

Second-Order Stochastic Optimization for Machine Learning in Linear Time

Naman Agarwal

*Computer Science Department
Princeton University
Princeton, NJ 08540, USA*

NAMANA@CS.PRINCETON.EDU

Brian Bullins

*Computer Science Department
Princeton University
Princeton, NJ 08540, USA*

BBULLINS@CS.PRINCETON.EDU

Elad Hazan

*Computer Science Department
Princeton University
Princeton, NJ 08540, USA*

EHAZAN@CS.PRINCETON.EDU

Editor: Tong Zhang

Abstract

First-order stochastic methods are the state-of-the-art in large-scale machine learning optimization owing to efficient per-iteration complexity. Second-order methods, while able to provide faster convergence, have been much less explored due to the high cost of computing the second-order information. In this paper we develop second-order stochastic methods for optimization problems in machine learning that match the per-iteration cost of gradient based methods, and in certain settings improve upon the overall running time over popular first-order methods. Furthermore, our algorithm has the desirable property of being implementable in time linear in the sparsity of the input data.

Keywords: second-order optimization, convex optimization, regression

1. Introduction

In recent literature stochastic first-order optimization has taken the stage as the primary workhorse for training learning models, due in large part to its affordable computational costs which are linear (in the data representation) per iteration. The main research effort devoted to improving the convergence rates of first-order methods have introduced elegant ideas and algorithms in recent years, including adaptive regularization (Duchi et al., 2011), variance reduction (Johnson and Zhang, 2013; Defazio et al., 2014), dual coordinate ascent (Shalev-Shwartz and Zhang, 2013), and many more.

In contrast, second-order methods have typically been much less explored in large scale machine learning (ML) applications due to their prohibitive computational cost per iteration which requires computation of the Hessian in addition to a matrix inversion. These operations are infeasible for large scale problems in high dimensions.

In this paper we propose a family of novel second-order algorithms, LiSSA (Linear time Stochastic Second-Order Algorithm) for convex optimization that attain fast convergence rates while also allowing for an implementation with linear time per-iteration cost, matching the running time of the best known gradient-based methods. Moreover, in the setting where the number of training examples m is much larger than the underlying dimension d , we show that our algorithm has provably faster running time than the best known gradient-based methods.

Formally, the main optimization problem we are concerned with is the empirical risk minimization (ERM) problem:

$$\min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x}) = \min_{\mathbf{x} \in \mathbb{R}^d} \left\{ \frac{1}{m} \sum_{k=1}^m f_k(\mathbf{x}) + R(\mathbf{x}) \right\}$$

where each $f_k(\mathbf{x})$ is a convex function and $R(\mathbf{x})$ is a convex regularizer. The above optimization problem is the standard objective minimized in most supervised learning settings. Examples include logistic regression, SVMs, etc. A common aspect of many applications of ERM in machine learning is that the loss function $f_i(\mathbf{x})$ is of the form $l(\mathbf{x}^\top \mathbf{v}_i, y_i)$ where (\mathbf{v}_i, y_i) is the i^{th} training example-label pair. We call such functions generalized linear models (GLM) and will restrict our attention to this case. We will assume that the regularizer is an ℓ_2 regularizer,¹ typically $\|\mathbf{x}\|^2$.

Our focus is second-order optimization methods (Newton’s method), where in each iteration, the underlying principle is to move to the minimizer of the second-order Taylor approximation at any point. Throughout the paper, we will let $\nabla^{-2}f(\mathbf{x}) \triangleq [\nabla^2 f(\mathbf{x})]^{-1}$. The update of Newton’s method at a point \mathbf{x}_t is then given by

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \nabla^{-2}f(\mathbf{x}_t)\nabla f(\mathbf{x}_t). \quad (1)$$

Certain desirable properties of Newton’s method include the fact that its updates are independent of the choice of coordinate system and that the Hessian provides the necessary regularization based on the curvature at the present point. Indeed, Newton’s method can be shown to eventually achieve quadratic convergence (Nesterov, 2013). Although Newton’s method comes with good theoretical guarantees, the complexity per step grows roughly as $\Omega(md^2 + d^\omega)$ (the former term for computing the Hessian and the latter for inversion, where $\omega \approx 2.37$ is the matrix multiplication constant), making it prohibitive in practice. Our main contribution is a suite of algorithms, each of which performs an approximate Newton update based on stochastic Hessian information and is implementable in linear $O(d)$ time. These algorithms match and improve over the performance of first-order methods in theory and give promising results as an optimization method on real world data sets. In the following we give a summary of our results. We propose two algorithms, LiSSA and LiSSA-Sample.

LiSSA: Algorithm 1 is a practical stochastic second-order algorithm based on a novel estimator of the Hessian inverse, leading to an efficient approximate Newton step (Equation 1). The estimator is based on the well known Taylor approximation of the inverse (Fact 2) and is described formally in Section 3.1. We prove the following informal theorem about LiSSA.

1. It can be seen that some of these assumptions can be relaxed in our analyses, but since these choices are standard in ML, we make these assumptions to simplify the discourse.

Theorem 1 (Informal) *LiSSA returns a point \mathbf{x}_t such that $f(\mathbf{x}_t) \leq \min_{\mathbf{x}^*} f(\mathbf{x}^*) + \varepsilon$ in total time*

$$\tilde{O} \left((m + S_1 \kappa) d \log \left(\frac{1}{\varepsilon} \right) \right)$$

where κ is the underlying condition number of the problem and S_1 is a bound on the variance of the estimator.

The precise version of the above theorem appears as Theorem 7. In theory, the best bound we can show for S_1 is $O(\kappa^2)$; however, in our experiments we observe that setting S_1 to be a small constant (often 1) is sufficient. We conjecture that S_1 can be improved to $O(1)$ and leave this for future work. If indeed S_1 can be improved to $O(1)$ (as is indicated by our experiments), LiSSA enjoys a convergence rate comparable to first-order methods. We provide a detailed comparison of our results with existing first-order and second-order methods in Section 1.2. Moreover, in Section 7 we present experiments on real world data sets that demonstrate that LiSSA as an optimization method performs well as compared to popular first-order methods. We also show that LiSSA runs in time proportional to input sparsity, making it an attractive method for high-dimensional sparse data.

LiSSA-Sample: This variant brings together efficient first-order algorithms with matrix sampling techniques (Li et al., 2013; Cohen et al., 2015) to achieve better runtime guarantees than the state-of-the-art in convex optimization for machine learning in the regime when $m > d$. Specifically, we prove the following theorem:

Theorem 2 (Informal) *LiSSA-Sample returns a point \mathbf{x}_t such that $f(\mathbf{x}_t) \leq \min_{\mathbf{x}^*} f(\mathbf{x}^*) + \varepsilon$ in total time*

$$\tilde{O} \left(m + \sqrt{\kappa d} \right) d \log^2 \left(\frac{1}{\varepsilon} \right) \log \log \left(\frac{1}{\varepsilon} \right).$$

The above result improves upon the best known running time for first-order methods achieved by acceleration when we are in the setting where $\kappa > m \gg d$. We discuss the implication of our bounds and further work in Section 1.3.

In all of our results stated above κ corresponds to the condition number of the underlying problem. In particular we assume some strong convexity for the underlying problem. This is a standard assumption which is usually enforced by the addition of the ℓ_2 regularizer. In stating our results formally we stress on the nuances between different notions of the condition number (ref. Section 2.1), and we state our results precisely with respect to these notions. In general, all of our generalization/relaxations of the condition number are smaller than $\frac{1}{\lambda}$ where λ is the regularization parameter, and this is usually taken to be the condition number of the problem. The condition of strong convexity has been relaxed in literature by introducing proximal methods. It is an interesting direction to adapt our results in those setting which we leave for future work.

We also remark that all of our results focus on the very high accuracy regime. In general the benefits of linear convergence and second-order methods can be seen to be effective only when considerably small error is required. This is also the case for recent advances in fast first-order methods where their improvement over stochastic gradient descent becomes apparent only in the high accuracy regime. Our experiments also demonstrate that second-order methods can improve upon fast first-order methods in the regime of very high accuracy.

While it is possible that this regime is less interesting for generalization, in this paper we focus on the optimization problem itself.

We further consider the special case when the function f is self-concordant. Self-concordant functions are a sub-class of convex functions which have been extensively studied in convex optimization literature in the context of interior point methods (Nemirovski, 2004). For self-concordant functions we propose an algorithm (Algorithm 5) which achieves linear convergence with running time guarantees independent of the condition number. We prove the formal running time guarantee as Theorem 25.

We believe our main contribution to be a demonstration of the fact that second-order methods are comparable to, or even better than, first-order methods in the large data regime, in both theory and practice.

1.1 Overview of Techniques

LiSSA: The key idea underlying LiSSA is the use of the Taylor expansion to construct a natural estimator of the Hessian inverse. Indeed, as can be seen from the description of the estimator in Section 3.1, the estimator we construct becomes unbiased in the limit as we include additional terms in the series. We note that this is not the case with estimators that were considered in previous works such as that of Erdogdu and Montanari (2015), and so we therefore consider our estimator to be more natural. In the implementation of the algorithm we achieve the optimal bias/variance trade-off by truncating the series appropriately.

An important observation underlying our linear time $O(d)$ step is that for GLM functions, $\nabla^2 f_i(\mathbf{x})$ has the form $\alpha \mathbf{v}_i \mathbf{v}_i^\top$ where α is a scalar dependent on $\mathbf{v}_i^\top \mathbf{x}$. A single step of LiSSA requires us to efficiently compute $\nabla^2 f_i(\mathbf{x}) \mathbf{b}$ for a given vector \mathbf{b} . In this case it can be seen that the matrix-vector product reduces to a vector-vector product, giving us an $O(d)$ time update.

LiSSA-Sample: LiSSA-Sample is based on Algorithm 2, which represents a general family of algorithms that couples the quadratic minimization view of Newton’s method with any efficient first-order method. In essence, Newton’s method allows us to reduce (up to log log factors) the optimization of a general convex function to solving intermediate quadratic or ridge regression problems. Such a reduction is useful in two ways.

First, as we demonstrate through our algorithm LiSSA-Sample, the quadratic nature of ridge regression problems allows us to leverage powerful sampling techniques, leading to an improvement over the running time of the best known accelerated first-order method. On a high level this improvement comes from the fact that when solving a system of m linear equations in d dimensions, a constant number of passes through the data is enough to reduce the system to $O(d \log(d))$ equations. We carefully couple this principle and the computation required with accelerated first-order methods to achieve the running times for LiSSA-Sample. The result for the quadratic sub-problem (ridge regression) is stated in Theorem 15, and the result for convex optimization is stated in Theorem 16.

The second advantage of the reduction to quadratic sub-problems comes from the observation that the intermediate quadratic sub-problems can potentially be better conditioned than the function itself, allowing us a better choice of the step size in practice. We define these local notions of condition number formally in Section 2.1 and summarize the typical benefits for such algorithms in Theorem 13. In theory this is not a significant improve-

ment; however, in practice we believe that this could be significant and lead to runtime improvements.²

To achieve the bound for LiSSA-Sample we extend the definition and procedure for sampling via leverage scores described by Cohen et al. (2015) to the case when the matrix is given as a sum of PSD matrices and not just rank one matrices. We reformulate and reprove the theorems proved by Cohen et al. (2015) in this context, which may be of independent interest.

1.2 Comparison with Related Work

In this section we aim to provide a short summary of the key ideas and results underlying optimization methods for large scale machine learning. We divide the summary into three high level principles: first-order gradient-based methods, second-order Hessian-based methods, and quasi-Newton methods. For the sake of brevity we will restrict our summary to results in the case when the objective is strongly convex, which as justified above is usually ensured by the addition of an appropriate regularizer. In such settings the main focus is often to obtain algorithms which have provably linear convergence and fast implementations.

First-Order Methods: First-order methods have dominated the space of optimization algorithms for machine learning owing largely to the fact that they can be implemented in time proportional to the underlying dimension (or sparsity). Gradient descent is known to converge linearly to the optimum with a rate of convergence that is dependent upon the condition number of the objective. In the large data regime, stochastic first-order methods, introduced and analyzed first by Robbins and Monro (1951), have proven especially successful. Stochastic gradient descent (SGD), however, converges sub-linearly even in the strongly convex setting. A significant advancement in terms of the running time of first-order methods was achieved recently by a clever merging of stochastic gradient descent with its full version to provide *variance reduction*. The representative algorithms in this space are SAGA (Roux et al., 2012; Defazio et al., 2014) and SVRG (Johnson and Zhang, 2013; Zhang et al., 2013). The key technical achievement of the above algorithms is to relax the running time dependence on m (the number of training examples) and κ (the condition number) from a product to a sum. Another algorithm which achieves similar running time guarantees is based on dual coordinate ascent, known as SDCA (Shalev-Shwartz and Zhang, 2013).

Further improvements over SAGA, SVRG and SDCA have been obtained by applying the classical idea of *acceleration* emerging from the seminal work of Nesterov (1983). The progression of work here includes an accelerated version of SDCA (Shalev-Shwartz and Zhang, 2016); APCG (Lin et al., 2014); Catalyst (Lin et al., 2015), which provides a generic framework to accelerate first order algorithms; and Katyusha (Allen-Zhu, 2016), which introduces the concept of negative momentum to extend acceleration for variance reduced algorithms beyond the strongly convex setting. The key technical achievement of accelerated methods in general is to reduce the dependence on condition number from linear to a square root. We summarize these results in Table 1.

2. While this is a difficult property to verify experimentally, we conjecture that this is a possible explanation for why LiSSA performs better than first-order methods on certain data sets and ranges of parameters.

LiSSA places itself naturally into the space of fast first-order methods by having a running time dependence that is comparable to SAGA/SVRG (ref. Table 1). In LiSSA-Sample we leverage the quadratic structure of the sub-problem for which efficient sampling techniques have been developed in the literature and use accelerated first-order methods to improve the running time in the case when the underlying dimension is much smaller than the number of training examples. Indeed, to the best of our knowledge LiSSA-Sample is the theoretically fastest known algorithm under the condition $m \gg d$. Such an improvement seems out of hand for the present first-order methods as it seems to strongly leverage the quadratic nature of the sub-problem to reduce its size. We summarize these results in Table 1.

Algorithm	Runtime	Reference
SVRG, SAGA, SDCA	$(md + O(\hat{\kappa}d)) \log\left(\frac{1}{\varepsilon}\right)$	Johnson and Zhang (2013) Zhang et al. (2013) Roux et al. (2012) Shalev-Shwartz and Zhang (2013)
LiSSA	$(md + O(\hat{\kappa}_l)S_1d) \log\left(\frac{1}{\varepsilon}\right)$	Corollary 8
AccSDCA, Catalyst Katyusha	$\tilde{O}\left(md + d\sqrt{\hat{\kappa}m}\right) \log\left(\frac{1}{\varepsilon}\right)$	Shalev-Shwartz and Zhang (2016) Lin et al. (2015) Allen-Zhu (2016)
LiSSA-Sample	$\tilde{O}\left(md + d\sqrt{\kappa_{sample}d}\right) \log^2\left(\frac{1}{\varepsilon}\right)$	Theorem 16

Table 1: Run time comparisons. Refer to Section 2.1 for definitions of the various notions of condition number.

Second-Order Methods: Second-order methods such as Newton’s method have classically been used in optimization in many different settings including development of interior point methods (Nemirovski, 2004) for general convex programming. The key advantage of Newton’s method is that it achieves a linear-quadratic convergence rate. However, naive implementations of Newton’s method have two significant issues, namely that the standard analysis requires the full Hessian calculation which costs $O(md^2)$, an expense not suitable for machine learning applications, and the matrix inversion typically requires $O(d^3)$ time. These issues were addressed recently by the algorithm NewSamp (Erdogdu and Montanari, 2015) which tackles the first issue by subsampling and the second issue by low-rank projections. We improve upon the work of Erdogdu and Montanari (2015) by defining a more *natural* estimator for the Hessian inverse and by demonstrating that the estimator can be computed in time proportional to $O(d)$. We also point the reader to the works of Martens (2010); Byrd et al. (2011) which incorporate the idea of taking samples of the Hessian; however, these works do not provide precise running time guarantees on their proposed algorithm based on problem specific parameters. Second-order methods have also enjoyed success in the distributed setting (Shamir et al., 2014).

Quasi-Newton Methods: The expensive computation of the Newton step has also been tackled via estimation of the curvature from the change in gradients. These algorithms are generally known as quasi-Newton methods stemming from the seminal BFGS algo-

rithm (Broyden, 1970; Fletcher, 1970; Goldfarb, 1970; Shanno, 1970). The book of Nocedal and Wright (2006) is an excellent reference for the algorithm and its limited memory variant (L-BFGS). The more recent work in this area has focused on stochastic quasi-Newton methods which were proposed and analyzed in various settings by Schraudolph et al. (2007); Mokhtari and Ribeiro (2014); Byrd et al. (2016). These works typically achieve sub-linear convergence to the optimum. A significant advancement in this line of work was provided by Moritz et al. (2016) who propose an algorithm based on L-BFGS by incorporating ideas from variance reduction to achieve linear convergence to the optimum in the strongly convex setting. Although the algorithm achieves linear convergence, the running time of the algorithm depends poorly on the condition number (as acknowledged by the authors). Indeed, in applications that interest us, the condition number is not necessarily a constant as is typically assumed to be the case for the theoretical results in Moritz et al. (2016).

Our key observation of linear time Hessian-vector product computations for machine learning applications provides evidence that in such instances, obtaining true Hessian information is efficient enough to alleviate the need for quasi-Newton information via gradients.

1.3 Discussion and Subsequent Work

In this section we provide a brief survey of certain technical aspects of our bounds which have since been improved by subsequent work.

An immediate improvement in terms of $S_2 \sim \kappa$ (in fact suggested in the original manuscript) was achieved by Bollapragada et al. (2016) via conjugate gradient on a sub-sampled Hessian which reduces this to $\sqrt{\kappa}$. A similar improvement can also be achieved in theory through the extensions of LiSSA proposed in the paper. As we show in Section 7, the worse dependence on condition number has an effect on the running time when κ is quite large.³ Accelerated first-order methods, such as APCG (Lin et al., 2014), outperform LiSSA in this regime. To the best of our knowledge second-order stochastic methods have so far not exhibited an improvement in that regime experimentally. We believe a more practical version of LiSSA-Sample could lead to improvements in this regime, leaving this as future work.

To the best of our knowledge the factor of $S_1 = \kappa^2$ that appears to reduce the variance of our estimator has yet not been improved despite it being $O(1)$ in our experiments. This is an interesting question to which partial answers have been provided in the analysis of Ye et al. (2017).

Significant progress has been made in the space of inexact Newton methods based on matrix sketching techniques. We refer the reader to the works of Pilanci and Wainwright (2015); Xu et al. (2016); Cohen (2016); Luo et al. (2016); Ye et al. (2017) and the references therein.

We would also like to comment on the presence of a warm start parameter $\frac{1}{\kappa M}$ in our proofs of Theorems 7 and 15. In our experiments the warm start we required would be quite small (often a few steps of gradient descent would be sufficient) to make LiSSA converge. The warm start does not affect the asymptotic results proven in Theorems 7 and 15 because getting to such a warm start is independent of ε . However, improving this warm start, especially in the context of Theorem 15, is left as interesting future work.

3. Equivalently, λ is small.

On the complexity side, Arjevani and Shamir (2016) proved lower bounds on the best running times achievable by second-order methods. In particular, they show that to get the faster rates achieved by LiSSA-Sample, it is necessary to use a non-uniform sampling based method as employed by LiSSA-Sample. We would like to remark that in theory, up to logarithmic factors, the running time of LiSSA-Sample is still the best achieved so far in the setting $m \gg d$. Some of the techniques and motivations from this work were also generalized by the authors to provide faster rates for a large family of non-convex optimization problems (Agarwal et al., 2017).

1.4 Organization of the Paper

The paper is organized as follows: we first present the necessary definitions, notations and conventions adopted throughout the paper in Section 2. We then describe our estimator for LiSSA, as well as state and prove the convergence guarantee for LiSSA in Section 3. After presenting a generic procedure to couple first-order methods with Newton’s method in 4, we present LiSSA-Sample and the associated fast quadratic solver in Section 5. We then present our results regarding self-concordant functions in Section 6. Finally, we present an experimental evaluation of LiSSA in Section 7.

2. Preliminaries

We adopt the convention of denoting vectors and scalars in lowercase, matrices in uppercase, and vectors in boldface. We will use $\|\cdot\|$ without a subscript to denote the ℓ_2 norm for vectors and the spectral norm for matrices. Throughout the paper we denote $\mathbf{x}^* \triangleq \operatorname{argmin}_{\mathbf{x} \in \mathcal{K}} f(\mathbf{x})$. A convex function f is defined to be α -strongly convex and β -smooth if, for all \mathbf{x}, \mathbf{y} ,

$$\nabla f(\mathbf{x})^\top (\mathbf{y} - \mathbf{x}) + \frac{\beta}{2} \|\mathbf{y} - \mathbf{x}\|^2 \geq f(\mathbf{y}) - f(\mathbf{x}) \geq \nabla f(\mathbf{x})^\top (\mathbf{y} - \mathbf{x}) + \frac{\alpha}{2} \|\mathbf{y} - \mathbf{x}\|^2.$$

The following is a well known fact about the inverse of a matrix A s.t. $\|A\| \leq 1$ and $A \succ 0$:

$$A^{-1} = \sum_{i=0}^{\infty} (I - A)^i. \tag{2}$$

2.1 Definition of Condition Numbers

We now define several measures for the condition number of a function f . The differences between these notions are subtle and we use them to precisely characterize the running time for our algorithms.⁴

For an α -strongly convex and β -smooth function f , the condition number of the function is defined as $\kappa(f) \triangleq \frac{\beta}{\alpha}$, or κ when the function is clear from the context. Note that by definition this corresponds to the following notion:

$$\kappa \triangleq \frac{\max_{\mathbf{x}} \lambda_{\max}(\nabla^2 f(\mathbf{x}))}{\min_{\mathbf{x}} \lambda_{\min}(\nabla^2 f(\mathbf{x}))}.$$

4. During initial reading we suggest the reader to skip the subtlety with these notions with the knowledge that they are all smaller than the pessimistic bound one can achieve by considering a value proportional to $O(\lambda^{-1})$, where λ is the coefficient of the ℓ_2 regularizer.

We define a slightly relaxed notion of condition number where the *max* moves out of the fraction above. We refer to this notion as a *local* condition number κ_l as compared to the *global* condition number κ defined above:

$$\kappa_l \triangleq \max_{\mathbf{x}} \frac{\lambda_{\max}(\nabla^2 f(\mathbf{x}))}{\lambda_{\min}(\nabla^2 f(\mathbf{x}))} .$$

It follows that $\kappa_l \leq \kappa$. The above notions are defined for any general function f , but in the case of functions of the form $f(\mathbf{x}) = \frac{1}{m} \sum_{k=1}^m f_k(\mathbf{x})$, a further distinction is made with respect to the component functions. We refer to such definitions of the condition number by $\hat{\kappa}$. In such cases one typically assumes the each component is bounded by $\beta_{\max}(\mathbf{x}) \triangleq \max_k \lambda_{\max}(\nabla^2 f_k(\mathbf{x}))$. The running times of algorithms like SVRG depend on the following notion of condition number:

$$\hat{\kappa} = \frac{\max_{\mathbf{x}} \beta_{\max}(\mathbf{x})}{\min_{\mathbf{x}} \lambda_{\min}(\nabla^2 f(\mathbf{x}))} .$$

Similarly, we define a notion of local condition number for $\hat{\kappa}$, namely

$$\hat{\kappa}_l \triangleq \max_{\mathbf{x}} \frac{\beta_{\max}(\mathbf{x})}{\lambda_{\min}(\nabla^2 f(\mathbf{x}))} ,$$

and it again follows that $\hat{\kappa}_l \leq \hat{\kappa}$.

For our (admittedly pessimistic) bounds on the variance we also need a per-component strong convexity bound $\alpha_{\min}(\mathbf{x}) \triangleq \min_k \lambda_{\min}(\nabla^2 f_k(\mathbf{x}))$. We can now define

$$\hat{\kappa}_l^{\max} \triangleq \max_{\mathbf{x}} \frac{\beta_{\max}(\mathbf{x})}{\alpha_{\min}(\mathbf{x})} .$$

Assumptions: In light of the previous definitions, we make the following assumptions about the given function $f(\mathbf{x}) = \frac{1}{m} \sum_{k=1}^m f_k(\mathbf{x})$ to make the analysis easier. We first assume that the regularization term has been divided equally and included in $f_k(\mathbf{x})$. We further assume that each $\nabla^2 f_k(\mathbf{x}) \preceq I$.⁵ We also assume that f is α -strongly convex and β -smooth, $\hat{\kappa}_l$ is the associated local condition number and $\nabla^2 f$ has a Lipschitz constant bounded by M .

We now collect key concepts and pre-existing results that we use for our analysis in the rest of the paper.

Matrix Concentration: The following lemma is a standard concentration of measure result for sums of independent matrices.⁶ An excellent reference for this material is by Tropp (2012).

5. The scaling is without loss of generality, even when looking at additive errors, as this gets picked up in the log-term due to the linear convergence.

6. The theorem in the reference states the inequality for the more nuanced bounded variance case. We only state the simpler bounded spectral norm case which suffices for our purposes.

Theorem 3 (Matrix Bernstein, Tropp, 2012) Consider a finite sequence $\{X_k\}$ of independent, random, Hermitian matrices with dimension d . Assume that

$$\mathbf{E}[X_k] = 0 \text{ and } \|X_k\| \leq R.$$

Define $Y = \sum_k X_k$. Then we have for all $t \geq 0$,

$$\Pr(\|Y\| \geq t) \leq d \exp\left(\frac{-t^2}{4R^2}\right).$$

Accelerated SVRG: The following theorem was proved by Lin et al. (2015).

Theorem 4 Given a function $f(\mathbf{x}) = \frac{1}{m} \sum_{k=1}^m f_k(\mathbf{x})$ with condition number κ , the accelerated version of SVRG via Catalyst (Lin et al., 2015) finds an ε -approximate minimum with probability $1 - \delta$ in time

$$\tilde{O}(md + \min(\sqrt{\kappa m}, \kappa)d) \log\left(\frac{1}{\varepsilon}\right).$$

Sherman-Morrison Formula: The following is a well-known expression for writing the inverse of rank one perturbations of matrices:

$$(A + \mathbf{v}\mathbf{v}^T)^{-1} = A^{-1} - \frac{A^{-1}\mathbf{v}\mathbf{v}^T A^{-1}}{1 + \mathbf{v}^T A^{-1}\mathbf{v}}.$$

3. LiSSA: Linear (time) Stochastic Second-Order Algorithm

In this section, we provide an overview of LiSSA (Algorithm 1) along with its main convergence results.

3.1 Estimators for the Hessian Inverse

Based on a recursive reformulation of the Taylor expansion (Equation 2), we may describe an unbiased estimator of the Hessian. For a matrix A , define A_j^{-1} as the first j terms in the Taylor expansion, i.e.,

$$A_j^{-1} \triangleq \sum_{i=0}^j (I - A)^i, \text{ or equivalently } A_j^{-1} \triangleq I + (I - A)A_{j-1}^{-1}.$$

Note that $\lim_{j \rightarrow \infty} A_j^{-1} \rightarrow A^{-1}$. Using the above recursive formulation, we now describe an unbiased estimator of $\nabla^{-2}f$ by first deriving an unbiased estimator $\tilde{\nabla}^{-2}f_j$ for $\nabla^{-2}f_j$.

Definition 5 (Estimator) Given j independent and unbiased samples $\{X_1 \dots X_j\}$ of the Hessian $\nabla^2 f$, define $\{\tilde{\nabla}^{-2}f_0, \dots, \tilde{\nabla}^{-2}f_j\}$ recursively as follows:

$$\tilde{\nabla}^{-2}f_0 = I \text{ and } \tilde{\nabla}^{-2}f_t = I + (I - X_t)\tilde{\nabla}^{-2}f_{t-1} \text{ for } t = 1, \dots, j.$$

It can be readily seen that $\mathbf{E}[\tilde{\nabla}^{-2}f_j] = \nabla^{-2}f_j$, and so $\mathbf{E}[\tilde{\nabla}^{-2}f_j] \rightarrow \nabla^{-2}f$ as $j \rightarrow \infty$, providing us with an unbiased estimator in the limit.

Algorithm 1 LiSSA: Linear (time) Stochastic Second-Order Algorithm

Input: $T, f(\mathbf{x}) = \frac{1}{m} \sum_{k=1}^m f_k(\mathbf{x}), S_1, S_2, T_1$
 $\mathbf{x}_1 = FO(f(\mathbf{x}), T_1)$
for $t = 1$ to T **do**
 for $i = 1$ to S_1 **do**
 $X_{[i,0]} = \nabla f(\mathbf{x}_t)$
 for $j = 1$ to S_2 **do**
 Sample $\tilde{\nabla}^2 f_{[i,j]}(\mathbf{x}_t)$ uniformly from $\{\nabla^2 f_k(\mathbf{x}_t) \mid k \in [m]\}$
 $X_{[i,j]} = \nabla f(\mathbf{x}_t) + (I - \tilde{\nabla}^2 f_{[i,j]}(\mathbf{x}_t))X_{[i,j-1]}$
 end for
 end for
 $X_t = 1/S_1 \left(\sum_{i=1}^{S_1} X_{[i,S_2]} \right)$
 $\mathbf{x}_{t+1} = \mathbf{x}_t - X_t$
end for
return \mathbf{x}_{T+1}

Remark 6 *One can also define and analyze a simpler (non-recursive) estimator based on directly sampling terms from the series in Equation (2). Theoretically, one can get similar guarantees for the estimator; however, empirically our proposed estimator exhibited better performance.*

3.2 Algorithm

Our algorithm runs in two phases: in the first phase it runs any efficient first-order method FO for T_1 steps to shrink the function value to the regime where we can then show linear convergence for our algorithm. It then takes approximate Newton steps based on the estimator from Definition 5 in place of the true Hessian inverse. We use two parameters, S_1 and S_2 , to define the Newton step. S_1 represents the number of unbiased estimators of the Hessian inverse we average to get better concentration for our estimator, while S_2 represents the depth to which we capture the Taylor expansion. In the algorithm, we compute the average Newton step directly, which can be computed in linear time as observed earlier, instead of estimating the Hessian inverse.

3.3 Main Theorem

In this section we present our main theorem which analyzes the convergence properties of LiSSA. Define $FO(M, \hat{\kappa}_l)$ to be the total time required by a first-order algorithm to achieve accuracy $\frac{1}{4\hat{\kappa}_l M}$.

Theorem 7 *Consider Algorithm 1, and set the parameters as follows: $T_1 = FO(M, \hat{\kappa}_l)$, $S_1 = O((\hat{\kappa}_l^{\max})^2 \ln(\frac{d}{\delta}))$, $S_2 \geq 2\hat{\kappa}_l \ln(4\hat{\kappa}_l)$. The following guarantee holds for every $t \geq T_1$ with probability $1 - \delta$,*

$$\|\mathbf{x}_{t+1} - \mathbf{x}^*\| \leq \frac{\|\mathbf{x}_t - \mathbf{x}^*\|}{2}.$$

Moreover, we have that each step of the algorithm takes at most $\tilde{O}(md + (\hat{\kappa}_l^{\max})^2 \hat{\kappa}_l d^2)$ time. Additionally if f is GLM, then each step of the algorithm can be run in time $\tilde{O}(md + (\hat{\kappa}_l^{\max})^2 \hat{\kappa}_l d)$.

As an immediate corollary, we obtain the following:

Corollary 8 *For a GLM function $f(\mathbf{x})$, Algorithm 1 returns a point \mathbf{x}_t such that with probability at least $1 - \delta$,*

$$f(\mathbf{x}_t) \leq \min_{\mathbf{x}^*} f(\mathbf{x}^*) + \varepsilon$$

in total time $\tilde{O}(m + (\hat{\kappa}_l^{\max})^2 \hat{\kappa}_l d \ln(\frac{1}{\varepsilon}))$ for $\varepsilon \rightarrow 0$.

In the above theorems \tilde{O} hides log factors of $\kappa, d, \frac{1}{\delta}$. We note that our bound $(\hat{\kappa}_l^{\max})^2$ on the variance is possibly pessimistic and can likely be improved to a more average quantity. However, since in our experiments setting the parameter $S_1 \sim O(1)$ suffices, we have not tried to optimize it further.

We now prove our main theorem about the convergence of LiSSA (Theorem 7).

Proof [Proof of Theorem 7]

Note that since we use a first-order algorithm to get a solution of accuracy at least $\frac{1}{4\hat{\kappa}_l M}$, we have that

$$\|\mathbf{x}_1 - \mathbf{x}^*\| \leq \frac{1}{4\hat{\kappa}_l M}. \quad (3)$$

As can be seen from Definition 5, a single step of our algorithm is equivalent to $\mathbf{x}_{t+1} = \mathbf{x}_t - \tilde{\nabla}^{-2} f(\mathbf{x}_t) \nabla f(\mathbf{x}_t)$, where $\tilde{\nabla}^{-2} f(\mathbf{x}_t)$ is the average of S_1 independent estimators $\tilde{\nabla}^{-2} f(\mathbf{x}_t)_{S_2}$. We now make use of the following lemma.

Lemma 9 *Let $\mathbf{x}_{t+1} = \mathbf{x}_t - \tilde{\nabla}^{-2} f(\mathbf{x}_t) \nabla f(\mathbf{x}_t)$, as per a single iteration of Algorithm 1, and suppose S_1, S_2 are as defined in Algorithm 1. Then if we choose $S_2 \geq 2\hat{\kappa}_l \ln(2\hat{\kappa}_l)$ we have the following guarantee on the convergence rate for every step with probability $1 - \delta$:*

$$\|\mathbf{x}_{t+1} - \mathbf{x}^*\| \leq \gamma \|\mathbf{x}_t - \mathbf{x}^*\| + M \|\nabla^{-2} f(\mathbf{x}_t)\| \|\mathbf{x}_t - \mathbf{x}^*\|^2$$

where $\gamma = 16\hat{\kappa}_l^{\max} \sqrt{\frac{\ln(d\delta^{-1})}{S_1}} + \frac{1}{16}$.

Substituting the values of S_1 and S_2 , combining Equation (3) and Lemma 9, and noting that $\|\nabla^{-2} f(\mathbf{x}_t)\| \leq \hat{\kappa}_l$, we have that at the start of the Newton phase the following inequality holds:

$$\|\mathbf{x}_{t+1} - \mathbf{x}^*\| \leq \frac{\|\mathbf{x}_t - \mathbf{x}^*\|}{4} + M\hat{\kappa}_l^{\max} \|\mathbf{x}_t - \mathbf{x}^*\|^2 \leq \frac{\|\mathbf{x}_t - \mathbf{x}^*\|}{2}.$$

It can be shown via induction that the above property holds for all $t \geq T_1$, which concludes the proof. \blacksquare

We now provide a proof of Lemma 9.

Proof [Proof of Lemma 9] Define $\chi(\mathbf{x}_t) = \int_0^1 \nabla^2 f(\mathbf{x}^* + \tau(\mathbf{x}_t - \mathbf{x}^*)) d\tau$. Note that $\nabla f(\mathbf{x}_t) = \chi(\mathbf{x}_t)(\mathbf{x}_t - \mathbf{x}^*)$. Following an analysis similar to that of Nesterov (2013), we have that

$$\begin{aligned} \|\mathbf{x}_{t+1} - \mathbf{x}^*\| &= \|\mathbf{x}_t - \mathbf{x}^* - \tilde{\nabla}^{-2} f(\mathbf{x}_t) \nabla f(\mathbf{x}_t)\| \\ &= \|\mathbf{x}_t - \mathbf{x}^* - \tilde{\nabla}^{-2} f(\mathbf{x}_t) \chi(\mathbf{x}_t) (\mathbf{x}_t - \mathbf{x}^*)\| \\ &\leq \|I - \tilde{\nabla}^{-2} f(\mathbf{x}_t) \chi(\mathbf{x}_t)\| \|\mathbf{x}_t - \mathbf{x}^*\|. \end{aligned}$$

Following from the previous equations, we have that

$$\frac{\|\mathbf{x}_{t+1} - \mathbf{x}^*\|}{\|\mathbf{x}_t - \mathbf{x}^*\|} \leq \|I - \tilde{\nabla}^{-2} f(\mathbf{x}_t) \chi(\mathbf{x}_t)\| = \underbrace{\|I - \nabla^{-2} f(\mathbf{x}_t) \chi(\mathbf{x}_t)\|}_a - \underbrace{\left(\tilde{\nabla}^{-2} f(\mathbf{x}_t) - \nabla^{-2} f(\mathbf{x}_t) \right) \chi(\mathbf{x}_t)}_b.$$

We now analyze the above two terms a and b separately:

$$\begin{aligned} \|a\| &= \|I - \nabla^{-2} f(\mathbf{x}_t) \chi(\mathbf{x}_t)\| \\ &\leq \|\nabla^{-2} f(\mathbf{x}_t) \int_0^1 (\nabla^2 f(\mathbf{x}_t) - \nabla^2 f(\mathbf{x}^* + \tau(\mathbf{x}_t - \mathbf{x}^*))) d\tau\| \\ &\leq M \|\nabla^{-2} f(\mathbf{x}_t)\| \|\mathbf{x}_t - \mathbf{x}^*\|. \end{aligned}$$

The second inequality follows from the Lipschitz bound on the Hessian. The second term can be bounded as follows:

$$\|b\| \leq \left(\left\| \left(\tilde{\nabla}^{-2} f(\mathbf{x}_t) - \nabla^{-2} f(\mathbf{x}_t) \right) \right\| \|\chi(\mathbf{x}_t)\| \right) \leq \gamma.$$

The previous claim follows from Lemma 10 which shows a concentration bound on the sampled estimator and by noting that due to our assumption on the function, we have that for all \mathbf{x} , $\|\nabla^2 f(\mathbf{x})\| \leq 1$, and hence $\|\chi(\mathbf{x})\| \leq 1$.

Putting the above two bounds together and using the triangle inequality, we have that

$$\frac{\|\mathbf{x}_{t+1} - \mathbf{x}^*\|}{\|\mathbf{x}_t - \mathbf{x}^*\|} \leq M \|\nabla^{-2} f(\mathbf{x}_t)\| \|\mathbf{x}_t - \mathbf{x}^*\| + \gamma$$

which concludes the proof. ■

Lemma 10 *Let $\tilde{\nabla}^{-2} f(\mathbf{x}_t)$ be the average of S_1 independent samples of $\tilde{\nabla}^{-2} f(\mathbf{x}_t)_{S_2}$, as defined in Definition 5 and used in the per-step update of Algorithm 1, and let $\nabla^2 f(\mathbf{x}_t)$ be the true Hessian. If we set $S_2 \geq 2\hat{\kappa}_l \ln(\hat{\kappa}_l S_1)$, then we have that*

$$\Pr \left(\left\| \tilde{\nabla}^{-2} f(\mathbf{x}_t) - \nabla^{-2} f(\mathbf{x}_t) \right\| > 16\hat{\kappa}_l^{\max} \sqrt{\frac{\ln(\frac{d}{\delta})}{S_1}} + 1/16 \right) \leq \delta.$$

Proof [Proof of Lemma 10] First note the following statement which is a straightforward implication of our construction of the estimator:

$$\mathbf{E}[\tilde{\nabla}^{-2} f(\mathbf{x}_t)] = \mathbf{E}[\tilde{\nabla}^{-2} f(\mathbf{x}_t)_{S_2}] = \sum_{i=0}^{S_2} (I - \nabla^2 f(\mathbf{x}_t))^i.$$

We also know from Equation (2) that for matrices X such that $\|X\| \leq 1$ and $X \succ 0$,

$$X^{-1} = \sum_{i=0}^{\infty} (I - X)^i.$$

Since we have scaled the function such that $\|\nabla^2 f_k\| \leq 1$, it follows that

$$\nabla^{-2} f(\mathbf{x}_t) = \mathbf{E} \left[\tilde{\nabla}^{-2} f(\mathbf{x}_t)_{S_2} \right] + \sum_{i=S_2+1}^{\infty} (I - \nabla^2 f(\mathbf{x}_t))^i. \quad (4)$$

Also note that since $\nabla^2 f(\mathbf{x}_t) \succeq \frac{I}{\hat{\kappa}_l}$, it follows that $\|I - \nabla^2 f(\mathbf{x}_t)\| \leq 1 - \frac{1}{\hat{\kappa}_l}$. Observing the second term in the above equation,

$$\begin{aligned} \left\| \sum_{i=S_2}^{\infty} (I - \nabla^2 f(\mathbf{x}_t))^i \right\| &\leq \| (I - \nabla^2 f(\mathbf{x}_t)) \|^{S_2} \left(\sum_{i=0}^{\infty} \|I - \nabla^2 f(\mathbf{x}_t)\|^i \right) \\ &\leq \left(1 - \frac{1}{\hat{\kappa}_l}\right)^{S_2} \left(\sum_{i=0}^{\infty} \left(1 - \frac{1}{\hat{\kappa}_l}\right)^i \right) \\ &\leq \left(1 - \frac{1}{\hat{\kappa}_l}\right)^{S_2 \hat{\kappa}_l} \\ &\leq \exp\left(-\frac{S_2}{\hat{\kappa}_l}\right) \hat{\kappa}_l. \end{aligned}$$

Since we have chosen $S_2 \geq 2\hat{\kappa}_l \ln(4\hat{\kappa}_l)$, we get that the above term is bounded by $\frac{1}{16}$. We will now show, using the matrix Bernstein inequality (Theorem 3), that the estimate $\tilde{\nabla}^{-2} f$ is concentrated around its expectation. To apply the inequality we first need to bound the spectral norm of each random variable. To that end we note that $\tilde{\nabla}^{-2} f_{S_2}$ has maximum spectral norm bounded by

$$\|\tilde{\nabla}^{-2} f_{S_2}\| \leq \sum_{i=0}^{S_2} (1 - 1/\hat{\kappa}_l^{\max})^i \leq \hat{\kappa}_l^{\max}.$$

We can now apply Theorem 3, which gives the following:

$$\Pr \left(\|\tilde{\nabla}^{-2} f - \mathbf{E}[\tilde{\nabla}^{-2} f]\| > \varepsilon \right) \leq 2d \exp\left(\frac{-\varepsilon^2 S_1}{64(\hat{\kappa}_l^{\max})^2}\right).$$

Setting $\varepsilon = 16\hat{\kappa}_l^{\max} \sqrt{\frac{\ln(\frac{d}{\delta})}{S_1}}$ gives us that the probability above is bounded by δ . Now putting together the bounds and Equation (4) we get the required result. \blacksquare

3.4 Leveraging Sparsity

A key property of real-world data sets is that although the input is a high dimensional vector, the number of non-zero entries is typically very low. The following theorem shows

that LiSSA can be implemented in a way to leverage the underlying sparsity of the data. Our key observation is that for GLM functions, the rank one Hessian-vector product can be performed in $O(s)$ time where s is the sparsity of the input \mathbf{x}_k .

Theorem 11 *For GLM functions Algorithm 1 returns a point \mathbf{x}_t such that with probability at least $1 - \delta$*

$$f(\mathbf{x}_t) \leq \min_{\mathbf{x}^*} f(\mathbf{x}^*) + \varepsilon$$

in total time $\tilde{O}(ms + (\hat{\kappa}_l^{\max})^2 \hat{\kappa}_l s) \ln(\frac{1}{\varepsilon})$ for $\varepsilon \rightarrow 0$.

We will prove the following theorem, from which Theorem 11 will immediately follow.

Theorem 12 *Consider Algorithm 1, let f be of the form described above, and let s be such that the number of non zero entries in \mathbf{x}_i is bounded by s . Then each step of the algorithm can be implemented in time $O(ms + (\kappa_l^{\max})^2 \kappa_l s)$.*

Proof [Proof of Theorem 12] Proof by induction. Let $c_0 = 1$, $d_0 = 1$, $\mathbf{v}_0 = \mathbf{0}$, and consider the update rules $c_{j+1} = 1 + (1 - \lambda)c_j$, $d_{j+1} = (1 - \lambda)d_j$, and $\mathbf{v}_{j+1} = \mathbf{v}_j - \frac{1}{(1-\lambda)d_j} \tilde{\nabla}^2 f_{[i,j+1]}(\mathbf{x})(c_j \nabla f(\mathbf{x}) + d_j \mathbf{v}_j)$, where λ is the regularization parameter. For the base case, note that $X_{[i,0]} = c_0 \nabla f(\mathbf{x}) + d_0 \mathbf{v}_0 = \nabla f(\mathbf{x})$, as is the case in Algorithm 1. Furthermore, suppose $X_{[i,j]} = c_j \nabla f(\mathbf{x}) + d_j \mathbf{v}_j$. Then we see that

$$\begin{aligned} X_{[i,j+1]} &= \nabla f(\mathbf{x}) + (I - \lambda I - \tilde{\nabla}^2 f_{[i,j+1]}(\mathbf{x})) X_{[i,j]} \\ &= \nabla f(\mathbf{x}) + ((1 - \lambda)I - \tilde{\nabla}^2 f_{[i,j+1]}(\mathbf{x}))(c_j \nabla f(\mathbf{x}) + d_j \mathbf{v}_j) \\ &= (1 + (1 - \lambda)c_j) \nabla f(\mathbf{x}) + (1 - \lambda)(d_j \mathbf{v}_j) - \tilde{\nabla}^2 f_{[i,j+1]}(\mathbf{x})(c_j \nabla f(\mathbf{x}) + d_j \mathbf{v}_j) \\ &= c_{j+1} \nabla f(\mathbf{x}) + d_{j+1} \mathbf{v}_{j+1}. \end{aligned}$$

Note that updating c_{j+1} and d_{j+1} takes constant time, and $\tilde{\nabla}^2 f_{[i,j+1]}(\mathbf{x})(c_j \nabla f(\mathbf{x}))$ and $\tilde{\nabla}^2 f_{[i,j+1]}(\mathbf{x}) \mathbf{v}_j$ can each be calculated in $O(s)$ time. It can also be seen that each product gives an s -sparse vector, so computing \mathbf{v}_{j+1} takes time $O(s)$. Since $\nabla f(\mathbf{x})$ can be calculated in $O(ms)$ time, and since \mathbf{v}_0 is 0-sparse which implies the number of non-zero entries of \mathbf{v}_j is at most js , it follows that the total time to calculate X_t is $O(ms + (\kappa_l^{\max})^2 \kappa_l s)$. \blacksquare

4. LiSSA: Extensions

In this section we first describe a family of algorithms which generically couple first-order methods as sub-routines with second-order methods. In particular, we formally describe the algorithm LiSSA-Quad (Algorithm 2) and provide its runtime guarantee (Theorem 13). The key idea underlying this algorithm is that Newton's method essentially reduces a convex optimization problem to solving intermediate quadratic sub-problems given by the second-order Taylor approximation at a point, i.e., the sub-problem Q_t given by

$$Q_t(\mathbf{y}) = f(\mathbf{x}_{t-1}) + \nabla f(\mathbf{x}_{t-1})^T \mathbf{y} + \frac{\mathbf{y}^T \nabla^2 f(\mathbf{x}_{t-1}) \mathbf{y}}{2}$$

Algorithm 2 LiSSA-Quad

Input: $T, f(\mathbf{x}) = \sum_{k=1}^m f_k(\mathbf{x}), ALG, ALG_{params}, T_1, \varepsilon$
 $\mathbf{x}_0 = ALG(f(\mathbf{x}), T_1)$
for $t = 1$ to T **do**
 $Q_t(\mathbf{y}) = \nabla f(\mathbf{x}_{t-1})^T \mathbf{y} + \frac{\mathbf{y}^T \nabla^2 f(\mathbf{x}_{t-1}) \mathbf{y}}{2}$
 $\mathbf{x}_t = A(Q_t, \varepsilon^2, A_{params})$
end for
return \mathbf{x}_T

where $\mathbf{y} \triangleq \mathbf{x} - \mathbf{x}_{t-1}$. The above ideas provide an alternative implementation of our estimator for $\nabla^{-2} f(\mathbf{x})$ used in LiSSA. Consider running gradient descent on the above quadratic Q_t , and let \mathbf{y}_t^i be the i^{th} step in this process. By definition we have that

$$\mathbf{y}_t^{i+1} = \mathbf{y}_t^i - \nabla Q_t(\mathbf{y}_t^i) = (I - \nabla^2 f(\mathbf{x}_{t-1})) \mathbf{y}_t^i - \nabla f(\mathbf{x}_{t-1}).$$

It can be seen that the above expression corresponds exactly to the steps taken in LiSSA (Algorithm 2, line 8), the difference being that we use a sample of the Hessian instead of the true Hessian. Therefore LiSSA can also be interpreted as doing a partial stochastic gradient descent on the quadratic Q_t . It is partial because we have a precise estimate of gradient of the function f and a stochastic estimate for the Hessian. We note that this is essential for the linear convergence guarantees we show for LiSSA.

The above interpretation indicates the possibility of using any first-order linearly convergent scheme for approximating the minimizer of the quadratic Q_t . In particular, consider any algorithm ALG that, given a convex quadratic function Q_t and an error value ε , produces a point \mathbf{y} such that

$$\|\mathbf{y} - \mathbf{y}_t^*\| \leq \varepsilon \tag{5}$$

with probability at least $1 - \delta_{ALG}$, where $\mathbf{y}_t^* = \operatorname{argmin} Q_t$. Let the total time taken by the algorithm ALG to produce the point be $T_{ALG}(\varepsilon, \delta_{ALG})$. For our applications we require ALG to be linearly convergent, i.e. T_{ALG} is proportional to $\log(\frac{1}{\varepsilon})$ with probability at least $1 - \delta_{ALG}$.

Given such an algorithm ALG , LiSSA-Quad, as described in Algorithm 2, generically implements the above idea, modifying LiSSA by replacing the inner loop with the given algorithm ALG . The following is a meta-theorem about the convergence properties of LiSSA-Quad.

Theorem 13 *Given the function $f(\mathbf{x}) = \sum f_i(\mathbf{x})$ which is α -strongly convex, let \mathbf{x}^* be the minimizer of the function and $\{\mathbf{x}_t\}$ be defined as in Algorithm 2. Suppose the algorithm ALG satisfies condition (5) with probability $1 - \delta_{ALG}$ under the appropriate setting of parameters ALG_{params} . Set the parameters in the algorithm as follows: $T_1 = T_{ALG}(1/4\alpha M)$, $T = \log \log(1/\varepsilon)$, $\delta_{ALG} = \delta/T$, where ε is the final error guarantee one wishes to achieve. Then we have that after T steps, with probability at least $1 - \delta$,*

$$\min_{t=\{1 \dots T\}} \|\mathbf{x}_t - \mathbf{x}^*\| \leq \varepsilon.$$

In particular, *LiSSA-Quad(ALG)* produces a point \mathbf{x} such that

$$\|\mathbf{x} - \mathbf{x}^*\| \leq \varepsilon$$

in total time $O(T_{ALG}(\varepsilon, \delta_{ALG}) \log \log(1/\varepsilon))$ with probability at least $1 - \delta$ for $\varepsilon \rightarrow 0$.

Note that for GLM functions, the $\nabla Q_t(\mathbf{y})$ at any point can be computed in time linear in d . In particular, a full gradient of Q_t can be computed in time $O(md)$ and a stochastic gradient (corresponding to a stochastic estimate of the Hessian) in time $O(d)$. Therefore, a natural choice for the algorithm *ALG* in the above are first-order algorithms which are linearly convergent, for example SVRG, SDCA, and Acc-SDCA. Choosing a first-order algorithm FO gives us a family of algorithms *LiSSA-Quad(FO)*, each with running time comparable to the running time of the underlying FO, up to logarithmic factors. The following corollary summarizes the typical running time guarantees for *LiSSA-Quad(FO)* when FO is Acc-SVRG.

Corollary 14 *Given a GLM function $f(\mathbf{x})$, if *ALG* is replaced by Acc-SVRG (Lin et al., 2015), then under a suitable setting of parameters, *LiSSA-Quad* produces a point \mathbf{x} such that*

$$f(\mathbf{x}) - f(\mathbf{x}^*) \leq \varepsilon$$

with probability at least $1 - \delta$, in total time $\tilde{O}(m + \min\{\sqrt{\hat{\kappa}_l m}, \hat{\kappa}_l\})d \log(1/\varepsilon) \log \log(1/\varepsilon)$.

Here the \tilde{O} above hides logarithmic factors in κ, d, δ , but not in ε . Note that the above running times depend upon the condition number $\hat{\kappa}_l$ which as described in Section 2 can potentially provide better dependence compared to its global counterpart. In practice this difference could lead to faster running time for *LiSSA-Quad* as compared to the underlying first-order algorithm FO. We now provide a proof for Theorem 13.

Proof [Proof of Theorem 13] We run the algorithm *A* to achieve accuracy ε^2 on each of the intermediate quadratic functions Q_t , and we set $\delta_A = \delta/T$ which implies via a union bound that for all $t \leq T$,

$$\|\mathbf{x}_{t+1} - \mathbf{x}_t^*\| \leq \varepsilon^2 \tag{6}$$

with probability at least $1 - \delta$.

Assume that for all $t < T$, $\|\mathbf{x}_t - \mathbf{x}^*\| \geq \varepsilon$ (otherwise the theorem is trivially true). Using the analysis of Newton's method as before, we get that for all $t \leq T$,

$$\begin{aligned} \|\mathbf{x}_{t+1} - \mathbf{x}^*\| &\leq \|\mathbf{x}_t^* - \mathbf{x}^*\| + \|\mathbf{x}_{t+1} - \mathbf{x}_t^*\| \\ &\leq \|\mathbf{x}_t - \nabla^{-2} f(\mathbf{x}_t) \nabla f(\mathbf{x}_t) - \mathbf{x}^*\| + \|\mathbf{x}_{t+1} - \mathbf{x}_t^*\| \\ &\leq \frac{M}{4\alpha} \|\mathbf{x}_t - \mathbf{x}^*\|^2 + \varepsilon^2 \\ &\leq \left(\frac{M}{4\alpha} + 1 \right) \|\mathbf{x}_t - \mathbf{x}^*\|^2 \end{aligned}$$

where the second inequality follows from the analysis in the proof of Theorem 7 and Equation (6). We know that $\|\mathbf{x}_0 - \mathbf{x}^*\| \leq \sqrt{\frac{\alpha}{M}}$ from the initial run of the first-order algorithm *FO*. Applying the above inductively and using the value of T prescribed by the theorem statement, we get that $\|\mathbf{x}_T - \mathbf{x}^*\| \leq \varepsilon$. ■

5. Runtime Improvement through Fast Quadratic Solvers

The previous section establishes the reduction from general convex optimization to quadratic functions. In this section we show how we can leverage the fact that for quadratic functions the running time for accelerated first-order methods can be improved in the regime when $\kappa > m \gg d$. In particular, we show the following theorem.

Theorem 15 *Given a vector $\mathbf{b} \in \mathbb{R}^d$ and a matrix $A = \sum A_i$ where each A_i is of the form $A_i = \mathbf{v}_i \mathbf{v}_i^T + \lambda I$ for some $\mathbf{v}_i \in \mathbb{R}^d$, $\|\mathbf{v}_i\| \leq 1$ and $\lambda \geq 0$ a fixed parameter, Algorithm 4 computes a vector $\tilde{\mathbf{v}}$ such that $\|A^{-1}\mathbf{b} - \tilde{\mathbf{v}}\| \leq \varepsilon$ with probability at least $1 - \delta$ in total time*

$$\tilde{O} \left(md \log \left(\frac{1}{\varepsilon} \right) + \left(d + \sqrt{\kappa_{\text{sample}}(A)d} \right) d \log^2 \left(\frac{1}{\varepsilon} \right) \right).$$

$\tilde{O}()$ contains factors logarithmic in $m, d, \kappa_{\text{sample}}(A), \|b\|, \delta$.

$\kappa_{\text{sample}}(A)$ is the condition number of an $O(d \log(d))$ sized sample of A and is formally defined in Equation (11). We can now use Algorithm 4 to compute an approximate Newton step by setting $A = \nabla^2 f(x)$ and $\mathbf{b} = \nabla f(x)$. We therefore propose LiSSA-Sample to be a variant of LiSSA-Quad where Algorithm 4 is used as the subroutine ALG and any first-order algorithm can be used in the initial phase. The following theorem bounding the running time of LiSSA-Sample follows immediately from Theorem 13 and Theorem 15.

Theorem 16 *Given a GLM function $f(\mathbf{x}) = \sum_i f_i(\mathbf{x})$, let $\mathbf{x}^* = \text{argmin} f(\mathbf{x})$. LiSSA-Sample produces a point \mathbf{x} such that*

$$\|\mathbf{x} - \mathbf{x}^*\| \leq \varepsilon$$

with probability at least $1 - \delta$ in total time

$$\tilde{O} \left(\left(md \log \left(\frac{1}{\varepsilon} \right) + \left(d + \sqrt{\kappa_{\text{sample}}(f)d} \right) d \log^2 \left(\frac{1}{\varepsilon} \right) \right) \log \log \left(\frac{1}{\varepsilon} \right) \right).$$

$\tilde{O}()$ contains factors logarithmic in $m, d, \kappa_{\text{sample}}(f), G, \delta$.

5.1 Fast Quadratic Solver - Outline

In this section we provide a short overview of Algorithm 4. To simplify the discussion, let's consider the case when we have to compute $A^{-1}\mathbf{b}$ for a $d \times d$ matrix A given as $A = \sum_{i=1}^m \mathbf{v}_i \mathbf{v}_i^T = VV^T$ where the i^{th} column of V is \mathbf{v}_i . The computation can be recast as minimization of a convex quadratic function $Q(\mathbf{y}) = \frac{\mathbf{y}^T A \mathbf{y}}{2} + \mathbf{b}^T \mathbf{y}$ and can be solved up to accuracy ε in total time $\left(m + \sqrt{\kappa(A)m} \right) d \log(1/\varepsilon)$ as can be seen from Theorem 4. Algorithm 4 improves upon the running time bound in the case when $m > d$. In the following we provide a high level outline of the procedure which is formally described as Algorithm 4.

- Given A we will compute a low complexity constant spectral approximation B of A . Specifically $B = \sum_{i=1}^{O(d \log(d))} \mathbf{u}_i \mathbf{u}_i^T$ and $B \preceq A \preceq 2B$. This is achieved by techniques developed in matrix sampling/sketching literature, especially those of Cohen et al. (2015). The procedure requires solving a constant number of $O(d \log(d))$ sized linear systems, which we do via Accelerated SVRG.
- We use B as a preconditioner and compute $BA^{-1}\mathbf{b}$ by minimizing the quadratic $\frac{\mathbf{y}^T AB^{-1}\mathbf{y}}{2} + \mathbf{b}^T \mathbf{y}$. Note that this quadratic is well conditioned and can be minimized using gradient descent. In order to compute the gradient of the quadratic which is given by $AB^{-1}\mathbf{y}$, we again use Accelerated SVRG to solve a linear system in B .
- Finally, we compute $A^{-1}\mathbf{b} = B^{-1}BA^{-1}\mathbf{b}$ using Accelerated SVRG to solve a linear system in B .

In the rest of the section we formally describe the procedure outlined and provide the necessary definitions. One key nuance we must take into account is the fact that based on our assumption we have included the regularization term into the component functions. Due to this, the Hessian does not necessarily look like a sum of rank one matrices. Of course, one can decompose the identity matrix that appears due to the regularizer as a sum of rank one matrices. However, note that the procedure above requires that each of the sub-samples must have good condition number too in order to solve linear systems on them with Accelerated SVRG. Therefore, the sub-samples generated must look like sub-samples formed from the Hessians of component functions. For this purpose we extend the procedure and the definition for leverage scores described by Cohen et al. (2015) to the case when the matrix is given as a sum of PSD matrices and not just rank one matrices. We reformulate and reprove the basic theorems proved by Cohen et al. (2015) in this context. To maintain computational efficiency of the procedure, we then make use of the fact that each of the PSD matrices actually is a rank one matrix plus the Identity matrix. We now provide the necessary preliminaries for the description of the algorithm and its analysis.

5.2 Preliminaries for Fast Quadratic Solver

For all the definitions and preliminaries below assume we are given a $d \times d$ PSD matrix $A \triangleq \sum_{i=1}^m A_i$ where A_i are also PSD matrices. Let $A \cdot B \triangleq Tr(B^T A)$ be the standard matrix dot product. Given two matrices A and B we say B is a λ -spectral approximation of A if $\frac{1}{\lambda}A \preceq B \preceq A$.

Definition 17 (Generalized Matrix Leverage Scores) *Define*

$$\tau_i(A) \triangleq A^+ \cdot A_i \tag{7}$$

$$\tau_i^B(A) \triangleq B^+ \cdot A_i. \tag{8}$$

Then we have the following facts:

Fact 18

$$\sum_{i=1}^n \tau_i(A) = Tr \left(\sum A^+ A_i \right) = Tr(A^+ A) = rank(A) \leq d. \tag{9}$$

Fact 19 *If B is a λ -spectral approximation of A , then $\tau_i(A) \leq \tau_i^B(A) \leq \lambda\tau_i(A)$.*

Given A of the form $A = \sum A_i$, we define a sample of A of size r as the following. Consider a subset of indices of size r , $I = \{i_1 \dots i_r\} \subseteq [m]$. For every such sample I and given a weight vector $\mathbf{w} \in \mathbb{R}^r$, we can associate the following matrix:

$$\text{Sample}(\mathbf{w}, I) \triangleq \sum_{j \in I} w_j A_j. \quad (10)$$

When the sample is unweighted, i.e., $\mathbf{w} = \mathbf{1}$, we will simply denote the above as $\text{Sample}(I)$. We can now define $\kappa_{\text{sample}}(A, r)$ to be

$$\kappa_{\text{sample}}(A, r) \triangleq \max_{I: |I| \geq r} \kappa(\text{Sample}(r, I)). \quad (11)$$

This is the worst-case condition number of an unweighted sub-sample of size r of the matrix A . For our results we will be concerned with a $\Omega(d \log(d))$ sized sub-sample, i.e., the quantity $\kappa_{\text{sample}}(A, O(d \log(d)))$. We remind the reader that by definition κ_{sample} for Hessians of the functions we are concerned with is always larger than $1/\lambda$, where λ is the coefficient of the ℓ_2 regularizer.

The following lemma is a generalization of the leverage score sampling lemma (Lemma 4, Cohen et al. 2015). The proof is very similar to the original proof by Cohen et al. (2015) and is included in the Appendix for completeness.

Lemma 20 (Spectral Approximation via Leverage Score Sampling) *Given an error parameter $0 < \varepsilon < 1$, let \mathbf{u} be a vector of leverage score overestimates, i.e., $\tau_i(A) \leq \mathbf{u}_i$, for all $i \in [m]$. Let α be a sampling rate parameter and let c be a fixed positive constant. For each matrix A_i , we define a sampling probability $p_i(\alpha) = \min\{1, \alpha \mathbf{u}_i c \log d\}$. Let I be a random sample of indices drawn from $[m]$ by sampling each index with probability $p_i(\alpha)$. Define the weight vector $\mathbf{w}(\alpha)$ to be the vector such that $\mathbf{w}(\alpha)_i = \frac{1}{p_i(\alpha)}$. By definition of weighted samples we have that*

$$\text{Sample}(\mathbf{w}(\alpha), I) = \sum_{i=1}^m \frac{1}{p_i(\alpha)} A_i \mathbf{1}_{x_i \sim p_i(\alpha)}(x_i = 1)$$

where x_i is a Bernoulli random variable with probability $p_i(\alpha)$.

If we set $\alpha = \varepsilon^{-2}$, $S = \text{Sample}(\mathbf{w}(\alpha), I)$ is formed by at most $\sum_i \min\{1, \alpha \mathbf{u}_i c \log(d)\} \leq \alpha c \log(d) \|\mathbf{u}\|_1$ entries in the above sum, and $\frac{1}{1+\varepsilon} S$ is a $\frac{1+\varepsilon}{1-\varepsilon}$ spectral approximation for A with probability at least $1 - d^{-c/3}$.

The following theorem is an analogue of the key theorem regarding uniform sampling (Theorem 1, Cohen et al. 2015). The proof is identical to the original proof and is also included in the appendix for completeness.

Theorem 21 (Uniform Sampling) *Given any $A = \sum_{i=1}^m A_i$ as defined above, let $S = \sum_{j=1}^r X_j$ be formed by uniformly sampling r matrices $X_1 \dots X_r \sim \{A_i\}$ without repetition. Define*

$$\tilde{\tau}_i^S(A) = \begin{cases} \tau_i^S(A) & \text{if } \exists j \text{ s.t. } X_j = A_i \\ \tau_i^{S+A_i}(A) & \text{otherwise} \end{cases}.$$

Then $\tilde{\tau}_i^S(A) \geq \tau_i(A)$ for all i , and

$$\mathbb{E} \left[\sum_{i=1}^n \tilde{\tau}_i^S(A) \right] \leq \frac{md}{r}.$$

Unlike the case of rank one matrices as was done by Cohen et al. (2015), it is not immediately clear how to efficiently compute $\tilde{\tau}_i$ as we cannot directly apply the Sherman-Morrison formula. In our case though, since we have that each $A_i = \mathbf{v}_i \mathbf{v}_i^T + \lambda I$ we define slightly different estimates which are efficiently computable and prove that they work as well.

Theorem 22 *Suppose we are given any $A = \sum_{i=1}^m A_i$ where each A_i is of the form $A_i = \mathbf{v}_i \mathbf{v}_i^T + \lambda I$. Let $S = \sum_{j=1}^m X_j$ be formed by uniformly sampling r matrices $X_1 \dots X_r \sim \{A_i\}$ without repetition. Define*

$$\hat{\tau}_i^S(A) = \begin{cases} \mathbf{v}_i^T S^+ \mathbf{v}_i + \frac{d}{r} & \text{if } \exists j \text{ s.t. } X_j = A_i \\ \frac{1}{1 + \frac{1}{\mathbf{v}_i^T (S + \lambda I) + \mathbf{v}_i^T}} + \frac{d}{r} & \text{otherwise} \end{cases}.$$

Then $\hat{\tau}_i^S(A) \geq \tau_i(A)$ for all i , and

$$\mathbb{E} \left[\sum_{i=1}^n \hat{\tau}_i^S(A) \right] \leq O \left(\frac{md}{r} \right).$$

Proof [Proof of Theorem 22] We will prove that

$$\frac{d}{m} \geq \hat{\tau}_i^S(A) - \tilde{\tau}_i^S(A) \geq 0. \quad (12)$$

Consider the case when $\exists j$ s.t. $X_j = A_i$. By definition of $\tilde{\tau}_i^S(A)$,

$$\begin{aligned} \tilde{\tau}_i^S(A) &= \tau_i^S(A) = S^+ \cdot A_i = S^+ \cdot (\mathbf{v}_i \mathbf{v}_i^T + \lambda I) = \mathbf{v}_i^T S^+ \mathbf{v}_i + \lambda S^+ \cdot \lambda I \\ &\leq \mathbf{v}_i^T S^+ \mathbf{v}_i + (r\lambda I)^+ \cdot I = \mathbf{v}_i^T S^+ \mathbf{v}_i + \frac{d}{r} = \hat{\tau}_i^S(A). \end{aligned}$$

The above follows by noting that $S \succeq r\lambda I$. It also follows from the definition that $\hat{\tau}_i^S(A) - \tilde{\tau}_i^S(A) \leq \frac{d}{r}$.

In the other case a similar inequality can be shown by noting via the Sherman-Morrison formula that

$$\mathbf{v}_i^T (S + \lambda I + \mathbf{v}_i \mathbf{v}_i^T)^+ \mathbf{v}_i = \mathbf{v}_i^T \left((S + \lambda I)^+ - \frac{(S + \lambda I)^+ \mathbf{v}_i \mathbf{v}_i^T (S + \lambda I)^+}{1 + \mathbf{v}_i^T (S + \lambda I)^+ \mathbf{v}_i} \right) \mathbf{v}_i = \frac{1}{1 + \frac{1}{\mathbf{v}_i^T (S + \lambda I)^+ \mathbf{v}_i}}.$$

This proves Equation (12). A direct application of Theorem 21 now finishes the proof. \blacksquare

Algorithm 3 REPEATED HALVING

- 1: **Input:** $A = \sum_{i=1}^m (\mathbf{v}_i \mathbf{v}_i^T + \lambda I)$
- 2: **Output:** B an $O(d \log(d))$ size weighted sample of A and $B \preceq A \preceq 2B$
- 3: Take a uniformly random unweighted sample of size $\frac{m}{2}$ of A to form A'
- 4: **if** A' has size $> O(d \log(d))$ **then**
- 5: Recursively compute an 2-spectral approximation \tilde{A}' of A'
- 6: **end if**
- 7: Compute estimates γ_i of generalized leverage scores $\{\hat{\tau}_i^{A'}(A)\}$ s.t. the following are satisfied

$$\begin{aligned} \gamma_i &\geq \hat{\tau}_i^{A'}(A) \\ \sum \gamma_i &\leq \sum 16 \hat{\tau}_i^{A'}(A) + 1 \end{aligned}$$

- 8: Use these estimates to sample matrices from A to form B
-

5.3 Algorithms

In the following we formally state the two sub-procedures: Algorithm 4, which solves the required system, and Algorithm 3, which is the sampling routine for reducing the size of the system.

We prove the following theorem regarding the above algorithm REPEATED HALVING (Algorithm 3).

Theorem 23 *REPEATED HALVING (Algorithm 3) produces an $\tilde{O}(d)$ sized sample B such that $\frac{B}{2} \preceq A \preceq 2B$ and can be implemented in total time*

$$\tilde{O} \left(md + d^2 + \sqrt{\kappa_{\text{sample}}(O(d \log(d)))d} \right).$$

The \tilde{O} in the above contains factors logarithmic in $m, d, \|A\|_F$. Note that the Frobenius norm of A is bounded by $d\|A\|$ which is bounded in our case by d .

Algorithm 4 Fast Quadratic Solver (FQS)

- 1: **Input:** $A = \sum_{i=1}^m (\mathbf{v}_i \mathbf{v}_i^T + \lambda I)$, \mathbf{b} , ε
 - 2: **Output :** $\tilde{\mathbf{v}}$ s.t. $\|A^{-1}\mathbf{b} - \tilde{\mathbf{v}}\| \leq \varepsilon$
 - 3: Compute B s.t. $2B \succeq A \succeq B$ using REPEATED HALVING(Algorithm 3)
 - 4: $Q(\mathbf{y}) = \frac{\mathbf{y}^T A B^{-1} \mathbf{y}}{2} + \mathbf{b}^T \mathbf{y}$
 - 5: Compute $\hat{\mathbf{y}}$ such that $\|\hat{\mathbf{y}} - \text{argmin } Q(\mathbf{y})\| \leq \frac{\varepsilon}{4\|B^{-1}\|}$
 - 6: Output $\tilde{\mathbf{v}}$ such that $\|B^{-1}\hat{\mathbf{y}} - \tilde{\mathbf{v}}\| \leq \varepsilon/2$
-

We first provide the proof of Theorem 15 using Theorem 23 and then provide the proof for Theorem 23. For the purpose of clarity of discourse we hide the terms that appear due to the probabilistic part of the lemmas. We can take a union bound to bound the total probability of failure, and those terms show up in the log factor.

Proof [Proof of Theorem 15] We will first prove correctness which follows from noting that $\operatorname{argmin}_{\mathbf{y}} Q(\mathbf{y}) = BA^{-1}b$. Therefore we have that

$$\varepsilon/2 \geq \|B^{-1}(\hat{\mathbf{y}} - \operatorname{argmin}_{\mathbf{y}} Q(\mathbf{y}))\| = \|B^{-1}\hat{\mathbf{y}} - A^{-1}b\|.$$

A simple application of the triangle inequality proves that $\|A^{-1}b - \tilde{\mathbf{v}}\| \leq \varepsilon$.

Running Time Analysis: We now prove that the algorithm can be implemented efficiently. We first note the following two simple facts: since B is a 2-approximation of A , $\kappa(B) \leq 2\kappa(A)$; furthermore, the quadratic $Q(\mathbf{y})$ is 1-strongly convex and 2-smooth. Next, we consider the running time of implementing step 5. For this purpose we will perform gradient descent over the quadratic $Q(\mathbf{y})$. Note that the gradient of $Q(\mathbf{y})$ is given by

$$\nabla Q(\mathbf{y}) = AB^{-1}\mathbf{y} - \mathbf{b}.$$

Let $\mathbf{y}_0 = 0$, $\tilde{\varepsilon} = \left(\frac{\varepsilon}{4\|B^{-1}\|}\right)^2$ and G_Q be a bound on the gradient of the quadratic $Q(\mathbf{y})$. Let \mathbf{v}_t and \mathbf{y}_{t+1} be such that

$$\|B^{-1}\mathbf{y}_{t-1} - \mathbf{v}_t\| \leq \frac{\tilde{\varepsilon}}{100\|A\|G_Q}$$

$$\mathbf{y}_{t+1} = \mathbf{y}_t - \eta(A\mathbf{v}_t - \mathbf{b}).$$

Define $h_t = Q(\mathbf{y}_t) - \min_{\mathbf{y}} Q(\mathbf{y})$. Using the standard analysis of gradient descent, we show that the following holds true for $t > 0$:

$$h_t \leq \max\{\tilde{\varepsilon}, (0.9)^t h_0\}.$$

This follows directly from the gradient descent analysis which we outline below. To make the analysis easier, we define a *true* gradient descent series:

$$\mathbf{z}_{t+1} = \mathbf{y}_t - \eta\nabla Q(\mathbf{y}_t).$$

Note that

$$\|\mathbf{z}_{t+1} - \mathbf{y}_{t+1}\| = \|A(\mathbf{v}_t - B^{-1}\mathbf{y}_t)\| \leq \frac{\tilde{\varepsilon}}{10G_Q}.$$

We now have that

$$\begin{aligned} h_{t+1} - h_t &= Q(\mathbf{y}_{t+1}) - Q(\mathbf{y}_t) \\ &\leq \langle \nabla Q(\mathbf{y}_t), \mathbf{y}_{t+1} - \mathbf{y}_t \rangle + \frac{\beta}{2} \|\mathbf{y}_{t+1} - \mathbf{y}_t\|^2 \\ &= \langle \nabla Q(\mathbf{y}_t), \mathbf{z}_{t+1} - \mathbf{y}_t \rangle + \langle \nabla Q(\mathbf{y}_t), \mathbf{y}_{t+1} - \mathbf{z}_{t+1} \rangle + \frac{\beta}{2} \|\mathbf{z}_{t+1} - \mathbf{y}_t + \mathbf{y}_{t+1} - \mathbf{z}_{t+1}\|^2 \\ &\leq -\eta\|\nabla Q(\mathbf{y}_t)\|^2 + \langle \nabla Q(\mathbf{y}_t), \mathbf{y}_{t+1} - \mathbf{z}_{t+1} \rangle + \beta\eta^2\|\nabla Q(\mathbf{y}_t)\|^2 + \beta\|\mathbf{y}_{t+1} - \mathbf{z}_{t+1}\|^2 \\ &\leq \frac{1}{\beta}\|\nabla Q(\mathbf{y}_t)\|^2 + \frac{2}{10\beta}((\|\nabla Q(\mathbf{y}_t)\| + 1)\tilde{\varepsilon}) \\ &\leq -\frac{\alpha}{\beta}h_t + \frac{\tilde{\varepsilon}}{50\beta} \end{aligned}$$

where $\beta \geq 1$ and $\alpha \leq 2$ are the strong convexity and smoothness parameters of $Q(\mathbf{y})$. Therefore, we have that

$$h_{t+1} \leq 0.75h_t + 0.08\tilde{\varepsilon}.$$

Using the inductive assumption that $\|h_t\| \leq \max\{\tilde{\varepsilon}, (0.9)^t h_0\}$, it follows easily that

$$h_{t+1} \leq \max\{\tilde{\varepsilon}, (0.9)^{t+1} h_0\}.$$

Using the above inequality, it follows that for $t \geq O(\log(\frac{1}{\tilde{\varepsilon}}))$, we have that h_t is less than or equal to $\tilde{\varepsilon}$. By the 1-strong convexity of $Q(\mathbf{y})$, we have that if $h_t \leq \tilde{\varepsilon}$, then

$$\|\mathbf{y}_t - \underset{\mathbf{y}}{\operatorname{argmin}} Q(\mathbf{y})\| \leq \sqrt{\tilde{\varepsilon}} \leq \frac{\varepsilon}{4\|B\|^{-1}}.$$

The running time of the above sub-procedure is bounded by the time to multiply a vector with A , which takes $O(md)$ time, and the time required to compute \mathbf{v}_t , which involves solving a linear system in B at each step. Finally, in step 6 we once again compute the solution of a linear system in B . Combining these we get that the total running time is

$$\tilde{O}(md + LIN(B, \hat{\varepsilon})) \log\left(\frac{1}{\tilde{\varepsilon}}\right)$$

where $\hat{\varepsilon} = \frac{\varepsilon}{400\|A\|\|B\|^{-1}G_Q} = \Omega\left(\frac{\varepsilon}{\kappa(A)G_Q}\right)$. Now we can bound $LIN(B, \varepsilon)$ by $\tilde{O}(d^2 + d\sqrt{\kappa(A)d} \log(1/\varepsilon))$ by using Acc-SVRG to solve the linear system and by noting that B is an $O(d \log(d))$ sized 2-approximation sample of A . The bound G_Q can be taken as a bound on the gradient of the quadratic at the start of the procedure, as the norm of gradient only shrinks along the gradient descent path. It is therefore enough to take $G_Q \leq \|b\|$, which finishes the proof. \blacksquare

Proof [Proof of Theorem 23]

The correctness of the algorithm is immediate from Theorems 22 and 20. The key challenge in the proof lies in proving that there is an efficient way to implement the algorithm, which we describe next.

Runtime Analysis: To analyze the running time of the algorithm we consider the running time of the computation required in step 7 of the algorithm. We will show that there exists an efficient algorithm to compute numbers γ_i such that

$$\begin{aligned} \gamma_i &\geq \hat{\tau}_i^{A'}(A) \\ \sum \gamma_i &\leq \sum 16\hat{\tau}_i^{A'}(A) + 8k\varepsilon\|A\|_F^2. \end{aligned}$$

First note that by recursion, \tilde{A}' is a 2-spectral approximation of A' , and so we have that

$$\hat{\tau}_i^{A'}(A) \leq \hat{\tau}_i^{\tilde{A}'}(A) \leq 2\hat{\tau}_i^{A'}(A).$$

We also know that A' is an $O(d \log(d))$ sized weighted sub-sample of A , and so it is of the form

$$\tilde{A}' = \sum_{i=1}^{O(d \log(d))} \mathbf{v}_i \mathbf{v}_i^T + \lambda I = \sum_{i=1}^{O(d \log(d))} b_i b_i^T$$

where such a decomposition can be achieved easily via decomposing the identity into canonical vectors. Therefore, any such \tilde{A}' can be written as BB^T , where the columns of B are b_i .

To compute $\hat{\tau}_i^{\tilde{A}'}(A)$, we need to compute the value of $\mathbf{v}_i^T(\tilde{A}')^+\mathbf{v}_i = \|B^T(\tilde{A}')^+\mathbf{v}_i\|_2^2$. Indeed, computing this directly for all vectors v_i may be inefficient, but one can compute a constant approximation using the following procedure outlined by Cohen et al. (2015).

To compute a constant approximation, we randomly sample a Gaussian matrix G consisting of k rows and compute instead the norms $\gamma'(i) = \|GB^T(\tilde{A}')^+\mathbf{v}_i\|_2^2$. By the Johnson-Lindenstrauss lemma, setting $k = O(\log(md))$, we have that with high probability,

$$1/2\|B^T(\tilde{A}')^+\mathbf{v}_i\|_2^2 \leq \gamma'(i) \leq 2\|B^T(\tilde{A}')^+\mathbf{v}_i\|_2^2.$$

Consider the following procedure for the computation of $\gamma'(i)$:

1. Compute the matrix $G' = GB^T$ first which takes $\tilde{O}(kd^2)$ time.
2. For each row G'_i of G' compute a vector G''_i such that $\|G''_i - G'_i\tilde{A}'^+\| \leq \varepsilon$. This takes a total time of $kLIN(\tilde{A}', \varepsilon)$.
3. For each v_i compute $\gamma''_i \triangleq \sum_{j=1}^k \langle G''_j, \mathbf{v}_i \rangle^2$. This takes a total time of $O(kmd)$.

Here, $LIN(S, \varepsilon)$ is the running time of solving a linear system in S with error ε . Therefore, substituting $k = O(\log(md))$, the total running time of the above algorithm is

$$O(md + d^2 + LIN(\tilde{A}', \varepsilon)).$$

It is now easy from the definitions of γ''_i and γ'_i to see that

$$\gamma''_i \in [\gamma'_i - k\varepsilon\|\mathbf{v}_i\|^2, \gamma'_i + k\varepsilon\|\mathbf{v}_i\|^2].$$

Setting $\gamma_i = 4(\gamma''_i + k\varepsilon\|\mathbf{v}_i\|^2)$, it follows from the definitions that

$$\begin{aligned} \gamma_i &\geq \hat{\tau}_i^{A'}(A) \\ \sum \gamma_i &\leq \sum 16\hat{\tau}_i^{A'}(A) + 8k\varepsilon\|A\|_F^2. \end{aligned}$$

Now setting $\varepsilon = \frac{1}{8k\|V\|_F^2}$ satisfies the required inequality for γ_i . This implies that when sampling from γ_i , we will have a 2-approximation, and the number of matrices will be bounded by $O(d \log(d))$.

To solve the linear systems required above we use Accelerated SVRG (Theorem 4) which ensures that the running time of $LIN(S, \varepsilon) \leq \tilde{O}(d^2 + d\sqrt{\kappa(S)d} \log(\frac{1}{\varepsilon}))$. This follows by noting that in our case S is a sum of $O(d \log(d))$ matrices. To bound the condition number of S , note that it is a 2-approximation of some unweighted sample of the matrix A of size greater than $O(d \log(d))$. Therefore, we have that $\kappa(S) \leq \kappa_{sample}(A)$.

Putting the above arguments together we get that the total running time of the procedure is bounded by

$$\tilde{O}(md + (d^2 + d\sqrt{\kappa_{sample}(A)d})),$$

which proves the theorem. ■

6. Condition Number Independent Algorithms

In this section we state our main result regarding self-concordant functions and provide the proof. The following is the definition of a self-concordant function:

Definition 24 (Self-Concordant Functions) *Let $\mathcal{K} \subseteq \mathbb{R}^n$ be a non-empty open convex set, and let $f : \mathcal{K} \mapsto \mathbb{R}$ be a C^3 convex function. Then, f is said to be self-concordant if*

$$|\nabla^3 f(\mathbf{x})[\mathbf{h}, \mathbf{h}, \mathbf{h}]| \leq 2(\mathbf{h}^\top \nabla^2 f(\mathbf{x}) \mathbf{h})^{3/2},$$

where we have

$$\nabla^k f(\mathbf{x})[\mathbf{h}_1, \dots, \mathbf{h}_k] \triangleq \frac{\partial^k}{\partial t_1 \dots \partial t_k} \Big|_{t_1=\dots=t_k} f(\mathbf{x} + t_1 \mathbf{h}_1 + \dots + t_k \mathbf{h}_k).$$

Our main theorem regarding self-concordant functions is as follows:

Theorem 25 *Let $0 < r < 1$, let $\gamma \geq 1$, and let ν be a constant depending on γ, r . Set $\eta = 10(1-r)^2$, $S_1 = c_r = \frac{50}{(1-r)^4}$, where c_r is a constant depending on r , and $T = \frac{f(\mathbf{x}_1) - f(\mathbf{x}^*)}{\nu}$. Then, after $t > T$ steps, the following linear convergence guarantee holds between two full gradient steps for Algorithm 5:*

$$\mathbf{E}[f(w_s) - f(w^*)] \leq \frac{1}{2} \mathbf{E}[f(w_{s-1}) - f(w^*)].$$

Moreover, the time complexity of each full gradient step is $O(md + c_r d^2)$, where c_r is a constant depending on r (independent of the condition number).

To describe the algorithm we first present the requisite definitions regarding self-concordant functions.

6.1 Self-Concordant Functions Preliminaries

An excellent reference for this material is the lecture notes on this subject by Nemirovski (2004).

A key object in the analysis of self-concordant functions is the notion of a Dikin ellipsoid, which is the unit ball around a point \mathbf{x} in the norm given by the Hessian $\|\cdot\|_{\nabla^2 f(\mathbf{x})}$ at the point. We will refer to this norm as the *local norm* around a point and denote it as $\|\cdot\|_{\mathbf{x}}$. Formally, we have:

Definition 26 (Dikin ellipsoid) *The Dikin ellipsoid of radius r centered at a point \mathbf{x} is defined as*

$$W_r(\mathbf{x}) \triangleq \{\mathbf{y} \mid \|\mathbf{y} - \mathbf{x}\|_{\nabla^2 f(\mathbf{x})} \leq r\}.$$

One of the key properties of self-concordant functions is that inside the Dikin ellipsoid, the function is well conditioned with respect to the local norm at the center, and furthermore, the function is smooth. The following lemmas makes these notions formal, and the proofs of these lemmas can be found in the lecture notes of Nemirovski (2004).

Lemma 27 (Nemirovski 2004) For all \mathbf{h} such that $\|\mathbf{h}\|_{\mathbf{x}} < 1$ we have that

$$(1 - \|\mathbf{h}\|_{\mathbf{x}})^2 \nabla^2 f(\mathbf{x}) \preceq \nabla^2 f(\mathbf{x} + \mathbf{h}) \preceq \frac{1}{(1 - \|\mathbf{h}\|_{\mathbf{x}})^2} \nabla^2 f(\mathbf{x}).$$

Lemma 28 (Nemirovski 2004) For all \mathbf{h} such that $\|\mathbf{h}\|_{\mathbf{x}} < 1$ we have that

$$f(\mathbf{x} + \mathbf{h}) \leq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{h} \rangle + \rho(\|\mathbf{h}\|_{\mathbf{x}})$$

where $\rho(x) \triangleq -\ln(1 - x) - x$.

Another key quantity which is used both as a potential function as well as a dampening for the step size in the analysis of Newton's method in general is the Newton decrement which is defined as $\lambda(\mathbf{x}) \triangleq \|\nabla f(\mathbf{x})\|_{\mathbf{x}}^* = \sqrt{\nabla f(\mathbf{x})^\top \nabla^{-2} f(\mathbf{x}) \nabla f(\mathbf{x})}$. The following lemma quantifies how $\lambda(\mathbf{x})$ behaves as a potential by showing that once it drops below 1, it ensures that the minimum of the function lies in the current Dikin ellipsoid. This is the property which we use crucially in our analysis.

Lemma 29 (Nemirovski 2004) If $\lambda(\mathbf{x}) < 1$ then

$$\|\mathbf{x} - \mathbf{x}^*\|_{\mathbf{x}} \leq \frac{\lambda(\mathbf{x})}{1 - \lambda(\mathbf{x})}.$$

6.2 Condition Number Independent Algorithms

In this section we describe an efficient linearly convergent method (Algorithm 5) for optimization of self-concordant functions for which the running time is independent of the condition number. We have not tried to optimize the complexity of the algorithms in terms of d as our main focus is to make it condition number independent.

The key idea here is the ellipsoidal view of Newton's method, whereby we show that after making a constant number of full Newton steps, one can identify an ellipsoid and a norm such that the function is well conditioned with respect to the norm in the ellipsoid. This is depicted in Figure 1.

At this point one can run any desired first-order algorithm. In particular, we choose SVRG and prove its fast convergence. Algorithm 6 (described in the appendix) states the modified SVRG routine for general norms used in Algorithm 5.

We now state and prove the following theorem regarding the convergence of Algorithm 5.

Proof It follows from Lemma 31 that at $t = \frac{f(\mathbf{x}_1) - f(\mathbf{x}^*)}{\nu}$, the minimizer is contained in the Dikin ellipsoid of radius r , where ν is a constant depending on γ, r . This fact, coupled with Lemma 30, shows that the function satisfies the following property with respect to $W_{\mathbf{x}_t}$:

$$\forall \mathbf{x} \in W_{\mathbf{x}_t} \quad (1 - r)^2 \nabla^2 f(\mathbf{x}_T) \preceq \nabla^2 f(\mathbf{x}) \preceq (1 - r)^{-2} \nabla^2 f(\mathbf{x}_T).$$

Using the above fact along with Lemma 32, and substituting for the parameters, concludes the proof. ■

We describe the Algorithm N-SVRG (Algorithm 6 in the appendix). Since the algorithm and the following lemmas are minor variations of their original versions, we include the proofs of the following lemmas in the appendix for completeness.




Figure 1: Visualization of the steps taken in Algorithm 5.

Lemma 30 *Let f be a self-concordant function over \mathcal{K} , let $0 < r < 1$, and consider $\mathbf{x} \in \mathcal{K}$. Let $W_r(\mathbf{x})$ be the Dikin ellipsoid of radius r , and let $\alpha = (1 - r)^2$ and $\beta = (1 - r)^{-2}$. Then, for all \mathbf{h} s.t. $\mathbf{x} + \mathbf{h} \in W_r(\mathbf{x})$,*

$$\alpha \nabla^2 f(\mathbf{x}) \preceq \nabla^2 f(\mathbf{x} + \mathbf{h}) \preceq \beta \nabla^2 f(\mathbf{x}).$$

Lemma 31 *Let f be a self-concordant function over \mathcal{K} , let $0 < r < 1$, let $\gamma \geq 1$, and consider following the damped Newton step as described in Algorithm 5 with initial point \mathbf{x}_1 . Then, the number of steps t of the algorithm before the minimizer of f is contained in the Dikin ellipsoid of radius r of the current iterate, i.e. $\mathbf{x}^* \in W_r(\mathbf{x}_t)$, is at most $t = \frac{f(\mathbf{x}_1) - f(\mathbf{x}^*)}{\nu}$, where ν is a constant depending on γ, r .*

Algorithm 5 ellipsoidCover

- 1: **Input** : Self-concordant $f(\mathbf{x}) = \sum_{i=1}^m f_i(\mathbf{x})$, $T, r \in \{0, 1\}$, initial point $\mathbf{x}_1 \in \mathcal{K}$, $\gamma > 1, S_1, \eta, T$
 - 2: Initialize: $\mathbf{x}_{\text{curr}} = \mathbf{x}_1$
 - 3: **while** $\lambda(\mathbf{x}_{\text{curr}}) > \frac{r}{1+r}$ **do**
 - 4: $step = \gamma(1 + \lambda(\mathbf{x}_{\text{curr}}))$
 - 5: $\mathbf{x}_{\text{curr}} = \mathbf{x}_{\text{curr}} - \frac{1}{step} \nabla^{-2} f(\mathbf{x}_{\text{curr}}) \nabla f(\mathbf{x}_{\text{curr}})$
 - 6: **end while**
 - 7: $\mathbf{x}_{T+1} = \mathbf{N} - \mathbf{SVRG}(W_r(\mathbf{x}), f(\mathbf{x}), \nabla^2 f(\mathbf{x}_{\text{curr}}), S_1, \eta, T)$
 - 8: **return** \mathbf{x}_{T+1}
-

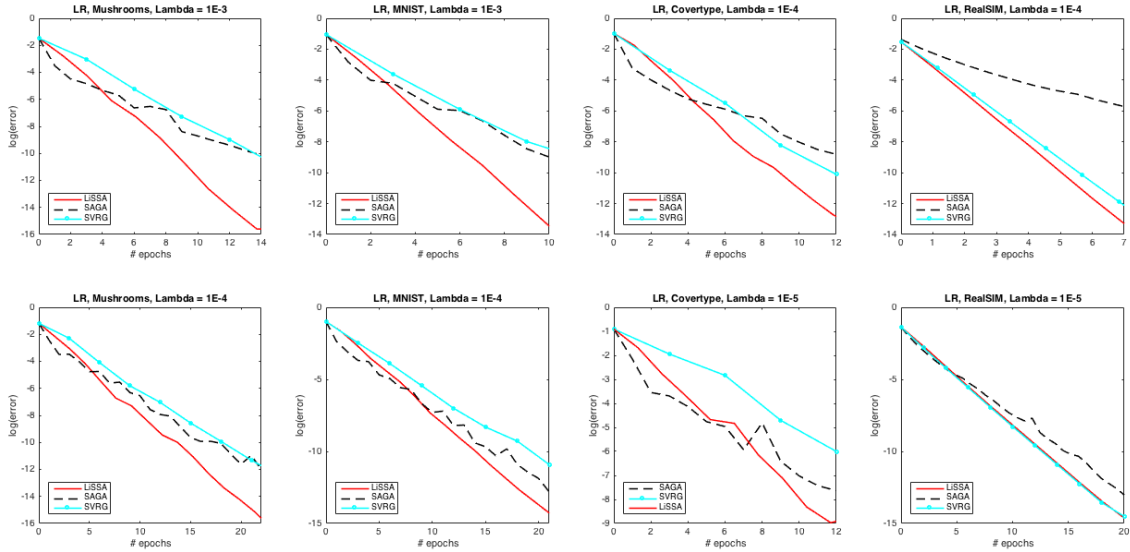


Figure 2: Performance of LiSSA as compared to a variety of related optimization methods for different data sets and choices of regularization parameter λ . $S_1 = 1$, $S_2 \sim \kappa \ln(\kappa)$.

Lemma 32 *Let f be a convex function. Suppose there exists a convex set \mathcal{K} and a positive semidefinite matrix A such that for all $\mathbf{x} \in \mathcal{K}$, $\alpha A \preceq \nabla^2 f(\mathbf{x}) \preceq \beta A$. Then the following holds between two full gradient steps of Algorithm 6:*

$$\mathbf{E}[f(\mathbf{x}_s) - f(\mathbf{x}^*)] \leq \left(\frac{1}{\alpha\eta(1 - 2\eta\beta)n} + \frac{2\eta\beta}{(1 - 2\eta\beta)} \right) \mathbf{E}[f(\mathbf{x}_{s-1}) - f(\mathbf{x}^*)].$$

7. Experiments

In this section we present experimental evaluation for our theoretical results.⁷ We perform the experiments for a classification task over two labels using the logistic regression (LR) objective function with the ℓ_2 regularizer. For all of the classification tasks we choose two values of λ : $\frac{1}{m}$ and $\frac{10}{m}$, where m is the number of training examples. We perform the above classification tasks over four data sets: MNIST, CoverType, Mushrooms, and RealSIM. Figure 2 displays the log-error achieved by LiSSA as compared to two standard first-order algorithms, SVRG and SAGA (Johnson and Zhang, 2013; Defazio et al., 2014), in terms of the number of passes over the data. Figure 3 presents the performance of LiSSA as compared to NewSamp (Erdogdu and Montanari, 2015) and standard Newton’s method with respect to both time and iterations.

7. Our code for LiSSA can be found here: https://github.com/brianbullins/lissa_code.




Figure 3: Convergence of LiSSA over time/iterations for logistic regression with MNIST, as compared to NewSamp and Newton’s method.

7.1 Experiment Details

In this section we describe our experiments and choice of parameters in detail. Table 2 provides details regarding the data sets chosen for the experiments. To make sure our functions are scaled such that the norm of the Hessian is bounded, we scale the above data set points to unit norm.

Table 2: A description of data sets used in the experiments.

DATA SET	M	D	REFERENCE
MNIST4-9	11791	784	LECUN AND CORTES (1998)
MUSHROOMS	8124	112	LICHMAN (2013)
COVERTYPE	100000	54	BLACKARD AND DEAN (1999)
REALSIM	72309	20958	MCCALLUM (1997)

7.2 Comparison with Standard Algorithms

In Figures 2 and 4 we present comparisons between the efficiency of our algorithm with different standard and popular algorithms. In both cases we plot $\log(\text{CurrentValue} - \text{OptimumValue})$. We obtained the optimum value for each case by running our algorithm for a long enough time until it converged to the point of machine precision.

Epoch Comparison: In Figure 2, we compare LiSSA with SVRG and SAGA in terms of the accuracy achieved versus the number of passes over the data. To compute the number of passes in SVRG and SAGA, we make sure that the inner stochastic gradient iteration in both the algorithms counts as exactly one pass. This is done because although it accesses gradients at two different points, one of them can be stored from before in both cases. The outer full gradient in SVRG counts as one complete pass over the data. We set the number of inner iterations of SVRG to $2m$ for the case when $\lambda = 1/m$, and we parameter tune the number of inner iterations when $\lambda = 10/m$. The stepsizes for all of the algorithms are parameter tuned by an exhaustive search over the parameters.

Time Comparison: For the comparison with respect to time (Figure 4), we consider the following algorithms: AdaGrad (Duchi et al., 2011), BFGS (Broyden, 1970; Fletcher, 1970;

Goldfarb, 1970; Shanno, 1970), gradient descent, SGD, SVRG (Johnson and Zhang, 2013) and SAGA (Defazio et al., 2014). The $\log(\text{Error})$ is plotted as a function of the time elapsed from the start of the run for each algorithm. We next describe our choice of parameters for the algorithms. For AdaGrad we used the faster diagonal scaling version proposed by Duchi et al. (2011). We implemented the basic version of BFGS with backtracking line search. In each experiment for gradient descent, we find a reasonable step size using parameter tuning. For stochastic gradient descent, we use the variable step size $\eta_t = \gamma/\sqrt{t}$ which is usually the prescribed step size, and we hand tune the parameter γ . The parameters for SVRG and SAGA were chosen in the same way as before.

Choice of Parameters for LiSSA: To pick the parameters for our algorithm, we observe that it exhibits smooth behavior even in the case of $S_1 = 1$, so this is used for the experiments. However, we observe that increasing S_2 has a positive effect on the convergence of the algorithm up to a certain point, as a higher S_2 leads to a larger per-iteration cost. This behavior is consistent with the theoretical analysis. We summarize the comparison between the per-iteration convergence and the value of S_2 in Figure 5. As the theory predicts S_2 to be of the order $\kappa \ln(\kappa)$, for our experiments we determine an estimate for κ and set S_2 to around $\kappa \ln(\kappa)$. This value is typically equal to m in our experiments. We observe that setting S_2 in this way resulted in the experimental results displayed in Figure 2.

Comparison with Second-Order Methods: Here we present details about the comparison between LiSSA, NewSamp (Erdogdu and Montanari, 2015), and standard Newton’s method, as displayed in Figure 3. We perform this experiment on the MNIST data set and show the convergence properties of all three algorithms over time as well as over iterations. We could not replicate the results of NewSamp on all of our data sets as it sometimes seems to diverge in our experiments. For logistic regression on the MNIST data set, we could get it to converge by setting the value of S_1 to be slightly higher. We observe as is predicted by the theory that when compared in terms of the number of iterations, NewSamp and LiSSA perform similarly, while Newton’s method performs the best as it attains a quadratic convergence rate. This can be seen in Figure 3. However, when we consider the performance in terms of time for these algorithms, we see that LiSSA has a significant advantage.

Comparison with Accelerated First-Order Methods: Here we present experimental results comparing LiSSA with a popular accelerated first-order method, APCG (Lin et al., 2014), as seen in Figure 6. We ran the experiment on the RealSIM data set with three settings of $\lambda = 10^{-5}, 10^{-6}, 10^{-7}$, to account for the high condition number setting. We observe a trend that can be expected from the runtime guarantees of the algorithms. When λ is not too low, LiSSA performs better than APCG, but as λ gets very low we see that APCG performs significantly better than LiSSA. This is not surprising when considering that the running time of APCG grows proportional to $\sqrt{\kappa m}$, whereas for LiSSA the running time can at best be proportional to κ . We note that for accelerated first-order methods to be useful, one needs the condition number to be quite large which is not often the case for applications. Nevertheless, we believe that an algorithm with running time guarantees similar to LiSSA-Sample can get significant gains in these settings, and we leave this exploration as future work.

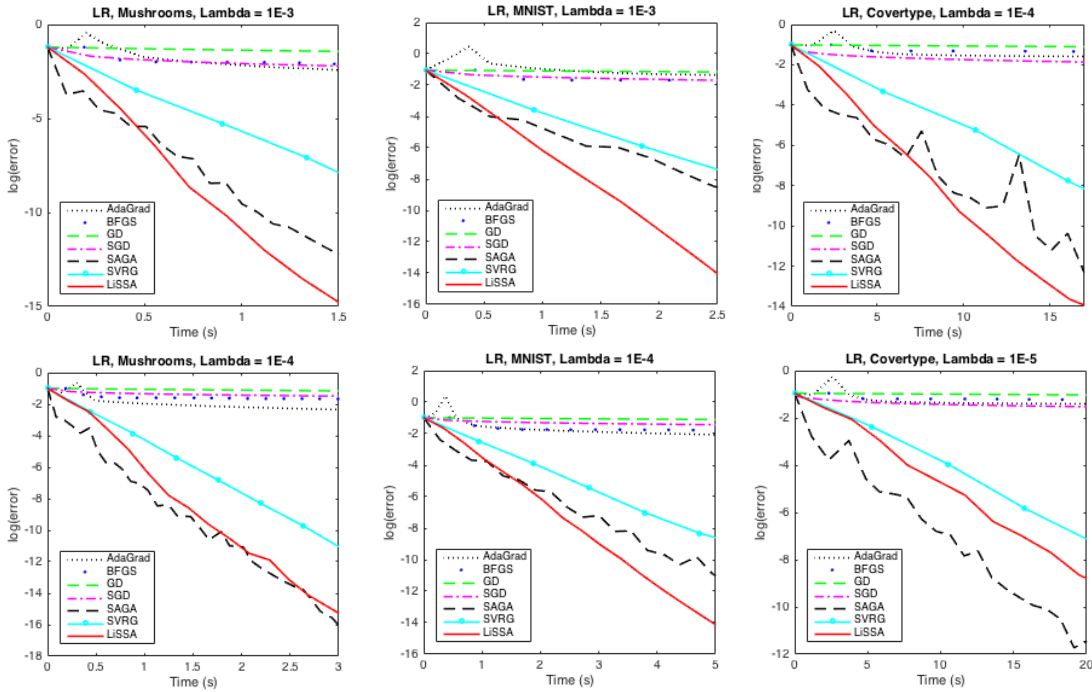


Figure 4: Performance of LiSSA as compared to a variety of related optimization methods for different data sets and choices of regularization parameter λ . $S_1 = 1$, $S_2 \sim \kappa \ln(\kappa)$.




Figure 5: Differing convergence rates for LiSSA based on different choices of the S_2 parameter.




Figure 6: Comparison with accelerated first-order methods in ill-conditioned settings.

Acknowledgments

The authors would like to thank Zeyuan Allen-Zhu, Dan Garber, Haipeng Luo and David McAllester for several illuminating discussions and suggestions. We would especially like to thank Aaron Sidford and Yin Tat Lee for discussions regarding the matrix sampling approach.

Appendix A. Remaining Proofs

Here we provide proofs for the remaining lemmas.

A.1 Proof of Lemma 20

Proof The proof is almost identical to the proof of Lemma 4 by Cohen et al. (2015). As in there we will use the following lemma on matrix concentration, which appears as Lemma 11 in the work of Cohen et al. (2015).

Lemma 33 (Cohen et al. 2015) *Let $Y_1 \dots Y_k$ be independent random positive semidefinite matrices of size $d \times d$. Let $Y = \sum Y_i$ and let $Z = E[Y]$. If $Y_i \preceq R \cdot Z$, then*

$$P \left[\sum Y_i \preceq (1 - \varepsilon)Z \right] \leq de^{\frac{\varepsilon^2}{2R}}$$

$$P \left[\sum Y_i \succeq (1 + \varepsilon)Z \right] \leq de^{\frac{\varepsilon^2}{3R}}.$$

For every matrix A_i choose $Y_i = \frac{A_i}{p_i}$ with probability p_i and 0 otherwise. Therefore we need to bound $\sum Y_i$. Note that $E[Y_i] = A$. We will now show that

$$\forall i \ Y_i \preceq \frac{1}{c \log d \varepsilon^{-2}} A$$

which will finish the proof by a direct application of Lemma 33. First we will show that

$$\frac{A_i}{\tau_i(A)} \preceq A.$$

We need to show that $\forall \mathbf{x} \ \mathbf{x}^T A_i \mathbf{x} \leq \tau_i(A) \mathbf{x}^T A \mathbf{x}$. By noting that the A_i are PSD we have that if $\mathbf{x}^T A \mathbf{x} = 0$ then $\forall i, \mathbf{x}^T A_i \mathbf{x} = 0$. We can now consider $\mathbf{x} = A^{+1/2} \mathbf{y}$. Therefore, we need to show that

$$\mathbf{y} A^{+1/2} A_i A^{+1/2} \mathbf{y} \leq \tau_i(A) \|\mathbf{y}\|^2.$$

This is true because the maximum eigenvalue of $A^{+1/2} A_i A^{+1/2}$ is bounded by its $\text{Trace}(A^{+1/2} A_i A^{+1/2})$ (due to positive semidefiniteness), which is equal to $\tau_i(A)$ by definition. Therefore when $p_i < 1$, i.e., $\alpha c \mathbf{u}_i \log(d) < 1$, the facts $\tau_i(A) \leq \mathbf{u}_i$ immediately provide

$$Y_i \preceq \frac{1}{c \log d \varepsilon^{-2}} A.$$

When $p_i = 1$ the above does not hold, but we can essentially replace $Y_i = A_i$ with $c \log d \varepsilon^{-2}$ variables each equal to $\frac{A_i}{c \log d \varepsilon^{-2}}$, each being sampled with probability 1. This does not change $E[\sum Y_i]$ but proves concentration. Now a direct application of Lemma 33 finishes the proof. Also note that a standard Chernoff bound proves the required bound on the sample size. ■

A.2 Proof of Lemma 21

Proof The proof is almost identical to the proof of Theorem 1 by Cohen et al. (2015). We include the proof here for completeness. Define S_i to be S if $\exists j$ such that $X_j = A_i$, and $S + A_i$ otherwise. Now since $S^i \preceq A$, we have that $\tilde{\tau}_i^S(A) \geq \tau_i(A)$. To bound the expected sum we break it down into two parts:

$$\sum_i \tilde{\tau}_i^S(A) = \sum_{i \in S} \tilde{\tau}_i^S(A) + \sum_{i \notin S} \tilde{\tau}_i^S(A).$$

The first part is a sum of the leverage scores of S , and so it is bounded by d . To bound the second term, consider a random process that first selects S , then selects a random $i \notin S$ and returns $\tilde{\tau}_i^S(A)$. There are always exactly $m - r$ such i , so the value returned by this random process is, in expectation, exactly equal to $\frac{1}{m-r} E[\sum_{i \notin S} \tilde{\tau}_i^S(A)]$. This random process is also equivalent to randomly selecting a set S' of size $r + 1$, then randomly choosing an $i \in S'$ and returning its leverage score. In expectation it is therefore equal to the average leverage score in S' . S' has size $r + 1$, and its leverage scores sum to its rank, so we can bound its average leverage score by $\frac{d}{r+1}$. Overall, we have that

$$E[\sum_i \tilde{\tau}_i^S(A)] = d + \frac{d(m-r)}{r+1} \leq O\left(\frac{md}{r}\right).$$

■

A.3 Proof of Ellipsoidal Cover Lemma

Proof [Proof of Lemma 31] Let $\lambda(\mathbf{x})$ be the Newton decrement at \mathbf{x} . By Lemma 29, we know that if $\lambda(\mathbf{x}) \leq \frac{r}{1+r} < 1$, then

$$\|\mathbf{x} - \mathbf{x}^*\|_{\mathbf{x}} \leq \frac{\lambda(\mathbf{x})}{1 - \lambda(\mathbf{x})} = r.$$

Consider a single iteration of the algorithm at time t . If $\lambda(\mathbf{x}_t) \leq \frac{r}{1+r}$, then we may conclude that $\mathbf{x}^* \in W_r(\mathbf{x}_t)$. Therefore, it is only when $\lambda(\mathbf{x}_t) > \frac{r}{1+r}$ that \mathbf{x}^* may not be contained within $W_r(\mathbf{x}_t)$. Since our update is of the form

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \frac{1}{\gamma(1 + \lambda(\mathbf{x}_t))} \nabla^{-2} f(\mathbf{x}_t) \nabla f(\mathbf{x}_t),$$

we have the following:

$$\begin{aligned} f(\mathbf{x}_{t+1}) &\leq f(\mathbf{x}_t) - \frac{1}{\gamma(1 + \lambda(\mathbf{x}_t))} \nabla f(\mathbf{x}_t)^\top ([\nabla^2 f(\mathbf{x}_t)]^{-1} \nabla f(\mathbf{x}_t)) + \rho \left(\frac{\lambda(\mathbf{x}_t)}{\gamma(1 + \lambda(\mathbf{x}_t))} \right) \\ &= f(\mathbf{x}_t) - \frac{(\lambda(\mathbf{x}_t))^2}{\gamma(1 + \lambda(\mathbf{x}_t))} + \rho \left(\frac{\lambda(\mathbf{x}_t)}{\gamma(1 + \lambda(\mathbf{x}_t))} \right) \\ &= f(\mathbf{x}_t) - \frac{(\lambda(\mathbf{x}_t))^2}{\gamma(1 + \lambda(\mathbf{x}_t))} - \ln \left(1 - \frac{\lambda(\mathbf{x}_t)}{\gamma(1 + \lambda(\mathbf{x}_t))} \right) - \frac{\lambda(\mathbf{x}_t)}{\gamma(1 + \lambda(\mathbf{x}_t))} \\ &= f(\mathbf{x}_t) - \frac{\lambda(\mathbf{x}_t)}{\gamma} + \ln \left(1 + \frac{\lambda(\mathbf{x}_t)}{\gamma + (\gamma - 1)\lambda(\mathbf{x}_t)} \right) \end{aligned}$$

where the first inequality follows from Lemma 28. It now follows that

$$\begin{aligned}
 f(\mathbf{x}_t) - f(\mathbf{x}_{t+1}) &\geq \frac{\lambda(\mathbf{x}_t)}{\gamma} - \ln \left(1 + \frac{\lambda(\mathbf{x}_t)}{\gamma + (\gamma - 1)\lambda(\mathbf{x}_t)} \right) \\
 &\geq \frac{\lambda(\mathbf{x}_t)}{\gamma} - \frac{\lambda(\mathbf{x}_t)}{\gamma + (\gamma - 1)\lambda(\mathbf{x}_t)} \\
 &= \lambda(\mathbf{x}_t) \left(\frac{1}{\gamma} - \frac{1}{\gamma + (\gamma - 1)\lambda(\mathbf{x}_t)} \right) \\
 &> \frac{r}{1+r} \left(\frac{1}{\gamma} - \frac{1}{\gamma + \frac{(\gamma-1)r}{1+r}} \right) \\
 &> 0
 \end{aligned}$$

where the second inequality comes from the fact that, for all $x \in \mathbb{R}$, $\ln(1+x) \leq 1+x$. Let $\nu = \frac{r}{1+r} \left(\frac{1}{\gamma} - \frac{1}{\gamma + \frac{(\gamma-1)r}{1+r}} \right)$. Then we see that after

$$\frac{f(\mathbf{x}_1) - f(\mathbf{x}^*)}{\nu}$$

steps, we can guarantee that we have arrived at \mathbf{x}_t such that $\mathbf{x}^* \in W_r(\mathbf{x}_t)$. ■

Algorithm 6 N-SVRG

Input: Convex set \mathcal{K} , $f(\mathbf{x}) = \frac{1}{m} \sum_{k=1}^m f_k(\mathbf{x})$, Norm $A \succeq 0$, update frequency S_1 , and learning rate η
 Initialize $\tilde{\mathbf{x}}_0$
for $s = 1, 2 \dots$ **do**
 $\tilde{\mathbf{x}} = \tilde{\mathbf{x}}_{s-1}$
 $\tilde{\mu} = \frac{1}{m} \sum_{k=1}^m \nabla f_k(\tilde{\mathbf{x}})$
 $\mathbf{x}_0 = \tilde{\mathbf{x}}$
 for $t = 1$ to S_1 **do**
 Randomly choose $i_t \in \{1 \dots m\}$
 $g = \nabla f_{i_t}(\mathbf{x}_{t-1}) - \nabla f_{i_t}(\tilde{\mathbf{x}}) + \tilde{\mu}$
 $\mathbf{x}_t = \Pi_{\mathcal{K}}^A(\mathbf{x}_{t-1} - \eta A^{-1}g)$
 end for
 $\tilde{\mathbf{x}}_s = \mathbf{x}_t$ for randomly chosen $t \in \{1, \dots, m\}$
end for

A.4 Proof of SVRG Lemma

Proof [Proof of Lemma 32] The proof follows the original proof of SVRG (Johnson and Zhang, 2013) with a few modifications to take the general norm into account. For any i , consider

$$g_i(\mathbf{x}) = f_i(\mathbf{x}) - f_i(\mathbf{x}^*) - \langle \nabla f_i(\mathbf{x}^*), \mathbf{x} - \mathbf{x}^* \rangle.$$

We know that $g_i(\mathbf{x}^*) = \min_{\mathbf{x}} g_i(\mathbf{x})$ since $\nabla g_i(\mathbf{x}^*) = 0$. Therefore,

$$\begin{aligned} 0 = g_i(\mathbf{x}^*) &\leq \min_{\eta} (g_i(\mathbf{x} - \eta A^{-1} \nabla g_i(\mathbf{x}^*))) \\ &\leq \min_{\eta} (g_i(\mathbf{x}) - \eta \|g_i(\mathbf{x})\|_{A^{-1}}^2 + 0.5\beta\eta^2 \|g_i(\mathbf{x})\|_{A^{-1}}^2) \\ &= g_i(\mathbf{x}) - \frac{1}{2\beta} \|g_i(\mathbf{x})\|_{A^{-1}}^2 . \end{aligned}$$

The second inequality follows from smoothness of g_i , where we note that g_i is as smooth as f_i . Summing the above inequality over i and setting \mathbf{x}^* to be the minimum of f so that $\nabla f(\mathbf{x}^*) = 0$, we get

$$\frac{\sum_{i=1}^m \|\nabla f_i(\mathbf{x}) - \nabla f_i(\mathbf{x}^*)\|_{A^{-1}}^2}{2m\beta} \leq f(\mathbf{x}) - f(\mathbf{x}^*).$$

Define $v_t = \nabla f_{i_t}(\mathbf{x}_{t-1}) - \nabla f_{i_t}(\tilde{\mathbf{x}}) + \tilde{\mu}$. Conditioned on \mathbf{x}_{t-1} , we can take the expectation with respect to i_t and obtain

$$\begin{aligned} \mathbf{E} [\|v_t\|_{A^{-1}}^2] &\leq 2\mathbf{E} [\|\nabla f_{i_t}(\mathbf{x}_{t-1}) - \nabla f_{i_t}(\mathbf{x}^*)\|_{A^{-1}}^2] + 2\mathbf{E} [\|\nabla f_{i_t}(\tilde{\mathbf{x}}) - \nabla f_{i_t}(\mathbf{x}^*) - \nabla f(\tilde{\mathbf{x}})\|_{A^{-1}}^2] \\ &\leq 2\mathbf{E} [\|\nabla f_{i_t}(\mathbf{x}_{t-1}) - \nabla f_{i_t}(\mathbf{x}^*)\|_{A^{-1}}^2] + 2\mathbf{E} [\|\nabla f_{i_t}(\tilde{\mathbf{x}}) - \nabla f_{i_t}(\mathbf{x}^*)\|_{A^{-1}}^2] \\ &\leq 4\beta (f(\mathbf{x}_{t-1}) - f(\mathbf{x}^*) + f(\tilde{\mathbf{x}}) - f(\mathbf{x}^*)). \end{aligned}$$

The above inequalities follow by noting the following three facts:

- $\|a + b\|_{A^{-1}}^2 \leq 2\|a\|_{A^{-1}}^2 + 2\|b\|_{A^{-1}}^2$.
- $\tilde{\mu} = \nabla f(\tilde{\mathbf{x}})$.
- $\mathbf{E} [\|X - \mathbf{E}X\|_{A^{-1}}^2] \leq \mathbf{E}\|X\|_{A^{-1}}^2$.

Now note that conditioned on \mathbf{x}_{t-1} , we have that $\mathbf{E}v_t = \nabla f(\mathbf{x}_{t-1})$, and so

$$\begin{aligned} \mathbf{E}\|\mathbf{x}_t - \mathbf{x}^*\|_A^2 &\leq \|\mathbf{x}_{t-1} - \eta A^{-1}v_t - \mathbf{x}^*\|_A^2 \\ &= \|\mathbf{x}_{t-1} - \mathbf{x}^*\|_A^2 - 2\eta \langle \mathbf{x}_{t-1} - \mathbf{x}^*, AA^{-1}\mathbf{E}v_t \rangle + \eta^2 \mathbf{E}\|v_t\|_{A^{-1}}^2 \\ &\leq \|\mathbf{x}_{t-1} - \mathbf{x}^*\|_A^2 - 2\eta \langle \mathbf{x}_{t-1} - \mathbf{x}^*, \nabla f(\mathbf{x}_{t-1}) \rangle + 4\beta\eta^2 (f(\mathbf{x}_{t-1}) - f(\mathbf{x}^*) + f(\tilde{\mathbf{x}}) - f(\mathbf{x}^*)) \\ &\leq \|\mathbf{x}_{t-1} - \mathbf{x}^*\|_A^2 - 2\eta (f(\mathbf{x}_{t-1}) - f(\mathbf{x}^*)) + 4\beta\eta^2 (f(\mathbf{x}_{t-1}) - f(\mathbf{x}^*) + f(\tilde{\mathbf{x}}) - f(\mathbf{x}^*)) \\ &= \|\mathbf{x}_{t-1} - \mathbf{x}^*\|_A^2 - 2\eta(1 - 2\eta\beta) (f(\mathbf{x}_{t-1}) - f(\mathbf{x}^*)) + 4\beta\eta^2 (f(\tilde{\mathbf{x}}) - f(\mathbf{x}^*)). \end{aligned}$$

The first inequality uses the Pythagorean inequality for norms, the second inequality uses the derived inequality for $\mathbf{E}\|v_t\|_{A^{-1}}^2$, and the third inequality uses the convexity of $f(\mathbf{x})$. Consider a fixed stage s , so that $\tilde{\mathbf{x}} = \tilde{\mathbf{x}}_{s-1}$ and $\tilde{\mathbf{x}}_s$ are selected after all of the updates have been completed. By summing the previous inequality over t and taking an expectation over the history, we get

$$\begin{aligned} \mathbf{E}\|\mathbf{x}_n - \mathbf{x}^*\|_A^2 + 2\eta(1 - 2\eta\beta)n\mathbf{E}[f(\mathbf{x}_s) - f(\mathbf{x}^*)] &\leq \mathbf{E}\|\mathbf{x}_0 - \mathbf{x}^*\|_A^2 + 4\beta n\eta^2 \mathbf{E}[f(\tilde{\mathbf{x}}) - f(\mathbf{x}^*)] \\ &= \|\tilde{\mathbf{x}} - \mathbf{x}^*\|_A^2 + 4\beta n\eta^2 \mathbf{E}[f(\tilde{\mathbf{x}}) - f(\mathbf{x}^*)] \\ &\leq \frac{\mathbf{E}[f(\tilde{\mathbf{x}}) - f(\mathbf{x}^*)]}{\alpha} + 4\beta n\eta^2 \mathbf{E}[f(\tilde{\mathbf{x}}) - f(\mathbf{x}^*)] \\ &= 2 \left(\frac{1}{\alpha} + 2\beta n\eta^2 \right) \mathbf{E}[f(\tilde{\mathbf{x}}) - f(\mathbf{x}^*)]. \end{aligned}$$

The second inequality uses the strong convexity property. Therefore, we have that

$$\mathbf{E}[f(\mathbf{x}_s) - f(\mathbf{x}^*)] \leq \left(\frac{1}{\alpha\eta(1 - 2\eta\beta)n} + \frac{2\eta\beta}{(1 - 2\eta\beta)} \right) \mathbf{E}[f(\mathbf{x}_{s-1}) - f(\mathbf{x}^*)].$$

■

References

- Naman Agarwal, Zeyuan Allen-Zhu, Brian Bullins, Elad Hazan, and Tengyu Ma. Finding approximate local minima faster than gradient descent. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1195–1199. ACM, 2017.
- Zeyuan Allen-Zhu. Katyusha: The first direct acceleration of stochastic gradient methods. *arXiv preprint arXiv:1603.05953*, 2016.
- Yossi Arjevani and Ohad Shamir. Oracle complexity of second-order methods for finite-sum problems. *arXiv preprint arXiv:1611.04982*, 2016.
- Jock A. Blackard and Denis J. Dean. Comparative accuracies of artificial neural networks and discriminant analysis in predicting forest cover types from cartographic variables. *Computers and Electronics in Agriculture*, 24(3):131–151, 1999.
- Raghu Bollapragada, Richard Byrd, and Jorge Nocedal. Exact and inexact subsampled Newton methods for optimization. *arXiv preprint arXiv:1609.08502*, 2016.
- Charles G. Broyden. The convergence of a class of double-rank minimization algorithms: 2. the new algorithm. *IMA Journal of Applied Mathematics*, 6(3):222–231, 1970.
- 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.
- Michael B. Cohen. Nearly tight oblivious subspace embeddings by trace inequalities. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 278–287. Society for Industrial and Applied Mathematics, 2016.
- Michael B. Cohen, Yin Tat Lee, Cameron Musco, Christopher Musco, Richard Peng, and Aaron Sidford. Uniform sampling for matrix approximation. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science*, pages 181–190. ACM, 2015.
- 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.

- John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul): 2121–2159, 2011.
- Murat A. Erdogdu and Andrea Montanari. Convergence rates of sub-sampled Newton methods. In *Advances in Neural Information Processing Systems*, pages 3034–3042, 2015.
- Roger Fletcher. A new approach to variable metric algorithms. *The Computer Journal*, 13(3):317–322, 1970.
- Donald Goldfarb. A family of variable-metric methods derived by variational means. *Mathematics of Computation*, 24(109):23–26, 1970.
- Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in Neural Information Processing Systems*, pages 315–323, 2013.
- Yann LeCun and Corinna Cortes. The MNIST database of handwritten digits. 1998. URL <http://yann.lecun.com/exdb/mnist/>.
- Mu Li, Gary L. Miller, and Richard Peng. Iterative row sampling. In *IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 127–136. IEEE, 2013.
- Moshe Lichman. UCI machine learning repository, 2013. URL <http://archive.ics.uci.edu/ml>.
- Hongzhou Lin, Julien Mairal, and Zaïd Harchaoui. A universal catalyst for first-order optimization. In *Advances in Neural Information Processing Systems*, pages 3384–3392, 2015.
- Qihang Lin, Zhaosong Lu, and Lin Xiao. An accelerated proximal coordinate gradient method. In *Advances in Neural Information Processing Systems*, pages 3059–3067, 2014.
- Haipeng Luo, Alekh Agarwal, Nicolò Cesa-Bianchi, and John Langford. Efficient second order online learning by sketching. In *Advances in Neural Information Processing Systems*, pages 902–910, 2016.
- James Martens. Deep learning via Hessian-free optimization. In *International Conference on Machine Learning*, pages 735–742, 2010.
- Andrew McCallum. Real-sim, 1997. Available at <https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/binary.html#real-sim>.
- Aryan Mokhtari and Alejandro Ribeiro. RES: regularized stochastic BFGS algorithm. *IEEE Transactions on Signal Processing*, 62(23):6089–6104, 2014.
- Philipp Moritz, Robert Nishihara, and Michael Jordan. A linearly-convergent stochastic L-BFGS algorithm. In *Artificial Intelligence and Statistics*, pages 249–258, 2016.
- Arkadi Nemirovski. Interior point polynomial time methods in convex programming. *Lecture notes*, 2004.

- Yurii Nesterov. A method of solving a convex programming problem with convergence rate $O(1/k^2)$. In *Soviet Mathematics Doklady*, volume 27, pages 372–376, 1983.
- Yurii Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Springer Science & Business Media, 2013.
- Jorge Nocedal and Stephen Wright. *Numerical Optimization*. Springer Science & Business Media, 2006.
- Mert Pilanci and Martin J. Wainwright. Newton sketch: A linear-time optimization algorithm with linear-quadratic convergence. *arXiv preprint arXiv:1505.02250*, 2015.
- Herbert Robbins and Sutton Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, pages 400–407, 1951.
- 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.
- Nicol N. Schraudolph, Jin Yu, and Simon Günter. A stochastic quasi-Newton method for online convex optimization. In *Artificial Intelligence and Statistics*, pages 436–443, 2007.
- 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 *International Conference on Machine Learning*, pages 1000–1008, 2014.
- David F. Shanno. Conditioning of quasi-Newton methods for function minimization. *Mathematics of Computation*, 24(111):647–656, 1970.
- Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12(4):389–434, 2012.
- 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.
- Haishan Ye, Luo Luo, and Zhihua Zhang. A unifying framework for convergence analysis of approximate Newton methods. *arXiv preprint arXiv:1702.08124*, 2017.
- Lijun Zhang, Mehrdad Mahdavi, and Rong Jin. Linear convergence with condition number independent access of full gradients. In *Advances in Neural Information Processing Systems*, pages 980–988, 2013.