta = \frac{1}{8L}$, and when the deviation ξ_t satisfies $\forall 1 \leq t \leq T$: $\|\xi_t\|^2 \leq \frac{2\lambda^2\eta^3\sqrt{\lambda\eta}}{15}(f(\tilde{w}_{s-1}) - f(w^*))$, then $\mathbb{E}[f(\tilde{w}_s) - f(w^*)] \leq \frac{1}{2}[f(\tilde{w}_{s-1}) - f(w^*)]$ when $T \geq \sqrt{\frac{8L}{\lambda}}$.

A direct consequence of Theorem 24 is the following iteration complexity of IMBA-SVRG.

Corollary 26 For IMBA-SVRG algorithm, If we choose stepsize as $\eta = \min\left\{\frac{b^2\lambda}{6400L^2}, \frac{1}{8L}\right\}$, and when at every state s , the deviation ξ_t at every iteration satisfies $\forall 1 \leq t \leq T$: $\|\xi_t\|^2 \leq \frac{2\lambda^2\eta^3\sqrt{\lambda\eta}}{15}(f(\tilde{w}_{s-1}) - f(w^*))$, then $\mathcal{O}(\log(\frac{1}{\epsilon}))$ full gradient evaluations and

$$\mathcal{O}\left(\max\left\{\frac{L}{\lambda b}, \sqrt{\frac{L}{\lambda}}\right\} \cdot \log\left(\frac{1}{\epsilon}\right)\right)$$

inexact update steps of form (24) are sufficient to ensure $\mathbb{E}f(\tilde{w}_s) - f(w^*) \leq \epsilon$.

Our convergence proof of IMBA-SVRG relies on the machinery of stochastic estimation sequence (Lin et al., 2014), which originated from the framework of estimation sequence developed in (Nesterov, 2004). The details are given in Appendix B.

The main difference between our analysis and that of (Lin et al., 2014) is that we no longer require the inequalities in estimation sequences hold almost surely. Instead we only make sure they hold in expectation, and thus the quadratic lower bound used to construct the stochastic estimation sequence is also different. IMBA-SVRG can be viewed as an inexact extension of accelerated minibatch SVRG (Nitanda, 2014). Although (Nitanda, 2014) also considered stochastic estimation sequences, here we allow error in the stochastic

Table 2: List of datasets used in the experiments.

Name	#Instances	#Features	Task
codrna	59,535	8	Classification
covtype	581,012	54	Classification
svmguid3	1,243	21	Classification
synthetic-c	10,000	1,000	Classification
cadata	20,460	8	Regression
spacega	3,107	6	Regression
synthetic-r	20,000	2,000	Regression
year	463,715	91	Regression

Algorithm 4 Practical implementation of MB-SVRP Algorithm used in Experiments.

Parameters: b .

Initialize $\eta = \frac{1}{L}$, $\nu = \frac{1-\sqrt{\lambda\eta}}{1+\sqrt{\lambda\eta}}$, $m = \lceil \frac{n}{b} \rceil$, $\bar{\lambda} = \frac{1}{\sqrt{b}}$, $\tilde{w}_0 = 0$.

Sampling Sampling b items from $[n]$ to form a minibatch \bar{B} .

for $s = 1, 2, \dots$ **do**

Calculate $\tilde{v} = \frac{1}{n} \sum_{i=1}^n \nabla f_i(\tilde{w}_{s-1})$.

Initialize $y_0 = w_0 = \tilde{w}_s$.

for $t = 1, 2, \dots, \lceil 2n/b \rceil$ **do**

Sampling b items from $[n]$ to form a minibatch B_t .

Initialize $w_t = y_{t-1}$.

Calculate $\tilde{u} = \frac{\eta}{b} \sum_{i \in B_t} \nabla f_i(y_{t-1}) + \eta \tilde{v} - \frac{\eta}{b} \sum_{i \in B_t} \nabla f_i(\tilde{w}_{s-1})$.

for $k = 1, 2, \dots, b$ **do**

Sampling Sampling an index i_k from \bar{B} .

Update (MB-SVRP-I): $w_t \leftarrow w_t - \eta (\nabla f_{i_k}(w_t) - \nabla f_{i_k}(y_{t-1}) + \bar{\lambda}(w_t - y_{t-1}) + \tilde{u})$.

Update (MB-SVRP-II): $w_t \leftarrow w_t - \eta (\nabla^2 f_{i_k}(y_{t-1})(w_t - y_{t-1}) + \bar{\lambda}(w_t - y_{t-1}) + \tilde{u})$.

end for

Update: $y_t = w_t + \nu(w_t - w_{t-1})$.

end for

Update $\tilde{w}_s = w_m$.

end for

Return \tilde{w}_s

gradient update, and thus need to construct a different estimation sequence which takes the inexactness into consideration. On the other hand, batch (accelerated) gradient methods with inexact first order oracle have been studied in (Schmidt et al., 2011; Villa et al., 2013; Devolder et al., 2014), and the analysis of IMBA-SVRG extends these results in the context of stochastic gradient methods with variance reduction.

5. Experiments

In this section we compare the proposed MB-SVRP algorithm to several state-of-the-art methods for minimizing (1). The datasets are summarized in Table 2, most of which can be download from the LibSVM website². We also considered a synthetic dataset for each task, where the data $\{x_i, b_i\}_{i=1}^n$ are i.i.d. drawn from the following model:

$$\text{Regression : } b_i = \langle x_i, \bar{w} \rangle + a_i, \quad x_i \sim \mathcal{N}(0, \Sigma), \quad a_i \sim \mathcal{N}(0, 1), \quad \forall i \in [n],$$

$$\text{Classification : } P(b_i = \pm 1) = \frac{\exp(b_i \langle x_i, \bar{w} \rangle)}{1 + \exp(b_i \langle x_i, \bar{w} \rangle)}, \quad x_i \sim \mathcal{N}(0, \Sigma), \quad \forall i \in [n],$$

where entries of \bar{w} are drawn i.i.d. from $\mathcal{N}(0, 1)$. To make the problem ill-conditioned with fast decaying spectrum, we set $\Sigma_{ij} = 2^{-|i-j|/500}, \forall i, j \in [n]$.

5.1 Empirical Results for Smooth Optimization

We consider ridge regression and logistic regression models for regression and classification problems, respectively. We normalize the dataset by $x_i \leftarrow x_i / (\max_i \|x_i\|^2)^{1/2}$ to ensure that the maximum norm of data points is 1. Note that the normalization will not change statistical property of the problem, the purpose here is to make the smoothness parameter fixed and invariant across different datasets, thus easier for us to set a suitable regularization parameter λ for the target condition number. We tried three settings of λ , as $1/n, 10^{-1}/n$ and $10^{-2}/n$ to represent different levels of regularization. It is expected that when $\lambda = 1/n$, the SVRG/SAGA/SDCA algorithms should converge very fast since the condition number is of the same order as sample size n , while for weak regularization case $\lambda = 10^{-2}/n$ these algorithms are expected to converge slowly.

We compare with SVRG (Johnson and Zhang, 2013), SDCA (Shalev-Shwartz and Zhang, 2013), and SAGA (Defazio et al., 2014), as they represent popular variance reduced optimization algorithms. We also compare with the minibatch version of SVRG and SAGA, in particular, we compare with a related minibatch accelerated SVRG (MB-SVRG) (Nittanda, 2014), which allows large minibatch size without slowing down convergence; while for SAGA, we apply the standard minibatch technique for SAGA and get the MB-SAGA algorithm. For quasi-Newton methods, we compare with L-BFGS with 10 as the limited memory size, and with line search of stepsize to satisfy the Wolfe condition (Wright and Nocedal, 1999). For the proposed MB-SVRP method, we implemented both Option I (MB-SVRP-I) and Option II (MB-SVRP-II), for ridge regression problem, these the two options are equivalent. At every iteration we simply run one pass of SVRG method initialized with y_{t-1} , to approximately solve (4). The implementation is summarized in Algorithm 4, where we fixed many parameters except the minibatch size b as the main tuning parameter. For SVRG and MB-SVRP method, we set $T = \lceil 2n/b \rceil$ as suggested in (Johnson and Zhang, 2013). The minibatch size of MB-SVRP are set to be $\max(\min((L/\lambda)^{1/3}, d), 40)$ inspired by theoretical analysis in Section 3, the minibatch size of MB-SVRG and MB-SAGA is set to be the same as MB-SVRP to demonstrate the direct comparison. In Section 5.3 we also performed a sensitivity study on the minibatch size b in MB-SVRP algorithm. Other parameters, such as stepsize in SVRG and SAGA, are tuned to give the fastest convergence.

We use primal objective suboptimality and the number of gradient evaluations as the evaluation metrics. Since gradient calculations account for the major computing part of

2. <https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/>

first-order optimization methods, it is used as a good estimate (Johnson and Zhang, 2013; Shalev-Shwartz and Zhang, 2013; Defazio et al., 2014) of the computation cost, and it can be evaluated independent of specific implementation and computing hardware. Figure 2 and 3 showed the results for ℓ_2 regularized logistic regression and ridge regression, respectively, where we plot how the objective suboptimality $f(w_t) - f(w^*)$ decreases as the number of gradient evaluations divided by sample size (a.k.a. number of effective passes) increases. We have the following observations:

- When the regularization parameter λ (thus strong convexity) is large enough, all methods (except L-BFGS and MB-SAGA) converges very fast, typically using less than 50 passes over data to converges to numerical precision. The slow convergence of MB-SAGA is due to the usage of minibatch which slow down its computation efficiency.
- The convergence of L-BFGS is typically the slowest. When λ is large, the advantages of variance-reduced stochastic methods over L-BFGS are significant.
- When the regularization parameter is small, the proposed MB-SVRP method starts to show advantages. In particular, when $\lambda = 10^{-2}/n$, MB-SVRP is substantially faster than all other methods compared in our study. We also observed there is only negligible performance difference between MB-SVRP-I and MB-SVRP-II.

5.2 Empirical Results for Composite Optimization

We also consider empirical comparisons in the setting of non-smooth composition optimization problems. We adopted the same datasets used for the smooth optimization problems, but considered the elastic-net regularized logistic regression and linear regression models (Zou and Hastie, 2005). Where the objective is in the form of (12) but with a elastic-net regularization:

$$g(w) = \frac{\lambda}{2} \|w\|^2 + \mu \|w\|_1,$$

where μ is set to be $10^{-1}/n$. We follow the same settings as in the previous section, and compare with prox-SVRG (Xiao and Zhang, 2014), prox-SAGA (Defazio et al., 2014), prox MB-SVRG (Nitanda, 2014), and prox MB-SAGA, but didn't compare with L-BFGS because the method is generally only applicable for smooth problems. The results are shown in Figure 4 and 5, and we have similar observations as those of the smooth optimization case: all variance-reduced methods converges reasonably fast when λ is large, but when λ becomes small, the proposed prox MB-SVRP method converges significantly faster because it effectively leveraged the higher-order information from minibatches.

5.3 Sensitivity Study on the Minibatch Size in MB-SVRP

We also performed a sensitivity study on the minibatch size b , to see how b will affect the convergence behavior of MB-SVRP method. The experimental setting follows Section 5.1 and 5.2, except that we varying b to be 16, 32, 64, 128 for MB-SVRP, and the regularization parameter λ is fixed to be $10^{-2}/n$. The results were shown in Figure 6, we observed that MB-SVRP method performed well on a wide range settings of minibatch sizes, and converges fast than SVRG in most cases.

6. Conclusion

In this paper, we proposed a novel minibatch stochastic optimization approach for regularized loss minimization in machine learning. This approach efficiently utilizes both variance-reduced first-order gradients and sub-sampled higher order information, and under suitable conditions, it can provably improve the iteration complexity over that of the previous state-of-the-art. Empirical experiments demonstrated that this approach performs well on a variety of smooth and composite optimization tasks in practice. The minibatch nature of the algorithm makes it useful for parallel and distributed computing environment, where additional speed up can be obtained by using many CPU cores for computation.

Acknowledgment

This work was performed when Jialei Wang was a research intern with Tong Zhang at Tencent AI Lab. The authors would like to thank Ji Liu and anonymous reviewers for their helpful feedbacks.




Figure 2: Comparison of various optimization algorithms for solving ℓ_2 regularized logistic regression problems.

MINIBATCH STOCHASTIC VARIANCE REDUCED PROXIMAL ITERATIONS




Figure 3: Comparison of various optimization algorithms for solving ridge regression problems.




Figure 4: Comparison of various optimization algorithms for solving elastic-net regularized logistic regression problems.

MINIBATCH STOCHASTIC VARIANCE REDUCED PROXIMAL ITERATIONS




Figure 5: Comparison of various optimization algorithms for solving elastic-net regularized regression problems.




Figure 6: Convergence of MB-SVRP algorithms with various minibatch size b (in parentheses), compared with SVRG.

Appendix

The appendix contains proofs of some theorems and lemmas stated in the main paper.

Appendix A. Proofs for Section 3

A.1 Proof of Proposition 3

Proof Let $w^* = \arg \min_w f(w)$, and $z^* = \arg \min_z f(\bar{H}_\lambda^{-1/2}(w^*)z)$, it is clear that $w^* = \bar{H}_\lambda^{-1/2}(w^*)z^*$, which is coincident with the exchange of variable rule $z = \bar{H}_\lambda^{1/2}(w^*)w$. Moreover, since

$$\begin{aligned} & \eta \left(\frac{1}{b} \sum_{i \in B_t} \nabla_z f_i(\bar{H}_\lambda^{-1/2}(w^*)z_{t-1}) + \frac{1}{n} \sum_{i=1}^n \nabla_z f_i(\bar{H}_\lambda^{-1/2}(w^*)\tilde{z}) - \frac{1}{b} \sum_{i \in B_t} \nabla_z f_i(\bar{H}_\lambda^{-1/2}(w^*)\tilde{z}) \right) \\ & \stackrel{\textcircled{1}}{=} \eta \bar{H}_\lambda^{-1/2}(w^*) \left(\frac{1}{b} \sum_{i \in B_t} \nabla f_i(\bar{H}_\lambda^{-1/2}(w^*)z_{t-1}) + \frac{1}{n} \sum_{i=1}^n \nabla f_i(\bar{H}_\lambda^{-1/2}(w^*)\tilde{z}) - \frac{1}{b} \sum_{i \in B_t} \nabla f_i(\bar{H}_\lambda^{-1/2}(w^*)\tilde{z}) \right) \\ & \stackrel{\textcircled{2}}{=} \eta \bar{H}_\lambda^{-1/2}(w^*) \left(\frac{1}{b} \sum_{i \in B_t} \nabla f_i(w_{t-1}) + \frac{1}{n} \sum_{i=1}^n \nabla f_i(\tilde{w}) - \frac{1}{b} \sum_{i \in B_t} \nabla f_i(\tilde{w}) \right), \end{aligned}$$

where at step $\textcircled{1}$ we have used the gradient chain rule, at step $\textcircled{2}$ we used the exchange of variables $z = \bar{H}_\lambda^{1/2}(w^*)w$. Substitute to (15), and multiplying both sides of above equation with $\bar{H}_\lambda^{-1/2}(w^*)$, we get

$$w_t \leftarrow w_{t-1} - \eta \bar{H}_\lambda(w_{t-1})^{-1} \left(\frac{1}{b} \sum_{i \in B_t} \nabla f_i(w_{t-1}) + \frac{1}{n} \sum_{i=1}^n \nabla f_i(\tilde{w}) - \frac{1}{b} \sum_{i \in B_t} \nabla f_i(\tilde{w}) \right),$$

which recovers the exact formulation of (14). \blacksquare

A.2 Proof of Lemma 4

Proof We first perform the following decomposition:

$$\|z_t - \bar{z}_t\| \leq \|z_t - z_t^*\| + \|z_t^* - \bar{z}_t\|, \quad (26)$$

then we can bound $\|z_t - z_t^*\|$ by

$$\begin{aligned} \|z_t - z_t^*\| &= \left\| \bar{H}_\lambda^{1/2}(w^*)(w_t - w_t^*) \right\| \leq \|\bar{H}_\lambda(w^*)\|^{1/2} \|w_t - w_t^*\| \leq \sqrt{L + \tilde{\lambda}} \|w_t - w_t^*\| \\ &\leq \sqrt{\frac{2(\tilde{f}_t(w_t) - \tilde{f}_t(w_t^*))(L + \tilde{\lambda})}{\lambda + \tilde{\lambda}}}, \end{aligned} \quad (27)$$

where in the last inequality we use the $\lambda + \tilde{\lambda}$ -strong convexity of $\tilde{f}_t(w)$. For the term $\|z_t^* - \bar{z}_t\|$, we can bounded it using Hessian Lipschitz property,

$$\begin{aligned} \|z_t^* - \bar{z}_t\| &= \left\| \eta \bar{H}_{\tilde{\lambda}}^{1/2}(w^*) (\bar{H}_{\tilde{\lambda}}(w^*)^{-1} - \bar{H}_{\tilde{\lambda}}(w_{t-1})^{-1}) \vartheta_t \right\| \\ &\leq \eta \left\| \left(\bar{H}_{\tilde{\lambda}}(w^*)^{-1/2} \bar{H}_{\tilde{\lambda}}(w_{t-1})^{-1} \right) (\bar{H}_{\tilde{\lambda}}(w^*) - \bar{H}_{\tilde{\lambda}}(w_{t-1})) \right\| \|\vartheta_t\| \\ &\leq \frac{\eta}{\tilde{\lambda}^{3/2}} M \|w_{t-1} - w^*\| \|\vartheta_t\|. \end{aligned} \quad (28)$$

Combining (26), (27) and (28) finishes the proof. \blacksquare

A.3 Proof of Lemma 7

Proof We first decompose the terms $\frac{\lambda_j}{\lambda_j + \tilde{\lambda}}$ into two parts, the first one contains indexes smaller or equal than k and the second part contains indexes larger than k , then we upper bound these two parts get obtain the stated results.

$$\begin{aligned} d_{\tilde{\lambda}} &= \sum_j \frac{\lambda_j}{\lambda_j + \tilde{\lambda}} = \min_k \left\{ \sum_{j \leq k} \frac{\lambda_j}{\lambda_j + \tilde{\lambda}} + \sum_{j > k} \frac{\lambda_j}{\lambda_j + \tilde{\lambda}} \right\} \\ &\leq \min_k \left\{ \sum_{j \leq k} \frac{\lambda_j + \tilde{\lambda}}{\lambda_j + \tilde{\lambda}} + \sum_{j > k} \frac{\lambda_j}{\tilde{\lambda}} \right\} = \min_k \left\{ k + \frac{\sum_{j > k} \lambda_j}{\tilde{\lambda}} \right\} \end{aligned}$$

\blacksquare

A.4 Proof of Lemma 9

Proof For $j = 1, \dots, b$, define random matrix

$$v_j = \frac{1}{b} \left(H_{\tilde{\lambda}}^{-1/2}(w^*) \ell''(w^{*\top} x_k; b_k) x_k x_k^\top H_{\tilde{\lambda}}^{-1/2}(w^*) - \frac{1}{n} \sum_{i=1}^n H_{\tilde{\lambda}}^{-1/2}(w^*) \ell''(w^{*\top} x_i; b_i) x_i x_i^\top H_{\tilde{\lambda}}^{-1/2}(w^*) \right),$$

with probability $1/n, \forall k \in [n]$. It is easy to check that $\mathbb{E}[v_j] = 0$, and

$$\begin{aligned} \|v_j\| &\leq \frac{1}{b} \left(\left\| H_{\tilde{\lambda}}^{-1/2}(w^*) \ell''(w^{*\top} x_k; b_k) x_k x_k^\top H_{\tilde{\lambda}}^{-1/2}(w^*) \right\| \right) \\ &\quad + \frac{1}{b} \left(\left\| \frac{1}{n} \sum_{i=1}^n H_{\tilde{\lambda}}^{-1/2}(w^*) \ell''(w^{*\top} x_i; b_i) x_i x_i^\top H_{\tilde{\lambda}}^{-1/2}(w^*) \right\| \right) \\ &\leq \frac{1}{b} \left(2 \max_{i \in [n]} \ell''(w^{*\top} x_i; b_i) x_i^\top H_{\tilde{\lambda}}^{-1}(w^*) x_i \right) \leq \frac{2\rho_{\tilde{\lambda}}^2}{b} \left(\frac{1}{n} \sum_{i=1}^n \ell''(w^{*\top} x_i; b_i) x_i^\top H_{\tilde{\lambda}}^{-1}(w^*) x_i \right) \\ &= \frac{2\rho_{\tilde{\lambda}}^2}{b} \text{tr} \left(H_{\tilde{\lambda}}^{-1}(w^*) \left(\frac{1}{n} \sum_{i=1}^n \ell''(w^{*\top} x_i; b_i) x_i x_i^\top \right) \right) = \frac{2\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}}{b}. \end{aligned}$$

for the expected second order moment, denote $\tilde{x}_k = (\ell''(w^{*\top} x_k; b_k))^{1/2} H_{\tilde{\lambda}}^{-1/2}(w^*) x_k$ we have

$$\begin{aligned}
 \|\mathbb{E}[v_j^2]\| &= \left\| \frac{1}{nb^2} \sum_{k=1}^n \left(\tilde{x}_k \tilde{x}_k^\top - \frac{1}{n} \sum_{i=1}^n \tilde{x}_i \tilde{x}_i^\top \right)^2 \right\| \\
 &= \left\| \frac{1}{nb^2} \sum_{k=1}^n \left((\tilde{x}_k \tilde{x}_k^\top)^2 - 2(\tilde{x}_k \tilde{x}_k^\top) \left(\frac{1}{n} \sum_{i=1}^n \tilde{x}_i \tilde{x}_i^\top \right) + \left(\frac{1}{n} \sum_{i=1}^n \tilde{x}_i \tilde{x}_i^\top \right)^2 \right) \right\| \\
 &= \left\| \frac{1}{nb^2} \sum_{k=1}^n (\tilde{x}_k \tilde{x}_k^\top)^2 - \frac{2}{b^2} \left(\frac{1}{n} \sum_{i=1}^n \tilde{x}_i \tilde{x}_i^\top \right)^2 + \frac{1}{b^2} \left(\frac{1}{n} \sum_{i=1}^n \tilde{x}_i \tilde{x}_i^\top \right)^2 \right\| \\
 &= \left\| \frac{1}{nb^2} \sum_{k=1}^n (\tilde{x}_k \tilde{x}_k^\top)^2 - \frac{1}{b^2} \left(\frac{1}{n} \sum_{i=1}^n \tilde{x}_i \tilde{x}_i^\top \right)^2 \right\| \\
 &\leq \left\| \frac{1}{nb^2} \sum_{k=1}^n (\tilde{x}_k \tilde{x}_k^\top)^2 \right\| \leq \left\| \frac{\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}}{b^2} \left(\frac{1}{n} \sum_{k=1}^n \tilde{x}_k \tilde{x}_k^\top \right) \right\| \\
 &= \frac{\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}}{b^2} \left\| H_{\tilde{\lambda}}^{-1/2}(w^*) \left(\frac{1}{n} \sum_{i=1}^n \ell''(w^{*\top} x_i; b_i) x_i x_i^\top \right) H_{\tilde{\lambda}}^{-1/2}(w^*) \right\| \\
 &\leq \frac{\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}}{b^2}.
 \end{aligned}$$

Thus

$$\left\| \sum_{j=1}^b \mathbb{E}[v_j^2] \right\| \leq \frac{\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}}{b},$$

using Lemma 35, we know

$$\begin{aligned}
 &P \left(\left\| H_{\tilde{\lambda}}^{-1/2}(w^*) \left(\frac{1}{b} \sum_{i \in \tilde{B}} \ell''(w^{*\top} x_i; b_i) x_i x_i^\top - \frac{1}{n} \sum_{i=1}^n \ell''(w^{*\top} x_i; b_i) x_i x_i^\top \right) H_{\tilde{\lambda}}^{-1/2}(w^*) \right\|_2 \geq t \right) \\
 &= P \left(\left\| \sum_{j=1}^b v_j \right\| \geq t \right) \leq d \exp \left(\frac{-t^2/2}{\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}/b + 2\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}} t/(3b)} \right),
 \end{aligned}$$

setting $t = 2 \log \left(\frac{d}{\delta} \right) \cdot \sqrt{\frac{\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}}{b}}$ we conclude the proof. \blacksquare

A.5 Proof of Lemma 10

Proof When $\tilde{\lambda} \geq \lambda$, we know

$$\begin{aligned}
 & \lambda_{\min}(\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*)H_{\lambda}(w)\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*)) \\
 & \geq \lambda_{\min}\left(\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*)\left(\frac{\lambda}{\tilde{\lambda}}\left(\frac{1}{n}\sum_{i=1}^n \ell''(w^\top x_i; b_i)x_i x_i^\top\right) + \lambda I\right)\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*)\right) \\
 & = \frac{\lambda}{\tilde{\lambda}}\lambda_{\min}(\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*)H_{\tilde{\lambda}}(w)\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*)).
 \end{aligned}$$

Further more we can lower bound $\lambda_{\min}(\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*)H_{\tilde{\lambda}}(w)\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*))$ by

$$\begin{aligned}
 & \lambda_{\min}(\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*)H_{\tilde{\lambda}}(w)\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*)) \\
 & = \lambda_{\min}(\bar{H}_{\tilde{\lambda}}^{-1}(w^*)H_{\tilde{\lambda}}(w)) \\
 & = \frac{1}{\lambda_{\max}(H_{\tilde{\lambda}}^{-1}(w)\bar{H}_{\tilde{\lambda}}(w^*))} \\
 & \geq \frac{1}{\lambda_{\max}(H_{\tilde{\lambda}}^{-1}(w)H_{\tilde{\lambda}}(w)) + \lambda_{\max}(H_{\tilde{\lambda}}^{-1}(w)(H_{\tilde{\lambda}}(w) - \bar{H}_{\tilde{\lambda}}(w^*)))} \\
 & \geq \frac{1}{1 + \lambda_{\max}(H_{\tilde{\lambda}}^{-1}(w)(H_{\tilde{\lambda}}(w^*) - \bar{H}_{\tilde{\lambda}}(w^*))) + \lambda_{\max}(H_{\tilde{\lambda}}^{-1}(w)(H_{\tilde{\lambda}}(w) - H_{\tilde{\lambda}}(w^*)))}, \tag{29}
 \end{aligned}$$

For the term $\lambda_{\max}(H_{\tilde{\lambda}}^{-1}(w)(H_{\tilde{\lambda}}(w^*) - \bar{H}_{\tilde{\lambda}}(w^*)))$, we have the following decomposition:

$$\begin{aligned}
 \lambda_{\max}(H_{\tilde{\lambda}}^{-1}(w)(H_{\tilde{\lambda}}(w^*) - \bar{H}_{\tilde{\lambda}}(w^*))) & \leq \lambda_{\max}(H_{\tilde{\lambda}}^{-1}(w^*)(H_{\tilde{\lambda}}(w^*) - \bar{H}_{\tilde{\lambda}}(w^*))) \\
 & \quad + \lambda_{\max}((H_{\tilde{\lambda}}^{-1}(w^*) - H_{\tilde{\lambda}}^{-1}(w))(H_{\tilde{\lambda}}(w^*) - \bar{H}_{\tilde{\lambda}}(w^*)))
 \end{aligned} \tag{30}$$

since with probability at least $1 - \delta$, we have

$$\lambda_{\max}(H_{\tilde{\lambda}}^{-1}(w^*)(H_{\tilde{\lambda}}(w^*) - \bar{H}_{\tilde{\lambda}}(w^*))) \leq 2 \log\left(\frac{d}{\delta}\right) \sqrt{\frac{\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}}{b}}, \tag{31}$$

and

$$\begin{aligned}
 & \lambda_{\max}((H_{\tilde{\lambda}}^{-1}(w^*) - H_{\tilde{\lambda}}^{-1}(w))(H_{\tilde{\lambda}}(w^*) - \bar{H}_{\tilde{\lambda}}(w^*))) \\
 & = \lambda_{\max}((H_{\tilde{\lambda}}(w^*) - H_{\tilde{\lambda}}(w))H_{\tilde{\lambda}}^{-1}(w)H_{\tilde{\lambda}}^{-1}(w^*)(H_{\tilde{\lambda}}(w^*) - \bar{H}_{\tilde{\lambda}}(w^*))) \\
 & \leq \lambda_{\max}(H_{\tilde{\lambda}}(w^*) - H_{\tilde{\lambda}}(w))2 \log\left(\frac{d}{\delta}\right) \sqrt{\frac{\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}}{b}} \cdot \frac{1}{\tilde{\lambda}} \\
 & \leq \frac{M \|w - w^*\|}{\tilde{\lambda}} \cdot 2 \log\left(\frac{d}{\delta}\right) \sqrt{\frac{\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}}{b}}, \tag{32}
 \end{aligned}$$

using the results of Lemma (9), and also Lipschitz Hessian property (Condition 1). Moreover,

$$\lambda_{\max}(H_{\tilde{\lambda}}^{-1}(w)(H_{\tilde{\lambda}}(w) - H_{\tilde{\lambda}}(w^*))) \leq \frac{\lambda_{\max}(H_{\tilde{\lambda}}(w) - H_{\tilde{\lambda}}(w^*))}{\tilde{\lambda}} \leq \frac{M \|w - w^*\|}{\tilde{\lambda}}. \quad (33)$$

Combining (29), (31), (32), (33) we have with probability at least $1 - \delta$

$$\min_{\|w - w^*\| \leq R} \lambda_{\min}(\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*)H_{\tilde{\lambda}}(w)\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*)) \geq \frac{\lambda}{\tilde{\lambda}} \cdot \frac{1}{1 + 2 \log\left(\frac{d}{\delta}\right) \sqrt{\frac{\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}}{\lambda}} + 2 \log\left(\frac{d}{\delta}\right) \sqrt{\frac{\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}}{\lambda}} \cdot \frac{MR}{\tilde{\lambda}} + \frac{MR}{\tilde{\lambda}}},$$

which finishes the proof. \blacksquare

A.6 Proof of Lemma 11

Proof First we perform the following decomposition:

$$\begin{aligned} \max_{i \in [n]} \max_{\|w - w^*\| \leq R} \ell''(w^\top x_i; b_i) x_i^\top \bar{H}_{\tilde{\lambda}}^{-1}(w^*) x_i &\leq \max_{i \in [n]} \max_{\|w - w^*\| \leq R} |\ell''(w^\top x_i; b_i) - \ell''(w^{*\top} x_i; b_i)| x_i^\top \bar{H}_{\tilde{\lambda}}^{-1}(w^*) x_i \\ &\quad + \max_{i \in [n]} \ell''(w^{*\top} x_i; b_i) x_i^\top \bar{H}_{\tilde{\lambda}}^{-1}(w^*) x_i, \end{aligned}$$

where

$$\begin{aligned} \max_{i \in [n]} \max_{\|w - w^*\| \leq R} |\ell''(w^\top x_i; b_i) - \ell''(w^{*\top} x_i; b_i)| x_i^\top \bar{H}_{\tilde{\lambda}}^{-1}(w^*) x_i &\leq \max_{i \in [n]} \bar{M} R x_i^\top \bar{H}_{\tilde{\lambda}}^{-1}(w^*) x_i \\ &\leq \frac{\bar{M} R D}{\tilde{\lambda}} = \frac{MR}{\tilde{\lambda}}. \end{aligned}$$

Thus we have

$$\max_{i \in [n]} \max_{\|w - w^*\| \leq R} \ell''(w^\top x_i; b_i) x_i^\top \bar{H}_{\tilde{\lambda}}^{-1}(w^*) x_i \leq \max_{i \in [n]} \ell''(w^{*\top} x_i; b_i) x_i^\top \bar{H}_{\tilde{\lambda}}^{-1}(w^*) x_i + \frac{MR}{\tilde{\lambda}}.$$

Then we relate $x_i^\top \bar{H}_{\tilde{\lambda}}^{-1}(w^*) x_i$ to $x_i^\top H_{\tilde{\lambda}}^{-1}(w^*) x_i$, for this purpose, we first upper bound $\lambda_{\max}\left(\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*)H_{\tilde{\lambda}}(w^*)\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*)\right)$. By Lemma 9, when $b \geq 16\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}} \log^2(d/\delta)$ we have

$$\left\| H_{\tilde{\lambda}}^{-1}(w^*)(H_{\tilde{\lambda}}(w^*) - \bar{H}_{\tilde{\lambda}}(w^*)) \right\|_2 \leq 2 \log\left(\frac{d}{\delta}\right) \sqrt{\frac{\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}}{b}} \leq \frac{1}{2},$$

thus

$$\begin{aligned} \lambda_{\max}\left(\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*)H_{\tilde{\lambda}}(w^*)\bar{H}_{\tilde{\lambda}}^{-1/2}(w^*)\right) &= \frac{1}{\lambda_{\min}\left(H_{\tilde{\lambda}}^{-1}(w^*)\bar{H}_{\tilde{\lambda}}(w^*)\right)} \\ &\leq \frac{1}{1 - \left\| H_{\tilde{\lambda}}^{-1}(w^*)(H_{\tilde{\lambda}}(w^*) - \bar{H}_{\tilde{\lambda}}(w^*)) \right\|_2} \\ &\leq \frac{1}{1 - 1/2} \leq 2. \end{aligned}$$

Then we bound $\max_{i \in [n]} \ell''(w^{*\top} x_i; b_i) x_i^\top \bar{H}_\lambda^{-1}(w^*) x_i$ through $\max_{i \in [n]} \ell''(w^{*\top} x_i; b_i) x_i^\top H_\lambda^{-1}(w^*) x_i$, because

$$\begin{aligned} x_i^\top \bar{H}_\lambda^{-1}(w^*) x_i &= x_i^\top H_\lambda^{-1}(w^*) x_i + x_i^\top \left(\bar{H}_\lambda^{-1}(w^*) - H_\lambda^{-1}(w^*) \right) x_i \\ &= x_i^\top H_\lambda^{-1}(w^*) x_i + x_i^\top H_\lambda^{-1/2}(w^*) \left(H_\lambda^{1/2}(w^*) \bar{H}_\lambda^{-1}(w^*) H_\lambda^{1/2}(w^*) - I \right) H_\lambda^{-1/2}(w^*) x_i \\ &\leq x_i^\top H_\lambda^{-1}(w^*) x_i + \lambda_{\max} \left(H_\lambda^{1/2}(w^*) \bar{H}_\lambda^{-1}(w^*) H_\lambda^{1/2}(w^*) - I \right) x_i^\top H_\lambda^{-1}(w^*) x_i \\ &\stackrel{\textcircled{1}}{\leq} 2x_i^\top H_\lambda^{-1}(w^*) x_i, \end{aligned}$$

where in step $\textcircled{1}$ we used the fact that

$$\begin{aligned} &\lambda_{\max} \left(H_\lambda^{1/2}(w^*) \bar{H}_\lambda^{-1}(w^*) H_\lambda^{1/2}(w^*) - I \right) \\ &\leq \max \left\{ \left| \lambda_{\max} \left(H_\lambda^{1/2}(w^*) \bar{H}_\lambda^{-1}(w^*) H_\lambda^{1/2}(w^*) \right) - 1 \right|, \left| \lambda_{\min} \left(H_\lambda^{1/2}(w^*) \bar{H}_\lambda^{-1}(w^*) H_\lambda^{1/2}(w^*) \right) - 1 \right| \right\} \\ &= \max \left\{ \left| \lambda_{\max} \left(\bar{H}_\lambda^{-1}(w^*) H_\lambda(w^*) \right) - 1 \right|, \left| \lambda_{\min} \left(\bar{H}_\lambda^{-1}(w^*) H_\lambda(w^*) \right) - 1 \right| \right\} \leq 1. \end{aligned}$$

Based on the definition of effective dimension (Definition 6) and condition of bounded statistical leverage (Assumption 8), we can bound the smoothness as

$$\begin{aligned} \max_{i \in [n]} \max_{\|w-w^*\| \leq R} \ell''(w^\top x_i; b_i) x_i^\top \bar{H}_\lambda^{-1}(w^*) x_i &\leq 2 \max_{i \in [n]} \ell''(w^{*\top} x_i; b_i) x_i^\top H_\lambda^{-1}(w^*) x_i + \frac{MR}{\lambda} \\ &\leq 2\rho_\lambda^2 \left(\frac{1}{n} \sum_{i=1}^n \ell''(w^{*\top} x_i; b_i) x_i^\top H_\lambda^{-1}(w^*) x_i \right) + \frac{MR}{\lambda} \\ &\stackrel{\textcircled{1}}{=} 2\rho_\lambda^2 \left(\frac{1}{n} \sum_{i=1}^n \text{tr}(\ell''(w^{*\top} x_i; b_i) x_i x_i^\top H_\lambda^{-1}(w^*)) \right) + \frac{MR}{\lambda}, \end{aligned}$$

where in step $\textcircled{1}$ we have used the fact that $\text{tr}(ABC) = \text{tr}(CAB)$ for any A, B, C . Furthermore we have

$$\begin{aligned} 2\rho_\lambda^2 \left(\frac{1}{n} \sum_{i=1}^n \text{tr}(\ell''(w^{*\top} x_i; b_i) x_i x_i^\top H_\lambda^{-1}(w^*)) \right) &= 2\rho_\lambda^2 \left(\text{tr} \left(\frac{1}{n} \sum_i \ell''(w^{*\top} x_i; b_i) x_i x_i^\top H_\lambda^{-1}(w^*) \right) \right) \\ &= 2\rho_\lambda^2 \left(\text{tr} \left(\left(\frac{1}{n} \sum_i \ell''(w^{*\top} x_i; b_i) x_i x_i^\top \right) H_\lambda^{-1}(w^*) \right) \right) \\ &= 2\rho_\lambda^2 \sum_{j=1}^d \frac{\lambda_j}{\lambda_j + \lambda} = 2\rho_\lambda^2 d_\lambda. \end{aligned}$$

Combining above analysis, we get

$$\max_{i \in [n]} \max_{\|w-w^*\| \leq R} \ell''(w^\top x_i; b_i) x_i^\top \bar{H}_\lambda^{-1}(w^*) x_i \leq 2\rho_\lambda^2 d_\lambda + \frac{MR}{\lambda},$$

which concludes the proof. \blacksquare

A.7 Proof of Theorem 13

Proof Since we choose $\tilde{\lambda}$ at the scale of $\tilde{\lambda} = \max\{\lambda, \frac{L}{b}\}$, then we can apply Lemma 10 and Lemma 11 to verify that the new condition number after “preconditioning” will be

$$\max\left\{4\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}, \frac{4L\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}}{\lambda b}\right\},$$

applying Corollary 26 we know as long as the inexactness condition (25) is satisfied, we requires $\mathcal{O}\left(\log\left(\frac{1}{\epsilon}\right)\right)$ full gradient evaluations and

$$\max\{C_{\text{sb}}, C_{\text{lb}}\} \cdot \frac{10}{9} \log(36) \log\left(\frac{1}{\epsilon}\right) \quad (34)$$

total calls of approximate minimization of (8) to ensure $\mathbb{E}f(\tilde{w}_s) - f(w^*) \leq \epsilon$, where the factors $C_{\text{sb}}, C_{\text{lb}}$ are

$$\begin{aligned} C_{\text{sb}} &= \mathcal{O}\left(\max\left\{\frac{240}{b}, \frac{240L}{\lambda b^2}\right\} \cdot \rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}\right), \\ C_{\text{lb}} &= \mathcal{O}\left(\max\left\{2\sqrt{6}, 2\sqrt{6}\sqrt{\frac{L}{\lambda b}}\right\} \cdot \sqrt{\rho_{\tilde{\lambda}}^2 d_{\tilde{\lambda}}}\right), \end{aligned}$$

which represent the cases of small and large minibatch sizes, respectively. Since $\tilde{\lambda} \geq \frac{L}{b}$, then the condition number of (8) is never larger than $\frac{L}{\tilde{\lambda}} \leq L \cdot \frac{b}{L} = b$, we know from Lemma 36, when applying SVRG to solve (8), to reach some point of which objective the suboptimality of (8) satisfies

$$\tilde{f}_t(w_t) - \min_w \tilde{f}_t(w) \leq \frac{1}{10^5} \cdot \left(\frac{\lambda}{L}\right)^7 \epsilon, \quad (35)$$

the following number of gradient calls sufficient:

$$C \cdot b \cdot \log\left(\frac{L}{\lambda}\right) \cdot \log\left(\frac{1}{\epsilon}\right), \quad (36)$$

for some universal constant C . On the other hand, for every subproblem (8), if (35) is satisfied, applying Lemma 4 and Lemma 33 we know the gradient error in IMBA-SVRG (Algorithm 3) in the preconditioned space can be upper bounded by

$$\|\xi_t\|^2 \leq \left(\frac{\lambda}{L}\right)^6 \cdot \frac{2\epsilon}{10^5} + \frac{\eta^2 M^2 R^2 L}{\tilde{\lambda}^4} (f(\tilde{w}_{s-1}) - f(w^*)),$$

thus condition (25) in Theorem 24 is satisfied. Combining (34) and (36) we know the total number of gradient calls:

$$n \cdot \log\left(\frac{1}{\epsilon}\right) + \frac{10C}{9} \log(36) \cdot \max\{C_{\text{sb}}, C_{\text{lb}}\} \cdot b \cdot \log\left(\frac{L}{\lambda}\right) \cdot \log^2\left(\frac{1}{\epsilon}\right) \quad (37)$$

is sufficient to obtain a solution such that $\mathbb{E}f(\tilde{w}_s) - f(w^*) \leq \epsilon$ is satisfied. Next we choose b such that the total iteration complexity of above expression is minimized. Start from here we will ignore the constants before these factors for simplicity, we know the term $\max\{C_{\text{sb}}, C_{\text{lb}}\} \cdot b$ is of order

$$\mathcal{O}\left(\max\left\{\rho_\lambda^2 d_\lambda, \frac{\rho_\lambda^2 d_\lambda L}{\lambda b}, \sqrt{\frac{\rho_\lambda^2 d_\lambda L b}{\lambda}}\right\}\right),$$

when $n \gtrsim \rho_\lambda^2 d_\lambda \left(\frac{L}{\lambda}\right)^{1/3}$, we can choose $b \asymp \rho_\lambda^2 d_\lambda \left(\frac{L}{\lambda}\right)^{1/3}$, then $\max\{C_{\text{sb}}, C_{\text{lb}}\} \cdot b$ is of order

$$\mathcal{O}\left(\max\left\{\rho_\lambda^2 d_\lambda, \left(\frac{L}{\lambda}\right)^{2/3}, \rho_\lambda^2 d_\lambda \left(\frac{L}{\lambda}\right)^{2/3}\right\}\right) = \mathcal{O}\left(\rho_\lambda^2 d_\lambda \left(\frac{L}{\lambda}\right)^{2/3}\right), \quad (38)$$

when $n \lesssim \rho_\lambda^2 d_\lambda \left(\frac{L}{\lambda}\right)^{1/3}$, we can choose $b \asymp n$, in which case $\max\{C_{\text{sb}}, C_{\text{lb}}\} \cdot b$ can be upper bounded by

$$\mathcal{O}\left(\max\left\{\frac{\rho_\lambda^2 d_\lambda L}{\lambda n}, \sqrt{\frac{\rho_\lambda^2 d_\lambda L n}{\lambda}}\right\}\right) = \mathcal{O}\left(\frac{\rho_\lambda^2 d_\lambda L}{\lambda n}\right), \quad (39)$$

Combining (37), (38) and (39), we know the total number of individual function gradient calls to reach ϵ -suboptimality is

$$\mathcal{O}\left(\tilde{\kappa} \cdot \log\left(\frac{L}{\lambda}\right) \cdot \log^2\left(\frac{1}{\epsilon}\right) + n \cdot \log\left(\frac{1}{\epsilon}\right)\right),$$

where

$$\tilde{\kappa} = \max\left\{\rho_\lambda^2 d_\lambda \left(\frac{L}{\lambda}\right)^{2/3}, \frac{\rho_\lambda^2 d_\lambda L}{\lambda n}\right\},$$

which finishes the proof. ■

A.8 Proof of Theorem 22

Proof Applying the theory of catalyst acceleration (Lemma 37) we know only $\mathcal{O}\left(\left(\sqrt{\frac{\lambda+\gamma}{\lambda}}\right) \log\left(\frac{1}{\epsilon}\right)\right)$ calls of MB-SVRP is sufficient to reach ϵ -objective suboptimality, as long as each iterate w_r satisfies

$$f(w_r) + \frac{\gamma}{2} \|w_r - z_{r-1}\|^2 \leq \min_w f(w) + \frac{\gamma}{2} \|w - z_{r-1}\|^2 + \frac{\lambda\epsilon}{3600(\lambda+\gamma)} \left(1 - \frac{9}{10} \sqrt{\frac{\lambda}{\lambda+\gamma}}\right).$$

Moreover, according to Theorem 13, the iteration complexity of solving a $\lambda + \gamma$ strongly convex problem (23) is

$$\tilde{\mathcal{O}}\left(n + \max\left\{\rho_\lambda^2 d_\lambda \left(\frac{L}{(\lambda+\gamma)}\right)^{2/3}, \frac{\rho_\lambda^2 d_\lambda L}{(\lambda+\gamma)n}\right\}\right),$$

combining these two results we know the total iteration complexity of Algorithm 2 can be upper bounded by

$$\tilde{\mathcal{O}} \left(\left(\sqrt{\frac{\lambda + \gamma}{\lambda}} \right) \left(n + \max \left\{ \left(\frac{\rho_\lambda^2 d_\lambda L}{\lambda + \gamma} \right)^{2/3}, \frac{\rho_\lambda^2 d_\lambda L}{(\lambda + \gamma)n} \right\} \right) \right),$$

When $\frac{\rho_\lambda^2 d_\lambda L}{\lambda} \geq n^{3/2}$, if we choose

$$\gamma = \frac{\rho_\lambda^2 d_\lambda L}{n^{3/2}} - \lambda,$$

we obtain the iteration complexity can be upper bounded by

$$\tilde{\mathcal{O}} \left(\left(\sqrt{\frac{\rho_\lambda^2 d_\lambda L}{n^{3/2} \lambda}} \right) \cdot n \right) = \tilde{\mathcal{O}} \left(\sqrt{\rho_\lambda^2 d_\lambda} \cdot n^{1/4} \left(\frac{L}{\lambda} \right)^{1/2} \right),$$

which concludes the proof. ■

Appendix B. Proofs for Section 4

In this appendix, we describe the stochastic estimation sequence approach for inexact accelerated minibatch SVRG in Section 4, and provide the detailed analysis.

B.1 Stochastic Estimation Sequences

We introduce the following definition.

Definition 27 (Stochastic estimation sequence) A sequence of pairs $\{V_t(w), \theta_t\}_{t \geq 0}$ is called an estimation sequence of the function $f(w)$ if $\theta_t > 0$ and for any $w \in \mathbb{R}^d, t \geq 0$ we have

$$V_t(w) \leq (1 - \theta_t)f(w) + \theta_t V_0(w), \quad (40)$$

if $\{V_t(w)\}_{t \geq 0}$ is a sequence of random functions, we call the sequence of pairs $\{V_t(w), \theta_t\}_{t \geq 0}$ a stochastic estimation sequence if (40) holds in expectation, i.e.

$$\mathbb{E}[V_t(w)] \leq (1 - \theta_t)f(w) + \theta_t V_0(w).$$

The following lemma is a generalized version of Lemma 1 in (Lin et al., 2014), which shows if we can construct upper bound of $\mathbb{E}[f(w_t)]$ using $\mathbb{E}[V_t(w)]$, then we get convergence of $\mathbb{E}[f(w_t)]$.

Lemma 28 Suppose $\{V_t(w), \theta_t\}_{t \geq 0}$ is a stochastic estimation sequence of the function $f(w)$. Let w^* be the minimizer of $f(w)$. If there are sequences of random variables $\{w_t\}_{t \geq 0}$ and $\{\epsilon_t\}_{t \geq 0}$ in \mathbb{R}^d , $\{\delta_t\}_{t \geq 0}$ in \mathbb{R} such that

$$\mathbb{E}[f(w_t)] \leq \min_w \{\mathbb{E}[V_t(w)]\} + \mathbb{E}[\delta_t] \quad (41)$$

holds for all $t \geq 0$, then

$$\mathbb{E}[f(w_t)] - f(w^*) \leq \theta_t(V_0(w^*) - f(w^*)) + \mathbb{E}[\delta_t].$$

The following lemma constructs a concrete stochastic estimation sequence.

Lemma 29 *Assume $f(w) = \frac{1}{n} \sum_{i=1}^n f_i(w)$, where each $f_i(w)$ is λ -strongly convex and L -smooth. Suppose that*

- $V_0(w)$ is an arbitrary deterministic function on \mathbb{R}^d ;
- $\{\alpha_t\}_{t \geq 0}$ is a sequence that satisfies $\alpha_t \in (0, 1), \forall t \geq 0$, and $\sum_{t=0}^{\infty} \alpha_t = \infty$.
- $\{y_t\}_{t=0}^{\infty}$ is an arbitrary sequence in \mathbb{R}^d ;
- Define $\{v_t\}_{t \geq 1}$ as:

$$v_t = \frac{1}{b} \sum_{i \in B_t} \nabla f_i(y_{t-1}) - \frac{1}{b} \sum_{i \in B_t} \nabla f_i(\tilde{w}_{s-1}) + \frac{1}{n} \sum_{i=1}^n \nabla f_i(\tilde{w}_{s-1}).$$

Define the sequence $\{w_t\}_{t \geq 0}$ and $\{V_t(w)\}_{t \geq 0}$ as follows. Let $y_0 = w_0$ and w_t be arbitrary vector in \mathbb{R}^d such that

$$w_t = y_{t-1} - \eta v_t + \xi_t,$$

where the stepsize η satisfies $0 < \eta \leq \frac{1}{L}$. Let

$$\begin{aligned} V_t(w) = & (1 - \alpha_{t-1})V_{t-1}(w) \\ & + \alpha_{t-1} \left(\frac{1}{b} \sum_{i \in B_t} f_i(y_{t-1}) + \left\langle v_t - \frac{\xi_t}{\eta}, w - y_{t-1} \right\rangle + \frac{\lambda}{4} \|w - y_{t-1}\|^2 - \frac{\|\xi_t\|^2}{\lambda \eta^2} \right), \end{aligned} \quad (42)$$

let $\theta_0 = 1$, and $\theta_t = (1 - \alpha_{t-1})\theta_{t-1}, \forall t \geq 1$. Then the sequence $\{V_t(w), \theta_t\}_{t \geq 0}$ is a stochastic estimation sequence of $f(w)$.

By above construction, we know $V_t(w)$ is a quadratic function, its minimizer and minimum value has the following iterative form:

Lemma 30 *Define $\forall t \geq 0$:*

$$V_t^* = \min_w V_t(w),$$

suppose we choose the function $V_0(w)$ in Lemma 29 as

$$V_0(w) = V_0^* + \frac{\lambda}{4} \|w - z_0\|^2,$$

with $V_0^* = f(z_0), \epsilon_0 = 0$ and $\delta_0 = 0$ for $z_0 = w_0$. Then the sequence $\{V_t(w)\}_{t \geq 0}$ defined in 42 can be written as

$$V_t(w) = V_t^* + \frac{\lambda}{4} \|w - z_t\|^2, \quad (43)$$

where the sequences $\{V_t^*\}, \{z_t\}$ are defined recursively as:

$$z_t = (1 - \alpha_{t-1})z_{t-1} + \alpha_{t-1}y_{t-1} - \frac{2\alpha_{t-1}}{\lambda} \left(v_t - \frac{\xi_t}{\eta} \right), \quad (44)$$

$$\begin{aligned} V_t^* = & (1 - \alpha_{t-1})V_{t-1}^* + \frac{\alpha_{t-1}(1 - \alpha_{t-1})\lambda}{4} \|z_{t-1} - y_{t-1}\|^2 + \alpha_{t-1}(1 - \alpha_{t-1}) \left\langle v_t - \frac{\xi_t}{\eta}, z_{t-1} - y_{t-1} \right\rangle \\ & - \frac{\alpha_{t-1}^2}{\lambda} \|v_t\|^2 + \frac{\alpha_{t-1}}{b} \sum_{i \in B_t} f_i(y_{t-1}) - \frac{(\alpha_{t-1} + \alpha_{t-1}^2) \|\xi_t\|^2}{\lambda\eta^2} + \frac{2\alpha_{t-1}^2}{\lambda\eta} \langle \xi_t, v_t \rangle. \end{aligned} \quad (45)$$

The following lemma establishes a connection between z_t, y_t, w_t when z_t is updated as (44), and y_t, w_t is updated as Algorithm 3.

Lemma 31 *Based on the updating rule of z_t in (44), if we choose $\alpha_t = \alpha = \sqrt{\frac{\lambda\eta}{2}}, \forall t \geq 0$, we have the following inequalities hold for all $t \geq 0$:*

$$z_t - y_t = \frac{1}{\alpha}(y_t - w_t).$$

The following lemma gives concrete expression of δ_t such that for the IMBA-SVRG algorithm, the condition (41) in Lemma 28 is satisfied.

Lemma 32 *Suppose we choose $\forall t \geq 0$*

$$\alpha_t = \alpha = \sqrt{\frac{\lambda\eta}{2}},$$

and the stepsize η satisfying $\eta \leq \frac{1}{8L}$, we have $\forall t \geq 0$

$$\mathbb{E}[f(w_t)] \leq \mathbb{E}[V_t^*] + \mathbb{E}[\delta_t],$$

where $\delta_0 = 0$ and $\forall t \geq 1$:

$$\delta_t = \left(1 - \sqrt{\frac{\lambda\eta}{2}} \right) \delta_{t-1} + \left(\eta \|v_t - \nabla f(y_{t-1})\|^2 - \frac{(1 - \sqrt{\lambda\eta/2})\lambda}{\sqrt{2\lambda\eta}} \|y_{t-1} - w_{t-1}\|^2 + \frac{1}{4\lambda^2\eta^3} \|\xi_t\|^2 \right).$$

The lemma below bounds the term $\mathbb{E} \left[\eta \|v_t - \nabla f(y_{t-1})\|^2 - \frac{(1 - \sqrt{\lambda\eta/2})\lambda}{\sqrt{2\lambda\eta}} \|y_{t-1} - w_{t-1}\|^2 \right]$.

Lemma 33 *Suppose B_t is constructed by uniform sampling with or without replacement, we have the following inequality holds:*

$$\begin{aligned} & \mathbb{E} \left[\eta \|v_t - \nabla f(y_{t-1})\|^2 - \frac{(1 - \sqrt{\lambda\eta/2})\lambda}{\sqrt{2\lambda\eta}} \|y_{t-1} - w_{t-1}\|^2 \right] \\ & \leq \frac{8\eta L}{b} (f(w_{t-1}) - f(w^*) + f(\tilde{w}_{s-1}) - f(w^*)) + \left(\frac{2\eta L^2}{b} - \frac{(1 - \sqrt{\lambda\eta/2})\lambda}{\sqrt{2\lambda\eta}} \right) \|y_{t-1} - w_{t-1}\|^2. \end{aligned}$$

Based on above lemma, by carefully choosing the stepsize η , we get the following lemma which is important to obtain iteration complexity for IMBA-SVRG.

Lemma 34 *If we choose the stepsize satisfying*

$$\eta \leq \min \left\{ \frac{b^2\lambda}{6400L^2}, \frac{1}{8L} \right\}$$

then the following inequality holds $\forall t \geq 0$:

$$\begin{aligned} \mathbb{E}f(w_t) - f(w^*) &\leq \left(1 - \sqrt{\frac{\lambda\eta}{2}}\right)^t (V_0(w^*) - f(w^*)) \\ &+ \mathbb{E} \left[\sum_{k=1}^t \left(1 - \sqrt{\frac{\lambda\eta}{2}}\right)^{t-k} \left(\frac{1}{4\lambda^2\eta^3} \|\xi_k\|^2 + \frac{8\eta L}{b} (f(w_{k-1}) + f(\tilde{w}_{s-1}) - 2f(w^*)) \right) \right], \end{aligned}$$

B.2 Proof of Lemma 28

Proof Since

$$\begin{aligned} \mathbb{E} [f(w_t)] &\leq \min_w \{\mathbb{E} [V_t(w)]\} + \mathbb{E} [\delta_t] \\ &\leq \mathbb{E} [V_t(w^*)] + \mathbb{E} [\delta_t] \\ &\leq (1 - \theta_t)f(w^*) + \theta_t V_0(w^*) + \mathbb{E} [\delta_t], \end{aligned}$$

subtracting both sides by $f(w^*)$ finishes the proof. ■

B.3 Proof of Lemma 29

Proof We proceed the proof by induction. When $t = 0$, we have

$$V_0(w) = (1 - \theta_0)f(w) + \theta_0 V_0(w) = V_0(w),$$

suppose

$$\mathbb{E} [V_{t-1}(w)] \leq (1 - \theta_{t-1})f(w) + \theta_{t-1}V_0(w)$$

is true for some $t \geq 1$, then for $\mathbb{E}[V_t(w)]$, since $\mathbb{E}[v_t] = \nabla f(y_{t-1})$, we have

$$\begin{aligned}
 \mathbb{E}[V_t(w)] &= (1 - \alpha_{t-1})\mathbb{E}[V_{t-1}(w)] \\
 &\quad + \alpha_{t-1}\mathbb{E}\left[\left(\frac{1}{b}\sum_{i \in B_t} f_i(y_{t-1}) + \left\langle v_t - \frac{\xi_t}{\eta}, w - y_{t-1} \right\rangle + \frac{\lambda}{4}\|w - y_{t-1}\|^2 - \frac{\|\xi_t\|^2}{\lambda\eta^2}\right)\right] \\
 &= (1 - \alpha_{t-1})\mathbb{E}[V_{t-1}(w)] \\
 &\quad + \alpha_{t-1}\left(f(y_{t-1}) + \langle \nabla f(y_{t-1}), w - y_{t-1} \rangle + \frac{\lambda}{4}\|w - y_{t-1}\|^2 - \left\langle \frac{\xi_t}{\eta}, w - y_{t-1} \right\rangle - \frac{\|\xi_t\|^2}{\lambda\eta^2}\right) \\
 &\stackrel{\textcircled{1}}{\leq} (1 - \alpha_{t-1})\mathbb{E}[V_{t-1}(w)] + \alpha_{t-1}\left(f(y_{t-1}) + \langle \nabla f(y_{t-1}), w - y_{t-1} \rangle + \frac{\lambda}{2}\|w - y_{t-1}\|^2\right) \\
 &\stackrel{\textcircled{2}}{\leq} (1 - \alpha_{t-1})(1 - \theta_{t-1})f(w) + (1 - \alpha_{t-1})\theta_{t-1}V_0(w) + \alpha_{t-1}f(w) \\
 &= (1 - (1 - \alpha_{t-1})\theta_{t-1})f(w) + (1 - \alpha_{t-1})\theta_{t-1}V_0(w) \\
 &= (1 - \theta_t)f(w) + \theta_t V_0(w),
 \end{aligned}$$

which concludes the proof, where at step $\textcircled{1}$ we used the inequality $-\langle v_1, v_2 \rangle - \|v_2\|^2 \leq \frac{\|v_1\|^2}{4}$, at step $\textcircled{2}$ we used the inductive hypothesis and the fact that $f(w)$ is λ -strongly convex. \blacksquare

B.4 Proof of Lemma 30

Proof We proceed the proof by induction, when $t = 0$ this is true by construction. Suppose for some $t \geq 1$ the following holds:

$$V_{t-1}(w) = V_{t-1}^* + \frac{\lambda}{4}\|w - z_{t-1}\|^2,$$

then at time t , we have

$$\begin{aligned}
 V_t(w) &= (1 - \alpha_{t-1})V_{t-1}(w) \\
 &\quad + \alpha_{t-1}\left(\frac{1}{b}\sum_{i \in B_t} f_i(y_{t-1}) + \left\langle v_t - \frac{\xi_t}{\eta}, w - y_{t-1} \right\rangle + \frac{\lambda}{4}\|w - y_{t-1}\|^2 - \frac{\|\xi_t\|^2}{\lambda\eta^2}\right) \\
 &= (1 - \alpha_{t-1})V_{t-1}^* + \frac{(1 - \alpha_{t-1})\lambda}{4}\|w - z_{t-1}\|^2 + \frac{\alpha_{t-1}\lambda}{4}\|w - y_{t-1}\|^2 + \alpha_{t-1}\left\langle v_t - \frac{\xi_t}{\eta}, w \right\rangle \\
 &\quad + \frac{\alpha_{t-1}}{b}\sum_{i \in B_t} f_i(y_{t-1}) - \alpha_{t-1}\left\langle v_t - \frac{\xi_t}{\eta}, y_{t-1} \right\rangle - \frac{\alpha_{t-1}\|\xi_t\|^2}{\lambda\eta^2},
 \end{aligned}$$

by first order optimality condition, it is clear that the minimizer of $V_t(w)$: z_t satisfies the following:

$$(1 - \alpha_{t-1})\lambda(z_t - z_{t-1}) + \alpha_{t-1}\lambda(z_t - y_{t-1}) + 2\alpha_{t-1}\left(v_t - \frac{\xi_t}{\eta}\right) = 0,$$

from above we obtain the recursive form of z_t as (44) defines. Plug in (44) to (42) we get

$$\begin{aligned}
 V_t^* = V_t(z_t) &= (1 - \alpha_{t-1})V_{t-1}^* + \frac{(1 - \alpha_{t-1})\lambda}{4} \|z_t - z_{t-1}\|^2 + \frac{\alpha_{t-1}\lambda}{4} \|z_t - y_{t-1}\|^2 + \alpha_{t-1} \left\langle v_t - \frac{\xi_t}{\eta}, z_t \right\rangle \\
 &\quad + \frac{\alpha_{t-1}}{b} \sum_{i \in B_t} f_i(y_{t-1}) - \alpha_{t-1} \left\langle v_t - \frac{\xi_t}{\eta}, y_{t-1} \right\rangle - \frac{\alpha_{t-1} \|\xi_t\|^2}{\lambda\eta^2} \\
 &= (1 - \alpha_{t-1})V_{t-1}^* + \frac{(1 - \alpha_{t-1})\lambda}{4} \left\| \alpha_{t-1}(y_{t-1} - z_{t-1}) - \frac{2\alpha_{t-1}}{\lambda} \left(v_t - \frac{\xi_t}{\eta} \right) \right\|^2 + \frac{\alpha_{t-1}}{b} \sum_{i \in B_t} f_i(y_{t-1}) \\
 &\quad + \frac{\alpha_{t-1}\lambda}{4} \left\| (1 - \alpha_{t-1})(z_{t-1} - y_{t-1}) - \frac{2\alpha_{t-1}}{\lambda} \left(v_t - \frac{\xi_t}{\eta} \right) \right\|^2 - \frac{\alpha_{t-1} \|\xi_t\|^2}{\lambda\eta^2} \\
 &\quad + \alpha_{t-1} \left\langle v_t - \frac{\xi_t}{\eta}, (1 - \alpha_{t-1})(z_{t-1} - y_{t-1}) - \frac{2\alpha_{t-1}}{\lambda} \left(v_t - \frac{\xi_t}{\eta} \right) \right\rangle \\
 &= (1 - \alpha_{t-1})V_{t-1}^* + \frac{\alpha_{t-1}(1 - \alpha_{t-1})\lambda}{4} \|z_{t-1} - y_{t-1}\|^2 + \alpha_{t-1}(1 - \alpha_{t-1}) \left\langle v_t - \frac{\xi_t}{\eta}, z_{t-1} - y_{t-1} \right\rangle \\
 &\quad - \frac{\alpha_{t-1}^2}{\lambda} \|v_t\|^2 + \frac{\alpha_{t-1}}{b} \sum_{i \in B_t} f_i(y_{t-1}) - \frac{(\alpha_{t-1} + \alpha_{t-1}^2) \|\xi_t\|^2}{\lambda\eta^2} + \frac{2\alpha_{t-1}^2}{\lambda\eta} \langle \xi_t, v_t \rangle,
 \end{aligned}$$

which verified (45). ■

B.5 Proof of Lemma 31

Proof We prove by induction, when $t = 0$ it is obviously true, suppose it is true for $t - 1, \forall t \geq 1$, i.e.

$$z_{t-1} - y_{t-1} = \frac{1}{\alpha}(y_{t-1} - w_{t-1}),$$

for iteration t , based on (44) we have

$$\begin{aligned}
 z_t - y_t &= (1 - \alpha)z_{t-1} + \alpha y_{t-1} - \frac{2\alpha}{\lambda} \left(v_t - \frac{\xi_t}{\eta} \right) - y_t \\
 &= (1 - \alpha)(z_{t-1} - y_{t-1}) + y_{t-1} - \frac{2\alpha}{\lambda} \left(v_t - \frac{\xi_t}{\eta} \right) - y_t \\
 &\stackrel{\textcircled{1}}{=} \frac{1 - \alpha}{\alpha}(y_{t-1} - w_{t-1}) + y_{t-1} - \frac{2\alpha}{\lambda} \left(v_t - \frac{\xi_t}{\eta} \right) - y_t \\
 &= \frac{1}{\alpha} \left(y_{t-1} - \frac{2\alpha^2}{\lambda} \left(v_t - \frac{\xi_t}{\eta} \right) \right) - \frac{1 - \alpha}{\alpha} w_{t-1} - y_t \\
 &\stackrel{\textcircled{2}}{=} \frac{1}{\alpha} w_t - \frac{1 - \alpha}{\alpha} w_{t-1} - y_t \\
 &\stackrel{\textcircled{3}}{=} \frac{1}{\alpha} \left(w_t + \frac{1 - \alpha}{1 + \alpha} (w_t - w_{t-1}) - w_t \right) \\
 &\stackrel{\textcircled{4}}{=} \frac{1}{\alpha} (y_t - w_t),
 \end{aligned}$$

which concludes the proof, where step ① used the inductive hypothesis, step ② used the update rule of w_t in Algorithm 3, step ③ and ④ used the update rule of y_t in Algorithm 3. ■

B.6 Proof of Lemma 32

Proof We prove by induction, when $t = 0$, it is true that

$$f(w_0) \leq V_0^* = f(z_0) = f(w_0),$$

suppose

$$\mathbb{E}[f(w_{t-1})] \leq \mathbb{E}[V_{t-1}^*] + \mathbb{E}[\delta_{t-1}], \quad (46)$$

for some $t \geq 1$, then based on smoothness, we know

$$f(w_t) \leq f(y_{t-1}) + \langle \nabla f(y_{t-1}), w_t - y_{t-1} \rangle + \frac{L}{2} \|w_t - y_{t-1}\|^2,$$

thus

$$\begin{aligned} V_t^* &= (1 - \alpha_{t-1})V_{t-1}^* + \frac{\alpha_{t-1}(1 - \alpha_{t-1})\lambda}{4} \|z_{t-1} - y_{t-1}\|^2 \\ &\quad + \alpha_{t-1}(1 - \alpha_{t-1}) \left\langle v_t - \frac{\xi_t}{\eta}, z_{t-1} - y_{t-1} \right\rangle \\ &\quad - \frac{\alpha_{t-1}^2}{\lambda} \|v_t\|^2 + \frac{\alpha_{t-1}}{b} \sum_{i \in B_t} f_i(y_{t-1}) - \frac{(\alpha_{t-1} + \alpha_{t-1}^2) \|\xi_t\|^2}{\lambda\eta^2} + \frac{2\alpha_{t-1}^2}{\lambda\eta} \langle \xi_t, v_t \rangle. \end{aligned}$$

Hence we have

$$\begin{aligned} \mathbb{E}[f(w_t) - V_t^*] &\leq \mathbb{E} \left[f(y_{t-1}) + \langle \nabla f(y_{t-1}), w_t - y_{t-1} \rangle + \frac{L}{2} \|w_t - y_{t-1}\|^2 \right] \\ &\quad - \mathbb{E} \left[(1 - \alpha)V_{t-1}^* - \frac{\alpha(1 - \alpha)\lambda}{4} \|z_{t-1} - y_{t-1}\|^2 - \alpha(1 - \alpha) \left\langle v_t - \frac{\xi_t}{\eta}, z_{t-1} - y_{t-1} \right\rangle \right] \\ &\quad + \mathbb{E} \left[\frac{\alpha^2}{\lambda} \|v_t\|^2 - \frac{\alpha}{b} \sum_{i \in B_t} f_i(y_{t-1}) + \frac{(\alpha + \alpha^2) \|\xi_t\|^2}{\lambda\eta^2} - \frac{2\alpha^2}{\lambda\eta} \langle \xi_t, v_t \rangle \right] \\ &\stackrel{\text{①}}{=} \mathbb{E} \left[(1 - \alpha)(f(y_{t-1}) - V_{t-1}^* + \langle v_t, w_{t-1} - y_{t-1} \rangle) - \langle \nabla f(y_{t-1}), \eta v_t - \xi_t \rangle + \frac{\eta}{2} \|v_t\|^2 \right] \\ &\quad + \mathbb{E} \left[\frac{L}{2} \|\eta v_t - \xi_t\|^2 - \frac{(1 - \alpha)\lambda}{4\alpha} \|y_{t-1} - w_{t-1}\|^2 + \frac{(\alpha + \alpha^2) \|\xi_t\|^2}{\lambda\eta^2} \right] \\ &\quad + \mathbb{E} \left[-\frac{2\alpha^2}{\lambda\eta} \langle \xi_t, v_t \rangle - (1 - \alpha) \left\langle \frac{\xi_t}{\eta}, w_{t-1} - y_{t-1} \right\rangle \right] \\ &\stackrel{\text{②}}{\leq} \mathbb{E} \left[(1 - \alpha)(f(y_{t-1}) - V_{t-1}^* + \langle v_t, w_{t-1} - y_{t-1} \rangle) - \langle \nabla f(y_{t-1}), \eta v_t - \xi_t \rangle + \frac{9\eta}{16} \|v_t\|^2 \right] \\ &\quad + \mathbb{E} \left[\frac{L}{2} \|\eta v_t - \xi_t\|^2 + \left(\frac{2\alpha}{\lambda\eta^2} + \frac{16\alpha^4}{\lambda^2\eta^3} + \frac{(1 - \alpha)}{2\lambda\eta^2} \right) \|\xi_t\|^2 \right] \\ &\quad + \mathbb{E} \left[\left(\frac{(1 - \alpha)\lambda}{2} - \frac{(1 - \alpha)\lambda}{4\alpha} \right) \|w_{t-1} - y_{t-1}\|^2 \right] \quad (47) \end{aligned}$$

where at step ① we used the fact of updating rule

$$w_t = y_{t-1} - \eta v_t + \xi_t,$$

as well as Lemma 31, and $\mathbb{E} \left[\frac{1}{b} \sum_{i \in B_t} f_i(y_{t-1}) \right] = f(y_{t-1})$; and at step ② we used the following two inequalities:

$$\begin{aligned} \left| (1-\alpha) \left\langle \frac{\xi_t}{\eta}, w_{t-1} - y_{t-1} \right\rangle \right| &\leq \frac{(1-\alpha)\lambda}{2} \|y_{t-1} - w_{t-1}\|^2 + \frac{(1-\alpha)}{2\lambda\eta^2} \|\xi_t\|^2, \\ \left| \frac{2\alpha^2}{\lambda\eta} \langle \xi_t, v_t \rangle \right| &\leq \frac{\eta}{16} \|v_t\|^2 + \frac{16\alpha^4}{\lambda^2\eta^3} \|\xi_t\|^2 \end{aligned}$$

Since

$$\begin{aligned} \mathbb{E} [f(y_{t-1}) - V_{t-1}^* + \langle v_t, w_{t-1} - y_{t-1} \rangle] &= \mathbb{E} [f(y_{t-1}) - V_{t-1}^* + \langle \nabla f(y_{t-1}), w_{t-1} - y_{t-1} \rangle] \\ &\stackrel{\text{①}}{\leq} \mathbb{E} \left[f(w_{t-1}) - \frac{\lambda}{2} \|w_{t-1} - y_{t-1}\|^2 - V_{t-1}^* \right] \\ &\stackrel{\text{②}}{\leq} \mathbb{E} \left[\delta_{t-1} - \frac{\lambda}{2} \|w_{t-1} - y_{t-1}\|^2 \right], \end{aligned} \quad (48)$$

where at step ① we used the λ -strong convexity of $f(w)$, at step ② we used the inductive hypothesis (46). Combining (47) and (48) we obtain

$$\begin{aligned} \mathbb{E} [f(w_t) - V_t^*] &\leq (1-\alpha)\delta_{t-1} - \mathbb{E} \left[\frac{(1-\alpha)\lambda}{4\alpha} \|y_{t-1} - w_{t-1}\|^2 \right] \\ &+ \mathbb{E} \left[\frac{L}{2} \|\eta v_t - \xi_t\|^2 + \frac{9\eta}{16} \|v_t\|^2 - \langle \nabla f(y_{t-1}), \eta v_t - \xi_t \rangle + \left(\frac{2\alpha}{\lambda\eta^2} + \frac{16\alpha^4}{\lambda^2\eta^3} + \frac{(1-\alpha)}{2\lambda\eta^2} \right) \|\xi_t\|^2 \right]. \end{aligned} \quad (49)$$

Next we bound the third term on the right hand side, since

$$\begin{aligned} &\frac{L}{2} \|\eta v_t - \xi_t\|^2 + \frac{9\eta}{16} \|v_t\|^2 - \langle \nabla f(y_{t-1}), \eta v_t + \xi_t \rangle \\ &\stackrel{\text{①}}{\leq} L\eta^2 \|v_t\|^2 + L \|\xi_t\|^2 + \frac{9\eta}{16} \|v_t\|^2 - \langle \nabla f(y_{t-1}), \eta v_t - \xi_t \rangle \\ &= \frac{\eta}{2} \|v_t - \nabla f(y_{t-1})\|^2 - \frac{\eta}{2} \|\nabla f(y_{t-1})\|^2 \\ &+ \left(L\eta^2 + \frac{\eta}{16} \right) \|v_t\|^2 + \langle \nabla f(y_{t-1}), \xi_t \rangle + L \|\xi_t\|^2 \\ &\stackrel{\text{②}}{\leq} \frac{\eta}{2} \|v_t - \nabla f(y_{t-1})\|^2 - \frac{\eta}{2} \|\nabla f(y_{t-1})\|^2 \\ &+ \left(2L\eta^2 + \frac{\eta}{8} \right) (\|v_t - \nabla f(y_{t-1})\|^2 + \|\nabla f(y_{t-1})\|^2) \\ &+ \frac{\eta}{8} \|\nabla f(y_{t-1})\|^2 + \frac{2}{\eta} \|\xi_t\|^2 + L \|\xi_t\|^2, \end{aligned}$$

where at step ① we used $\|v_1 + v_2\|^2 \leq 2(\|v_1\|^2 + \|v_2\|^2)$, and at step ② we used it again and also the inequality $2|\langle v_1, v_2 \rangle| \leq \|v_1\|^2 + \|v_2\|^2$. Thus we know when $\eta \leq \frac{1}{8L}$, we have

$$\frac{L}{2} \|\eta v_t - \xi_t\|^2 + \frac{9\eta}{16} \|v_t\|^2 - \langle \nabla f(y_{t-1}), \eta v_t + \xi_t \rangle \quad (50)$$

$$\begin{aligned} &\leq \left(\frac{\eta}{2} + 2L\eta^2 + \frac{\eta}{8} \right) \|v_t - \nabla f(y_{t-1})\|^2 + \left(L + \frac{2}{\eta} \right) \|\xi_t\|^2 \\ &\quad - \left(\frac{\eta}{2} - \frac{\eta}{8} - \frac{\eta}{8} - 2L\eta^2 \right) \|\nabla f(y_{t-1})\|^2 \\ &\leq \eta \|v_t - \nabla f(y_{t-1})\|^2 + \frac{3}{\eta} \|\xi_t\|^2. \end{aligned} \quad (51)$$

Combining (49) and (51) we obtain

$$\begin{aligned} \mathbb{E} [f(w_t) - V_t^*] &\leq (1 - \alpha)\delta_{t-1} + \mathbb{E} \left[\eta \|v_t - \nabla f(y_{t-1})\|^2 - \frac{(1 - \alpha)\lambda}{4\alpha} \|y_{t-1} - w_{t-1}\|^2 \right] \\ &\quad + \mathbb{E} \left[\left(\frac{2\alpha}{\lambda\eta^2} + \frac{16\alpha^4}{\lambda^2\eta^3} + \frac{(1 - \alpha)}{2\lambda\eta^2} + \frac{3}{\eta} \right) \|\xi_t\|^2 \right], \end{aligned} \quad (52)$$

also since

$$\alpha = \sqrt{\frac{\lambda\eta}{2}} \leq \frac{1}{4} \sqrt{\frac{\lambda}{L}} \leq \frac{1}{4},$$

we know

$$\frac{2\alpha}{\lambda\eta^2} + \frac{16\alpha^4}{\lambda^2\eta^3} + \frac{(1 - \alpha)}{2\lambda\eta^2} + \frac{3}{\eta} \leq \frac{1}{16\lambda^2\eta^3} + \frac{1}{16\lambda^2\eta^3} + \frac{1}{16\lambda^2\eta^3} + \frac{3}{64\lambda^2\eta^3} \leq \frac{1}{4\lambda^2\eta^3},$$

combining the inequality above and (52), then substituting $\alpha = \sqrt{\frac{\lambda\eta}{2}}$ finishes the proof. ■

B.7 Proof of Lemma 33

Proof The proof follows the strategy in (Johnson and Zhang, 2013; Nitanda, 2014). First, based on the minibatch sampling (here for simplicity we only consider sampling with replacement, for sampling without replacement, the bound below can be tightened by a factor

of $\frac{n-b}{n-1}$, see, e.g. Section 2.8 of (Lohr, 2009)) we know

$$\begin{aligned}
 \mathbb{E} \left[\|v_t - \nabla f(y_{t-1})\|^2 \right] &= \frac{1}{b^2} \mathbb{E} \left[\sum_{j \in B_t} \|\nabla f_j(y_{t-1}) - \nabla f(y_{t-1}) - (\nabla f_j(\tilde{w}_{s-1}) - \nabla f(\tilde{w}_{s-1}))\|^2 \right] \\
 &\stackrel{\textcircled{1}}{\leq} \frac{1}{b^2} \mathbb{E} \left[\sum_{j \in B_t} \|\nabla f_j(y_{t-1}) - \nabla f_j(\tilde{w}_{s-1})\|^2 \right] \\
 &\stackrel{\textcircled{2}}{\leq} \frac{2}{b^2} \mathbb{E} \left[\sum_{j \in B_t} (\|\nabla f_j(y_{t-1}) - \nabla f_j(w_{t-1})\|^2 + \|\nabla f_j(\tilde{w}_{s-1}) - \nabla f_j(w_{t-1})\|^2) \right] \\
 &\stackrel{\textcircled{3}}{\leq} \frac{4}{b^2} \mathbb{E} \left[\sum_{j \in B_t} (\|\nabla f_j(\tilde{w}_{s-1}) - \nabla f_j(w^*)\|^2 + \|\nabla f_j(w^*) - \nabla f_j(w_{t-1})\|^2) \right] \\
 &\quad + \frac{2L^2}{b} \|y_{t-1} - w_{t-1}\|^2 \\
 &\stackrel{\textcircled{4}}{\leq} \frac{8L}{b} (f(w_{t-1}) - f(w^*) + f(\tilde{w}_{s-1}) - f(w^*)) + \frac{2L^2}{b} \|y_{t-1} - w_{t-1}\|^2,
 \end{aligned}$$

where at step $\textcircled{1}$ we used $\mathbb{E} \|v - \mathbb{E}v\|^2 \leq \mathbb{E} \|v\|^2$, at step $\textcircled{2}$ we used $\|v_1 + v_2\|^2 \leq 2(\|v_1\|^2 + \|v_2\|^2)$, at step $\textcircled{3}$ we used it again along with the L -smoothness of $f_j(w)$, at step $\textcircled{4}$ we used standard results in SVRG analysis, e.g. Lemma 1 of (Xiao and Zhang, 2014). Substituting above into the term

$$\mathbb{E} \left[\eta \|v_t - \nabla f(y_{t-1})\|^2 - \frac{(1 - \sqrt{\lambda\eta/2})\lambda}{\sqrt{2\lambda\eta}} \|y_{t-1} - w_{t-1}\|^2 \right] \text{ concludes the proof. } \blacksquare$$

B.8 Proof of Lemma 34

Proof Combining Lemma 32 and 33, we have

$$\begin{aligned}
 \delta_t &= \left(1 - \sqrt{\frac{\lambda\eta}{2}} \right) \delta_{t-1} \tag{53} \\
 &\quad + \left(\eta \|v_t - \nabla f(y_{t-1})\|^2 - \frac{(1 - \sqrt{\lambda\eta/2})\lambda}{\sqrt{2\lambda\eta}} \|y_{t-1} - w_{t-1}\|^2 + \frac{1}{4\lambda^2\eta^3} \|\xi_t\|^2 \right) \\
 &\leq \left(1 - \sqrt{\frac{\lambda\eta}{2}} \right) \delta_{t-1} + \left(\frac{1}{4\lambda^2\eta^3} \|\xi_t\|^2 + \frac{8\eta L}{b} (f(w_{t-1}) + f(\tilde{w}_{s-1}) - 2f(w^*)) \right) \\
 &\quad + \left(\frac{2\eta L^2}{b} - \frac{(1 - \sqrt{\lambda\eta/2})\lambda}{\sqrt{2\lambda\eta}} \right) \|y_{t-1} - w_{t-1}\|^2. \tag{54}
 \end{aligned}$$

First we verify the factor before $\|y_{t-1} - w_{t-1}\|^2$ is non positive, since $1 - \eta\lambda \geq 1 - \frac{\lambda}{8L} \geq \frac{1}{2}$, then

$$\frac{2\eta L^2}{b} \cdot \frac{\sqrt{2\lambda\eta}}{(1 - \sqrt{\eta\lambda/2})\lambda} \leq \frac{4\eta^{3/2}L^2}{b\sqrt{\lambda}} \leq \min \left\{ \frac{b^2\lambda}{128000L}, \frac{L^{1/2}}{4b\sqrt{\lambda}} \right\} \leq 1,$$

where the last inequality is true because

$$\frac{b^2\lambda}{128000L} \left(\frac{L^{1/2}}{4b\sqrt{\lambda}} \right)^2 \leq 1,$$

thus we must have $\min \left\{ \frac{b^2\lambda}{128000L}, \frac{L^{1/2}}{4b\sqrt{\lambda}} \right\} \leq 1$ otherwise it leads to a contradiction. Thus $\frac{2\eta L^2}{b} - \frac{(1-\sqrt{\eta\lambda/2})\lambda}{\sqrt{2\eta\lambda}} \leq 0$, combining with (54), we have

$$\delta_t \leq \left(1 - \sqrt{\frac{\lambda\eta}{2}} \right) \delta_{t-1} + \left(\frac{1}{4\lambda^2\eta^3} \|\xi_t\|^2 + \frac{8\eta L}{b} (f(w_{t-1}) + f(\tilde{w}_{s-1}) - 2f(w^*)) \right),$$

applying above inequality recursively, we get

$$\delta_t \leq \sum_{k=1}^t \left(1 - \sqrt{\frac{\lambda\eta}{2}} \right)^{t-k} \left(\frac{1}{4\lambda^2\eta^3} \|\xi_k\|^2 + \frac{8\eta L}{b} (f(w_{k-1}) + f(\tilde{w}_{s-1}) - 2f(w^*)) \right),$$

then applying Lemma 28 and Lemma 32 concludes the proof. \blacksquare

B.9 Proof of Theorem 24

Proof We prove this theorem by induction, note that when $b \leq 20\sqrt{\frac{2L}{\lambda}}$, it is clear that

$$\frac{8\eta L}{b} \leq \frac{b^2\lambda}{6400L^2} \cdot \frac{8L}{b} = \frac{\lambda b}{800L} = \frac{\sqrt{\eta\lambda}}{10},$$

and when $b \geq 20\sqrt{\frac{2L}{\lambda}}$,

$$\frac{8\eta L}{b} = \frac{1}{b} \leq \frac{\sqrt{\lambda}}{20\sqrt{2L}} = \frac{\sqrt{\eta\lambda}}{10},$$

thus in both cases we have $\frac{8\eta L}{b} \leq \frac{\sqrt{\eta\lambda}}{10}$, thus

$$\begin{aligned} \sum_{k=1}^t \left(1 - \sqrt{\frac{\lambda\eta}{2}} \right)^{t-k} \left(\frac{8\eta L}{b} \right) &\leq \frac{1}{10} \sum_{k=1}^t \left(1 - \sqrt{\frac{\lambda\eta}{2}} \right)^{t-k} \sqrt{\lambda\eta} \\ &\leq \frac{1}{10} \sum_{k=1}^{\infty} \left(1 - \sqrt{\frac{\lambda\eta}{2}} \right)^k \sqrt{\lambda\eta} \leq \frac{\sqrt{2}}{10}. \end{aligned}$$

By Lemma 34 we know

$$\mathbb{E}f(w_t) - f(w^*) \leq \left(1 - \sqrt{\frac{\lambda\eta}{2}} \right)^t (V_0(w^*) - f(w^*)) + \frac{\sqrt{2}(f(\tilde{w}_{s-1}) - f(w^*))}{10} \quad (55)$$

$$+ \mathbb{E} \left[\sum_{k=1}^t \left(1 - \sqrt{\frac{\lambda\eta}{2}} \right)^{t-k} \left(\frac{1}{4\lambda^2\eta^3} \|\xi_k\|^2 + \frac{\sqrt{\lambda\eta}}{10} (f(w_{k-1}) - f(w^*)) \right) \right], \quad (56)$$

also when $\forall k \in [t]$,

$$\|\xi_k\|^2 \leq \frac{2\lambda^2\eta^3\sqrt{\lambda\eta}}{15}(f(\tilde{w}_{s-1}) - f(w^*)),$$

we know

$$\begin{aligned} \sum_{k=1}^t \left(1 - \sqrt{\frac{\lambda\eta}{2}}\right)^{t-k} \left(\frac{1}{4\lambda^2\eta^3} \|\xi_k\|^2\right) &\leq \frac{1}{30} \sum_{k=1}^t \left(1 - \sqrt{\frac{\lambda\eta}{2}}\right)^{t-k} \sqrt{\lambda\eta}(f(\tilde{w}_{s-1}) - f(w^*)) \\ &\leq \frac{1}{30} \sum_{k=1}^{\infty} \left(1 - \sqrt{\frac{\lambda\eta}{2}}\right)^k \sqrt{\lambda\eta}(f(\tilde{w}_{s-1}) - f(w^*)) \\ &\leq \frac{(f(\tilde{w}_{s-1}) - f(w^*))}{20}. \end{aligned} \quad (57)$$

Combining (56) and (57), we have

$$\begin{aligned} \mathbb{E}f(w_t) - f(w^*) &\leq \left(1 - \sqrt{\frac{\lambda\eta}{2}}\right)^t (V_0(w^*) - f(w^*)) + \frac{(f(\tilde{w}_{s-1}) - f(w^*))}{5} \\ &\quad + \mathbb{E} \left[\sum_{k=1}^t \left(1 - \sqrt{\frac{\lambda\eta}{2}}\right)^{t-k} \left(\frac{\sqrt{\lambda\eta}}{10}(f(w_{k-1}) - f(w^*))\right) \right], \end{aligned}$$

denote A_t be the right hand side on above inequality, i.e.

$$\begin{aligned} A_t &= \left(1 - \sqrt{\frac{\lambda\eta}{2}}\right)^t (V_0(w^*) - f(w^*)) + \frac{(f(\tilde{w}_{s-1}) - f(w^*))}{5} \\ &\quad + \mathbb{E} \left[\sum_{k=1}^t \left(1 - \sqrt{\frac{\lambda\eta}{2}}\right)^{t-k} \left(\frac{\sqrt{\lambda\eta}}{10}(f(w_{k-1}) - f(w^*))\right) \right] \\ &= \left(1 - \sqrt{\frac{\lambda\eta}{2}}\right) \left(\left(1 - \sqrt{\frac{\lambda\eta}{2}}\right)^{t-1} (V_0(w^*) - f(w^*)) + \frac{(f(\tilde{w}_{s-1}) - f(w^*))}{5} \right) \\ &\quad + \left(1 - \sqrt{\frac{\lambda\eta}{2}}\right) \mathbb{E} \left[\sum_{k=1}^{t-1} \left(1 - \sqrt{\frac{\lambda\eta}{2}}\right)^{t-1-k} \left(\frac{\sqrt{\lambda\eta}}{10}(f(w_{k-1}) - f(w^*))\right) \right] \\ &\quad + \frac{\sqrt{\lambda\eta}}{5\sqrt{2}}(f(\tilde{w}_{s-1}) - f(w^*)) + \frac{\sqrt{\lambda\eta}}{10}(f(w_{t-1}) - f(w^*)) \\ &\stackrel{\ominus}{\leq} \left(1 - \sqrt{\frac{\lambda\eta}{2}}\right) A_{t-1} + \frac{\sqrt{\lambda\eta}}{5\sqrt{2}}A_0 + \frac{\sqrt{\lambda\eta}}{10}A_{t-1} \\ &= \left(1 - \frac{9\sqrt{\lambda\eta}}{10}\right) A_{t-1} + \frac{\sqrt{\lambda\eta}}{5\sqrt{2}}A_0, \end{aligned}$$

where at step ① we used the fact that $f(\tilde{w}_{s-1}) - f(w^*) \leq A_0$, and $f(w_{t-1}) - f(w^*) \leq A_{t-1}$, applying above inequality recursively we get

$$\begin{aligned} A_t &\leq \left(1 - \frac{9\sqrt{\lambda\eta}}{10}\right)^t A_0 + \frac{\sqrt{\lambda\eta}}{5\sqrt{2}} \sum_{k=0}^{t-1} \left(1 - \frac{9\sqrt{\lambda\eta}}{10}\right)^k A_0 \\ &\leq \left(1 - \frac{9\sqrt{\lambda\eta}}{10}\right)^t A_0 + \frac{\sqrt{\lambda\eta}}{5\sqrt{2}} \sum_{k=0}^{\infty} \left(1 - \frac{9\sqrt{\lambda\eta}}{10}\right)^k A_0 \leq \left(\left(1 - \frac{9\sqrt{\lambda\eta}}{10}\right)^t + \frac{2}{9}\right) A_0, \end{aligned}$$

thus we know when $t \geq \frac{10}{9\sqrt{\lambda\eta}} \log(36)$, we have

$$A_t \leq \left(\frac{1}{36} + \frac{2}{9}\right) A_0 = \frac{1}{4} A_0,$$

also since

$$A_0 = V_0(w^*) - f(w^*) = f(\tilde{w}_{s-1}) - f(w^*) + \frac{\lambda}{2} \|\tilde{w}_{s-1} - w^*\|^2 \stackrel{\text{①}}{\leq} 2(f(\tilde{w}_{s-1}) - f(w^*)),$$

where at step ① we used the λ -strong convexity of $f(w)$. Combine above two inequality we get when $t \geq \frac{10}{9\sqrt{\lambda\eta}} \log(36)$, the expected objective suboptimality is halved:

$$\mathbb{E}f(w_t) - f(w^*) \leq A_t \leq \frac{1}{4} A_0 \leq \frac{1}{2} (f(\tilde{w}_{s-1}) - f(w^*)),$$

which concludes the proof. ■

Appendix C. Collections of Tools in the Analysis

Lemma 35 (*Matrix Bernstein, Theorem 1.4 of (Tropp, 2012) rephrased*) Let X_1, \dots, X_k be some independent, self-adjoint random matrices with dimension d , and assume each random matrix satisfies:

$$\mathbb{E}X_k = 0 \quad \text{and} \quad \|X_k\| \leq R \quad \text{almost surely.}$$

Then, for all $t \geq 0$,

$$P\left(\left\|\sum_k X_k\right\| \geq t\right) \leq d \exp\left(\frac{-t^2/2}{\sigma^2 + Rt/3}\right)$$

where

$$\sigma^2 := \left\|\sum_k \mathbb{E}(X_k^2)\right\|.$$

Lemma 36 (*Iteration complexity of SVRG, Corollary 1 of (Xiao and Zhang, 2014) rephrased*) If we apply SVRG to any finite-sum optimization objective $f(w)$ where each individual function is λ -strongly convex and L -smooth, then there are universal constant C

such that for any target accuracy ε , and success probability $1 - \delta$, SVRG is able to find a solution that satisfies ε objective suboptimality using

$$C \cdot \left(n + \frac{L}{\lambda} \right) \log \left(\frac{\varepsilon_0}{\delta \varepsilon} \right)$$

first order oracle calls of individual functions, where n is the total number of individual functions in $f(w)$, ε_0 is the initial objective suboptimality.

Lemma 37 (Iteration complexity of catalyst acceleration, Theorem 3.1 of (Lin et al., 2015a) rephrased) For any λ -strongly convex and L -smooth function $f(w)$, If the minimization step of (23) satisfies

$$f(w_r) - \frac{\gamma}{2} \|w_r - z_{r-1}\| - \min_w \left(f(w) + \frac{\gamma}{2} \|w - z_{r-1}\| \right) \leq \frac{2}{9} (f(w_0) - f(w^*)) \left(1 - \frac{9}{10} \sqrt{\frac{\lambda}{\lambda + \gamma}} \right)^r,$$

then if we initialize $\nu_0 = \sqrt{\frac{\lambda}{\lambda + \gamma}}$ and set ν_r such that $\nu_r^2 = (1 - \nu_r)\nu_{r-1}^2 + (\lambda\nu_r)/(\lambda + \gamma)$, then the sequences $\{w_r\}$ in Algorithm 2 satisfies

$$f(w_r) - f(w^*) \leq \frac{800(\lambda + \gamma)}{\lambda} \left(1 - \frac{9}{10} \sqrt{\frac{\lambda}{\lambda + \gamma}} \right)^{r+1} (f(w_0) - f(w^*)).$$

References

- Naman Agarwal, Brian Bullins, and Elad Hazan. Second order stochastic optimization in linear time. *arXiv preprint arXiv:1602.03943*, 2016.
- Zeyuan Allen-Zhu. Katyusha: The first direct acceleration of stochastic gradient methods. *arXiv preprint arXiv:1603.05953*, 2016.
- Zeyuan Allen-Zhu and Elad Hazan. Optimal black-box reductions between optimization objectives. In *Advances in Neural Information Processing Systems*, pages 1614–1622, 2016.
- Yossi Arjevani and Ohad Shamir. Oracle complexity of second-order methods for finite-sum problems. *arXiv preprint arXiv:1611.04982*, 2016.
- Raghu Bollapragada, Richard Byrd, and Jorge Nocedal. Exact and inexact subsampled newton methods for optimization. *arXiv preprint arXiv:1609.08502*, 2016.
- Léon Bottou, Frank E Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *arXiv preprint arXiv:1606.04838*, 2016.
- Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

- Richard H Byrd, Gillian M Chin, Will Neveitt, and Jorge Nocedal. On the use of stochastic hessian information in optimization methods for machine learning. *SIAM Journal on Optimization*, 21(3):977–995, 2011.
- Richard H Byrd, Samantha L Hansen, Jorge Nocedal, and Yoram Singer. A stochastic quasi-newton method for large-scale optimization. *SIAM Journal on Optimization*, 26(2):1008–1031, 2016.
- Samprit Chatterjee and Ali S Hadi. *Sensitivity analysis in linear regression*, volume 327. John Wiley & Sons, 2009.
- Andrew Cotter, Ohad Shamir, Nati Srebro, and Karthik Sridharan. Better mini-batch algorithms via accelerated gradient methods. In *Advances in neural information processing systems*, pages 1647–1655, 2011.
- Aaron Defazio. A simple practical accelerated method for finite sums. In *Advances In Neural Information Processing Systems*, pages 676–684, 2016.
- Aaron Defazio, Francis Bach, and Simon Lacoste-Julien. Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. In *Advances in Neural Information Processing Systems*, pages 1646–1654, 2014.
- Ofer Dekel, Ran Gilad-Bachrach, Ohad Shamir, and Lin Xiao. Optimal distributed online prediction using mini-batches. *Journal of Machine Learning Research*, 13(Jan):165–202, 2012.
- Olivier Devolder, François Glineur, and Yurii Nesterov. First-order methods of smooth convex optimization with inexact oracle. *Mathematical Programming*, 146(1-2):37–75, 2014.
- Petros Drineas and Michael W Mahoney. Lectures on randomized numerical linear algebra. *arXiv preprint arXiv:1712.08880*, 2017.
- Murat A. Erdogdu and Andrea Montanari. Convergence rates of sub-sampled newton methods. In *Advances in Neural Information Processing Systems*, pages 3052–3060, 2015.
- Roy Frostig, Rong Ge, Sham Kakade, and Aaron Sidford. Un-regularizing: approximate proximal point and faster stochastic algorithms for empirical risk minimization. In *Proceedings of The 32nd International Conference on Machine Learning*, pages 2540–2548, 2015.
- Alon Gonen, Francesco Orabona, and Shai Shalev-Shwartz. Solving ridge regression using sketched preconditioned svrg. In *Proceedings of The 33rd International Conference on Machine Learning*, pages 1397–1405, 2016.
- Robert Gower, Donald Goldfarb, and Peter Richtárik. Stochastic block bfgs: squeezing more curvature out of data. In *International Conference on Machine Learning*, pages 1869–1878, 2016.

- Daniel Hsu, Sham M Kakade, and Tong Zhang. Random design analysis of ridge regression. *Foundations of Computational Mathematics*, 3(14):569–600, 2014.
- Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in Neural Information Processing Systems*, pages 315–323, 2013.
- Jakub Konečný, Jie Liu, Peter Richtárik, and Martin Takáč. Mini-batch semi-stochastic gradient descent in the proximal setting. *IEEE Journal of Selected Topics in Signal Processing*, 10(2):242–255, 2016.
- Guanghui Lan. An optimal method for stochastic composite optimization. *Mathematical Programming*, 133(1):365–397, 2012.
- Mu Li, Tong Zhang, Yuqiang Chen, and Alexander J. Smola. Efficient mini-batch training for stochastic optimization. In *Proc. of the 20th ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining (SIGKDD 2014)*, pages 661–670, 2014.
- Hongzhou Lin, Julien Mairal, and Zaid Harchaoui. A universal catalyst for first-order optimization. In *Advances in Neural Information Processing Systems*, pages 3384–3392, 2015a.
- Qihang Lin, Xi Chen, and Javier Peña. A sparsity preserving stochastic gradient methods for sparse regression. *Computational Optimization and Applications*, 58(2):455–482, 2014.
- Qihang Lin, Zhaosong Lu, and Lin Xiao. An accelerated randomized proximal coordinate gradient method and its application to regularized empirical risk minimization. *SIAM Journal on Optimization*, 25(4):2244–2273, 2015b.
- Dong C Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical programming*, 45(1):503–528, 1989.
- Sharon Lohr. *Sampling: design and analysis*. Nelson Education, 2009.
- Michael W Mahoney et al. Randomized algorithms for matrices and data. *Foundations and Trends® in Machine Learning*, 3(2):123–224, 2011.
- Philipp Moritz, Robert Nishihara, and Michael Jordan. A linearly-convergent stochastic l-bfgs algorithm. In *Artificial Intelligence and Statistics*, pages 249–258, 2016.
- Yurii Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*, volume 87. Springer Science & Business Media, 2004.
- Atsushi Nitanda. Stochastic proximal gradient descent with acceleration techniques. In *Advances in Neural Information Processing Systems*, pages 1574–1582, 2014.
- Mert Pilanci and Martin J Wainwright. Newton sketch: A near linear-time optimization algorithm with linear-quadratic convergence. *SIAM Journal on Optimization*, 27(1):205–245, 2017.

- Zheng Qu, Peter Richtárik, Martin Takáč, and Olivier Fercoq. Sdna: Stochastic dual newton ascent for empirical risk minimization. *arXiv preprint arXiv:1502.02268*, 2015.
- Farbod Roosta-Khorasani and Michael W Mahoney. Sub-sampled newton methods i: globally convergent algorithms. *arXiv preprint arXiv:1601.04737*, 2016a.
- Farbod Roosta-Khorasani and Michael W Mahoney. Sub-sampled newton methods ii: Local convergence rates. *arXiv preprint arXiv:1601.04738*, 2016b.
- Nicolas L Roux, Mark Schmidt, and Francis R Bach. A stochastic gradient method with an exponential convergence rate for finite training sets. In *Advances in Neural Information Processing Systems*, pages 2663–2671, 2012.
- Mark Schmidt, Nicolas L Roux, and Francis R Bach. Convergence rates of inexact proximal-gradient methods for convex optimization. In *Advances in neural information processing systems*, pages 1458–1466, 2011.
- Shai Shalev-Shwartz and Tong Zhang. Stochastic dual coordinate ascent methods for regularized loss minimization. *Journal of Machine Learning Research*, 14(Feb):567–599, 2013.
- Shai Shalev-Shwartz and Tong Zhang. Accelerated proximal stochastic dual coordinate ascent for regularized loss minimization. *Mathematical Programming*, 155(1-2):105–145, 2016.
- Ohad Shamir, Nathan Srebro, and Tong Zhang. Communication-efficient distributed optimization using an approximate Newton-type method. In *Proc. of the 31st Int. Conf. Machine Learning (ICML 2014)*, pages 1000–1008, 2014.
- Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.
- Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12(4):389–434, 2012.
- Silvia Villa, Saverio Salzo, Luca Baldassarre, and Alessandro Verri. Accelerated and inexact forward-backward algorithms. *SIAM Journal on Optimization*, 23(3):1607–1633, 2013.
- Jialei Wang, Weiran Wang, and Nathan Srebro. Memory and communication efficient distributed stochastic optimization with minibatch prox. In *Proceedings of The 30th Conference on Learning Theory*, 2017a.
- Xiao Wang, Shiqian Ma, Donald Goldfarb, and Wei Liu. Stochastic quasi-newton methods for nonconvex stochastic optimization. *SIAM Journal on Optimization*, 27(2):927–956, 2017b.
- Blake E Woodworth and Nati Srebro. Tight complexity bounds for optimizing composite objectives. In *Advances in Neural Information Processing Systems*, pages 3639–3647, 2016.
- Stephen Wright and Jorge Nocedal. Numerical optimization. *Springer Science*, 35:67–68, 1999.

- Lin Xiao and Tong Zhang. A proximal stochastic gradient method with progressive variance reduction. *SIAM Journal on Optimization*, 24(4):2057–2075, 2014.
- Peng Xu, Jiyan Yang, Farbod Roosta-Khorasani, Christopher Ré, and Michael W Mahoney. Sub-sampled newton methods with non-uniform sampling. In *Advances in Neural Information Processing Systems*, pages 3000–3008, 2016.
- Tianbao Yang, Rong Jin, Shenghuo Zhu, and Qihang Lin. On data preconditioning for regularized loss minimization. *Machine Learning*, 103(1):57–79, 2016.
- Yuchen Zhang and Lin Xiao. Stochastic primal-dual coordinate method for regularized empirical risk minimization. In *Proceedings of The 32nd International Conference on Machine Learning*, pages 353–361, 2015.
- Yuchen Zhang, John Duchi, and Martin Wainwright. Divide and conquer kernel ridge regression: A distributed algorithm with minimax optimal rates. *The Journal of Machine Learning Research*, 16(1):3299–3340, 2015.
- Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.