

Statistical inference using SGD

Tianyang Li¹

lty@cs.utexas.edu

Liu Liu¹

liuliu@utexas.edu

Anastasios Kyrillidis²

anastasios.kyrillidis@ibm.com

Constantine Caramanis¹

constantine@utexas.edu

¹ The University of Texas at Austin

² IBM T.J. Watson Research Center, Yorktown Heights

Abstract

We present a novel method for frequentist statistical inference in M -estimation problems, based on stochastic gradient descent (SGD) *with a fixed step size*: we demonstrate that the average of such SGD sequences can be used for statistical inference, after proper scaling. An intuitive analysis using the Ornstein-Uhlenbeck process suggests that such averages are asymptotically normal. To show the merits of our scheme, we apply it to both synthetic and real data sets, and demonstrate that its accuracy is comparable to classical statistical methods, while requiring potentially far less computation.

1 Introduction

In M -estimation, the minimization of empirical risk functions (RFs) provides point estimates of the model parameters. Statistical inference then seeks to assess the quality of these estimates; *e.g.*, by obtaining confidence intervals or solving hypothesis testing problems. Within this context, a classical result in statistics states that the asymptotic distribution of the empirical RF's minimizer is normal, centered around the population RF's minimizer (van der Vaart, 2000). Thus, given the mean and covariance of this normal distribution, we can infer a range of values, along with probabilities, that allows us to quantify the probability that this interval includes the true minimizer.

The Bootstrap (Efron, 1982; Efron and Tibshirani, 1994) is a classical tool for obtaining estimates of the mean and covariance of this distribution. The Bootstrap operates by generating samples from this distribution (usually, by re-sampling with or without replacement from the entire data set) and repeating the estimation procedure over these different re-samplings. As parameter dimensionality and data size grow, the Bootstrap becomes increasingly –even prohibitively– expensive.

In this context, we follow a different path: we show that inference can also be accomplished by directly using stochastic gradient descent (SGD), both for point estimates and inference, *with a fixed step size over the data set*. It is well-established that fixed step-size SGD is by and large the dominant method used for large scale data analysis. We prove, and also demonstrate empirically, that *the average of SGD*

sequences, obtained by minimizing RFs, can also be used for statistical inference. Unlike the Bootstrap, our approach does not require creating many large-size subsamples from the data, neither re-running SGD from scratch for each of these subsamples. Our method only uses first order information from gradient computations, and does not require any second order information. Both of these are important for large scale problems, where re-sampling many times, or computing Hessians, may be computationally prohibitive.

Outline and main contributions: This paper studies and analyzes a simple, *fixed step size*¹, SGD-based algorithm for inference in M -estimation problems. Our algorithm produces samples, whose covariance converges to the covariance of the M -estimate, without relying on bootstrap-based schemes, and also avoiding direct and costly computation of second order information. Much work has been done on the asymptotic normality of SGD, as well as on the Stochastic Gradient Langevin Dynamics (and variants) in the Bayesian setting. As we discuss in detail in Section 4, this is the first work to provide finite sample inference results, using fixed step size, and without imposing overly restrictive assumptions on the convergence of fixed step size SGD.

The remainder of the paper is organized as follows. In the next section, we define the inference problem for M -estimation, and recall basic results of asymptotic normality and how these are used. Section 3 is the main body of the paper: we provide the algorithm for creating bootstrap-like samples, and also provide the main theorem of this work. As the details are involved, we provide an intuitive analysis of our algorithm and explanation of our main results, using an asymptotic Ornstein-Uhlenbeck process approximation for the SGD process (Kushner and Huang, 1981; Pflug, 1986; Benveniste, Priouret, and Métivier, 1990; Kushner and Yin, 2003; Mandt, Hoffman, and Blei, 2016), and we postpone the full proof until the appendix. We specialize our main theorem to the case of linear regression (see supplementary material), and also that of logistic regression. For logistic regression in particular, we require a somewhat different

¹*Fixed step size* means we use the same step size every iteration, but the step size is smaller with more total number of iterations. In contrast, *constant step size* means the step size is constant no matter how many iterations taken.

approach, as the logistic regression objective is not strongly convex. In Section 4, we present related work and elaborate how this work differs from existing research in the literature. Finally, in the experimental section, we provide parts of our numerical experiments that illustrate the behavior of our algorithm, and corroborate our theoretical findings. We do this using synthetic data for linear and logistic regression, and also by considering the Higgs detection (Baldi, Sadowski, and Whiteson, 2014) and the LIBSVM Splice data sets. A considerably expanded set of empirical results is deferred to the appendix.

Supporting our theoretical results, our empirical findings suggest that the SGD inference procedure produces results similar to bootstrap while using far fewer operations. Thereby, we produce a more efficient inference procedure applicable in large scale settings, where other approaches fail.

2 Statistical inference for M -estimators

Consider the problem of estimating a set of parameters $\theta^* \in \mathbb{R}^p$ using n samples $\{X_i\}_{i=1}^n$, drawn from some distribution P on the sample space \mathcal{X} . In frequentist inference, we are interested in estimating the minimizer θ^* of the population risk:

$$\theta^* = \operatorname{argmin}_{\theta \in \mathbb{R}^p} \mathbb{E}_P[f(\theta; X)] = \operatorname{argmin}_{\theta \in \mathbb{R}^p} \int_{\mathcal{X}} f(\theta; x) dP(x), \quad (1)$$

where we assume that $f(\cdot; x) : \mathbb{R}^p \rightarrow \mathbb{R}$ is real-valued and convex; further, we will use $\mathbb{E} \equiv \mathbb{E}_P$, unless otherwise stated. In practice, the distribution P is unknown. We thus estimate θ^* by solving an empirical risk minimization (ERM) problem, where we use the estimate $\hat{\theta}$:

$$\hat{\theta} = \operatorname{argmin}_{\theta \in \mathbb{R}^p} \frac{1}{n} \sum_{i=1}^n f(\theta; X_i). \quad (2)$$

Statistical inference consists of techniques for obtaining information beyond point estimates $\hat{\theta}$, such as confidence intervals. These can be performed if there is an asymptotic limiting distribution associated with $\hat{\theta}$ (Wasserman, 2013). Indeed, under standard and well-understood regularity conditions, the solution to M -estimation problems satisfies asymptotic normality. That is, the distribution $\sqrt{n}(\hat{\theta} - \theta^*)$ converges weakly to a normal distribution:

$$\sqrt{n}(\hat{\theta} - \theta^*) \rightarrow \mathcal{N}(0, H^{*-1} G^* H^{*-1}), \quad (3)$$

where

$$H^* = \mathbb{E}[\nabla^2 f(\theta^*; X)],$$

and

$$G^* = \mathbb{E}[\nabla f(\theta^*; X) \cdot \nabla f(\theta^*; X)^\top];$$

see also Theorem 5.21 in (van der Vaart, 2000). We can therefore use this result, as long as we have a good estimate of the covariance matrix: $H^{*-1} G^* H^{*-1}$. The central goal of this paper is obtaining accurate estimates for $H^{*-1} G^* H^{*-1}$.




Figure 1: Our SGD inference procedure

A naive way to estimate $H^{*-1} G^* H^{*-1}$ is through the empirical estimator $\hat{H}^{-1} \hat{G} \hat{H}^{-1}$ where:

$$\begin{aligned} \hat{H} &= \frac{1}{n} \sum_{i=1}^n \nabla^2 f(\hat{\theta}; X_i) \quad \text{and} \\ \hat{G} &= \frac{1}{n} \sum_{i=1}^n \nabla f(\hat{\theta}; X_i) \nabla f(\hat{\theta}; X_i)^\top. \end{aligned} \quad (4)$$

Beyond calculating² \hat{H} and \hat{G} , this computation requires an inversion of \hat{H} and matrix-matrix multiplications in order to compute $\hat{H}^{-1} \hat{G} \hat{H}^{-1}$ —a key computational bottleneck in high dimensions. Instead, our method uses SGD to directly estimate $\hat{H}^{-1} \hat{G} \hat{H}^{-1}$.

3 Statistical inference using SGD

Consider the optimization problem in (2). For instance, in maximum likelihood estimation (MLE), $f(\theta; X_i)$ is a negative log-likelihood function. For simplicity of notation, we use $f_i(\theta)$ and $f(\theta)$ for $f(\theta; X_i)$ and $\frac{1}{n} \sum_{i=1}^n f(\theta; X_i)$, respectively, for the rest of the paper.

The SGD algorithm with a fixed step size η , is given by the iteration

$$\theta_{t+1} = \theta_t - \eta g_s(\theta_t), \quad (5)$$

where $g_s(\cdot)$ is an unbiased estimator of the gradient, *i.e.*, $\mathbb{E}[g_s(\theta) \mid \theta] = \nabla f(\theta)$, where the expectation is w.r.t. the stochasticity in the $g_s(\cdot)$ calculation. A classical example of an unbiased estimator of the gradient is $g_s(\cdot) \equiv \nabla f_i(\cdot)$, where i is a uniformly random index over the samples X_i .

Our inference procedure uses the average of t consecutive SGD iterations. In particular, the algorithm proceeds as follows: Given a sequence of SGD iterates, we use the first SGD iterates $\theta_{-b}, \theta_{-b+1}, \dots, \theta_0$ as a burn in period; we discard these iterates. Next, for each “segment” of $t + d$ iterates, we

²In the case of maximum likelihood estimation, we have $H^* = G^*$ —which is called Fisher information. Thus, the covariance of interest is $H^{*-1} = G^{*-1}$. This can be estimated either using \hat{H} or \hat{G} .

use the first t iterates to compute $\bar{\theta}_t^{(i)} = \frac{1}{t} \sum_{j=1}^t \theta_j^{(i)}$ and discard the last d iterates, where i indicates the i -th segment. This procedure is illustrated in Figure 1. As the final empirical minimum $\hat{\theta}$, we use in practice $\hat{\theta} \approx \frac{1}{R} \sum_{i=1}^R \bar{\theta}_t^{(i)}$ (Bubeck, 2015).

Some practical aspects of our scheme are discussed below.

Step size η selection and length t : Theorem 1 below is consistent only for SGD with fixed step size that depends on the number of samples taken. Our experiments, however, demonstrate that choosing a constant (large) η gives equally accurate results with significantly reduced running time. We conjecture that a better understanding of t 's and η 's influence requires stronger bounds for SGD with constant step size. Heuristically, calibration methods for parameter tuning in subsampling methods ((Politis, Romano, and Wolf, 2012), Ch. 9) could be used for hyper-parameter tuning in our SGD procedure. We leave the problem of finding maximal (provable) learning rates for future work.

Discarded length d : Based on the analysis of mean estimation in the appendix, if we discard d SGD iterates in every segment, the correlation between consecutive $\theta^{(i)}$ and $\theta^{(i+1)}$ is of the order of $C_1 e^{-C_2 \eta d}$, where C_1 and C_2 are data dependent constants. This can be used as a rule of thumb to reduce correlation between samples from our SGD inference procedure.

Burn-in period b : The purpose of the burn-in period b , is to ensure that samples are generated when SGD iterates are sufficiently close to the optimum. This can be determined using heuristics for SGD convergence diagnostics. Another approach is to use other methods (e.g., SVRG (Johnson and Zhang, 2013)) to find the optimum, and use a relatively small b for SGD to reach stationarity, similar to Markov Chain Monte Carlo burn-in.

Statistical inference using $\bar{\theta}_t^{(i)}$ and $\hat{\theta}$: Similar to ensemble learning (Opitz and Maclin, 1999), we use $i = 1, 2, \dots, R$ estimators for statistical inference:

$$\theta^{(i)} = \hat{\theta} + \sqrt{\frac{K_s \cdot t}{n}} (\bar{\theta}_t^{(i)} - \hat{\theta}). \quad (6)$$

Here, K_s is a scaling factor that depends on how the stochastic gradient g_s is computed. We show examples of K_s for mini batch SGD in linear regression and logistic regression in the corresponding sections. Similar to other resampling methods such as bootstrap and subsampling, we use quantiles or variance of $\theta^{(1)}, \theta^{(2)}, \dots, \theta^{(R)}$ for statistical inference.

3.1 Theoretical guarantees

Next, we provide the main theorem of our paper. Essentially, this provides conditions under which our algorithm is guaranteed to succeed, and hence has inference capabilities.

Theorem 1. *For a differentiable convex function $f(\theta) = \frac{1}{n} \sum_{i=1}^n f_i(\theta)$, with gradient $\nabla f(\theta)$, let $\hat{\theta} \in \mathbb{R}^p$ be its minimizer, according to (2), and denote its Hessian at $\hat{\theta}$ by $H := \nabla^2 f(\hat{\theta}) = \frac{1}{n} \cdot \sum_{i=1}^n \nabla^2 f_i(\hat{\theta})$. Assume that $\forall \theta \in \mathbb{R}^p$, f satisfies:*

$$(F_1) \text{ Weak strong convexity: } (\theta - \hat{\theta})^\top \nabla f(\theta) \geq \alpha \|\theta - \hat{\theta}\|_2^2, \text{ for constant } \alpha > 0,$$

$$(F_2) \text{ Lipschitz gradient continuity: } \|\nabla f(\theta)\|_2 \leq L \|\theta - \hat{\theta}\|_2, \text{ for constant } L > 0,$$

$$(F_3) \text{ Bounded Taylor remainder: } \|\nabla f(\theta) - H(\theta - \hat{\theta})\|_2 \leq E \|\theta - \hat{\theta}\|_2^2, \text{ for constant } E > 0,$$

$$(F_4) \text{ Bounded Hessian spectrum at } \hat{\theta}: 0 < \lambda_L \leq \lambda_i(H) \leq \lambda_U < \infty, \forall i.$$

Furthermore, let $g_s(\theta)$ be a stochastic gradient of f , satisfying:

$$(G_1) \mathbb{E}[g_s(\theta) | \theta] = \nabla f(\theta),$$

$$(G_2) \mathbb{E}[\|g_s(\theta)\|_2^2 | \theta] \leq A \|\theta - \hat{\theta}\|_2^2 + B,$$

$$(G_3) \mathbb{E}[\|g_s(\theta)\|_2^4 | \theta] \leq C \|\theta - \hat{\theta}\|_2^4 + D,$$

$$(G_4) \|\mathbb{E}[g_s(\theta)g_s(\theta)^\top | \theta] - G\|_2 \leq A_1 \|\theta - \hat{\theta}\|_2 + A_2 \|\theta - \hat{\theta}\|_2^2 + A_3 \|\theta - \hat{\theta}\|_2^3 + A_4 \|\theta - \hat{\theta}\|_2^4,$$

where $G = \mathbb{E}[g_s(\hat{\theta})g_s(\hat{\theta})^\top | \hat{\theta}]$ and, for positive, data dependent constants A, B, C, D, A_i , for $i = 1, \dots, 4$.

Assume that $\|\theta_1 - \hat{\theta}\|_2^2 = O(\eta)$; then for sufficiently small step size $\eta > 0$, the average SGD sequence, $\bar{\theta}_t$, satisfies:

$$\begin{aligned} & \left\| t\mathbb{E}[(\bar{\theta}_t - \hat{\theta})(\bar{\theta}_t - \hat{\theta})^\top] - H^{-1}GH^{-1} \right\|_2 \\ & \lesssim \sqrt{\eta} + \sqrt{\frac{1}{t\eta} + t\eta^2}. \end{aligned} \quad (7)$$

We provide the full proof in the appendix, and also we give precise (data-dependent) formulas for the above constants. For ease of exposition, we leave them as constants in the expressions above. Further, in the next section, we relate a continuous approximation of SGD to Ornstein-Uhlenbeck process (Robbins and Monro, 1951) to give an intuitive explanation of our results.

Discussion. For linear regression, assumptions (F_1) , (F_2) , (F_3) , and (F_4) are satisfied when the empirical risk function is not degenerate. In mini batch SGD using sampling with replacement, assumptions (G_1) , (G_2) , (G_3) , and (G_4) are satisfied. Linear regression's result is presented in Corollary 2 in the appendix.

For logistic regression, assumption (F_1) is not satisfied because the empirical risk function in this case is strictly but not strongly convex. Thus, we cannot apply Theorem 1 directly. Instead, we consider the use of SGD on the *square of the empirical risk function plus a constant*; see eq. (11) below. When the empirical risk function is not degenerate, (11) satisfies assumptions (F_1) , (F_2) , (F_3) , and (F_4) . We cannot directly use vanilla SGD to minimize (11), instead we describe a modified SGD procedure for minimizing (11) in Section 3.3, which satisfies assumptions (G_1) , (G_2) , (G_3) , and (G_4) . We believe that this result is of interest by its own. We present the result specialized for logistic regression in Corollary 1.

Note that Theorem 1 proves consistency for SGD with fixed step size, requiring $\eta \rightarrow 0$ when $t \rightarrow \infty$. However, we empirically observe in our experiments that a sufficiently large constant η gives better results. We conjecture that the average of consecutive iterates in SGD with *larger constant step size* converges to the optimum and we consider it for future work.

3.2 Intuitive interpretation via the Ornstein-Uhlenbeck process approximation

Here, we describe a continuous approximation of the discrete SGD process and relate it to the Ornstein-Uhlenbeck process (Robbins and Monro, 1951), to give an intuitive explanation of our results. In particular, under regularity conditions, the stochastic process $\Delta_t = \theta_t - \hat{\theta}$ asymptotically converges to an Ornstein-Uhlenbeck process $\Delta(t)$, (Kushner and Huang, 1981; Pflug, 1986; Benveniste, Priouret, and Métivier, 1990; Kushner and Yin, 2003; Mandt, Hoffman, and Blei, 2016) that satisfies:

$$d\Delta(T) = -H\Delta(T) dT + \sqrt{\eta}G^{\frac{1}{2}} dB(T), \quad (8)$$

where $B(T)$ is a standard Brownian motion. Given (8), $\sqrt{t}(\bar{\theta}_t - \hat{\theta})$ can be approximated as

$$\begin{aligned} \sqrt{t}(\bar{\theta}_t - \hat{\theta}) &= \frac{1}{\sqrt{t}} \sum_{i=1}^t (\theta_i - \hat{\theta}) \\ &= \frac{1}{\eta\sqrt{t}} \sum_{i=1}^t (\theta_i - \hat{\theta})\eta \approx \frac{1}{\eta\sqrt{t}} \int_0^{t\eta} \Delta(T) dT, \end{aligned} \quad (9)$$

where we use the approximation that $\eta \approx dT$. By rearranging terms in (8) and multiplying both sides by H^{-1} , we can rewrite the stochastic differential equation (8) as $\Delta(T) dT = -H^{-1} d\Delta(T) + \sqrt{\eta}H^{-1}G^{\frac{1}{2}} dB(T)$. Thus, we have

$$\begin{aligned} \int_0^{t\eta} \Delta(T) dT &= \\ -H^{-1}(\Delta(t\eta) - \Delta(0)) + \sqrt{\eta}H^{-1}G^{\frac{1}{2}}B(t\eta). \end{aligned} \quad (10)$$

After plugging (10) into (9) we have

$$\begin{aligned} \sqrt{t}(\bar{\theta}_t - \hat{\theta}) &\approx \\ -\frac{1}{\eta\sqrt{t}}H^{-1}(\Delta(t\eta) - \Delta(0)) + \frac{1}{\sqrt{t\eta}}H^{-1}G^{\frac{1}{2}}B(t\eta). \end{aligned}$$

When $\Delta(0) = 0$, the variance $\text{Var}[-1/\eta\sqrt{t} \cdot H^{-1}(\Delta(t\eta) - \Delta(0))] = O(1/t\eta)$. Since $1/\sqrt{t\eta} \cdot H^{-1}G^{\frac{1}{2}}B(t\eta) \sim \mathcal{N}(0, H^{-1}GH^{-1})$, when $\eta \rightarrow 0$ and $t\eta \rightarrow \infty$, we conclude that

$$\sqrt{t}(\bar{\theta}_t - \hat{\theta}) \sim \mathcal{N}(0, H^{-1}GH^{-1}).$$

3.3 Logistic regression

We next apply our method to logistic regression. We have n samples $(X_1, y_1), (X_2, y_2), \dots, (X_n, y_n)$ where $X_i \in \mathbb{R}^p$ consists of features and $y_i \in \{+1, -1\}$ is the label. We estimate θ of a linear classifier $\text{sign}(\theta^T X)$ by:

$$\hat{\theta} = \underset{\theta \in \mathbb{R}^p}{\text{argmin}} \frac{1}{n} \sum_{i=1}^n \log(1 + \exp(-y_i \theta^T X_i)).$$

We cannot apply Theorem 1 directly because the empirical logistic risk is not strongly convex; it does not satisfy

assumption (F_1) . Instead, we consider the convex function

$$f(\theta) = \frac{1}{2} \left(c + \frac{1}{n} \sum_{i=1}^n \log(1 + \exp(-y_i \theta^T X_i)) \right)^2, \quad \text{where } c > 0 \text{ (e.g., } c = 1). \quad (11)$$

The gradient of $f(\theta)$ is a product of two terms

$$\begin{aligned} \nabla f(\theta) &= \underbrace{\left(c + \frac{1}{n} \sum_{i=1}^n \log(1 + \exp(-y_i \theta^T X_i)) \right)}_{\Psi} \times \\ &\quad \underbrace{\nabla \left(\frac{1}{n} \sum_{i=1}^n \log(1 + \exp(-y_i \theta^T X_i)) \right)}_{\Upsilon}. \end{aligned}$$

Therefore, we can compute $g_s = \Psi_s \Upsilon_s$, using two independent random variables satisfying $\mathbb{E}[\Psi_s | \theta] = \Psi$ and $\mathbb{E}[\Upsilon_s | \theta] = \Upsilon$. For Υ_s , we have $\Upsilon_s = \frac{1}{S_\Upsilon} \sum_{i \in I_\Upsilon^s} \nabla \log(1 + \exp(-y_i \theta^T X_i))$, where I_Υ^s are S_Υ indices sampled from $[n]$ uniformly at random with replacement. For Ψ_s , we have $\Psi_s = c + \frac{1}{S_\Psi} \sum_{i \in I_\Psi^s} \log(1 + \exp(-y_i \theta^T X_i))$, where I_Ψ^s are S_Ψ indices uniformly sampled from $[n]$ with or without replacement. Given the above, we have $\nabla f(\theta)^T (\theta - \hat{\theta}) \geq \alpha \|\theta - \hat{\theta}\|_2^2$ for some constant α by the generalized self-concordance of logistic regression (Bach, 2010, 2014), and therefore the assumptions are now satisfied.

For convenience, we write $k(\theta) = \frac{1}{n} \sum_{i=1}^n k_i(\theta)$ where $k_i(\theta) = \log(1 + \exp(-y_i \theta^T X_i))$. Thus $f(\theta) = (k(\theta) + c)^2$, $\mathbb{E}[\Psi_s | \theta] = k(\theta) + c$, and $\mathbb{E}[\Upsilon_s | \theta] = \nabla k(\theta)$.

Corollary 1. Assume $\|\theta_1 - \hat{\theta}\|_2^2 = O(\eta)$; also $S_\Psi = O(1)$, $S_\Upsilon = O(1)$ are bounded. Then, we have

$$\left\| t\mathbb{E}[(\bar{\theta}_t - \hat{\theta})(\bar{\theta}_t - \hat{\theta})^T] - H^{-1}GH^{-1} \right\|_2 \lesssim \sqrt{\eta} + \sqrt{\frac{1}{t\eta} + t\eta^2},$$

where $H = \nabla^2 f(\hat{\theta}) = (c + k(\hat{\theta}))\nabla^2 k(\hat{\theta})$. Here, $G = \frac{1}{S_\Upsilon} K_G(\hat{\theta}) \frac{1}{n} \sum_{i=1}^n \nabla k_i(\hat{\theta}) k_i(\hat{\theta})^T$ with $K_G(\theta) = \mathbb{E}[\Psi(\theta)^2]$ depending on how indexes are sampled to compute Ψ_s :

- with replacement: $K_G(\theta) = \frac{1}{S_\Psi} (\frac{1}{n} \sum_{i=1}^n (c + k_i(\theta))^2) + \frac{S_\Psi - 1}{S_\Psi} (c + k(\theta))^2$,
- no replacement: $K_G(\theta) = \frac{1 - \frac{S_\Psi - 1}{n-1}}{S_\Psi} (\frac{1}{n} \sum_{i=1}^n (c + k_i(\theta))^2) + \frac{S_\Psi - 1}{S_\Psi} \frac{n}{n-1} (c + k(\theta))^2$.

Quantities other than t and η are data dependent constants.

As with the results above, in the appendix we give data-dependent expressions for the constants. Simulations suggest that the term $t\eta^2$ in our bound is an artifact of our analysis.

Because in logistic regression the estimate's covariance is $\frac{(\nabla^2 k(\hat{\theta}))^{-1}}{n} \left(\frac{\sum_{i=1}^n \nabla k_i(\hat{\theta}) \nabla k_i(\hat{\theta})^T}{n} \right) (\nabla^2 k(\hat{\theta}))^{-1}$, we set the

scaling factor $K_s = \frac{(c+k(\hat{\theta}))^2}{K_G(\hat{\theta})}$ in (6) for statistical inference.

Note that $K_s \approx 1$ for sufficiently large S_Ψ .

4 Related work

Bayesian inference: First and second order iterative optimization algorithms—including SGD, gradient descent, and

variants—naturally define a Markov chain. Based on this principle, most related to this work is the case of stochastic gradient Langevin dynamics (SGLD) for Bayesian inference—namely, for sampling from the posterior distributions—using a variant of SGD (Welling and Teh, 2011; Bubeck, Eldan, and Lehec, 2015; Mandt, Hoffman, and Blei, 2016, 2017). We note that, here as well, the vast majority of the results rely on using a decreasing step size. Very recently, (Mandt, Hoffman, and Blei, 2017) uses a heuristic approximation for Bayesian inference, and provides results for fixed step size.

Our problem is different in important ways from the Bayesian inference problem. In such parameter estimation problems, the covariance of the estimator only depends on the gradient of the likelihood function. This is not the case, however, in general frequentist M -estimation problems (e.g., linear regression). In these cases, the covariance of the estimator depends both on the gradient and Hessian of the empirical risk function. For this reason, without second order information, SGLD methods are poorly suited for general M -estimation problems in frequentist inference. In contrast, our method exploits properties of averaged SGD, and computes the estimator’s covariance without second order information.

Connection with Bootstrap methods: The classical approach for statistical inference is to use the bootstrap (Efron and Tibshirani, 1994; Shao and Tu, 2012). Bootstrap samples are generated by replicating the entire data set by resampling, and then solving the optimization problem on each generated set of the data. We identify our algorithm and its analysis as an alternative to bootstrap methods. Our analysis is also specific to SGD, and thus sheds light on the statistical properties of this very widely used algorithm.

Connection with stochastic approximation methods: It has been long observed in stochastic approximation that under certain conditions, SGD displays asymptotic normality for both the setting of *decreasing step size*, e.g., (Ljung, Pflug, and Walk, 2012; Polyak and Juditsky, 1992), and more recently, (Toulis and Airoldi, 2014; Chen et al., 2016); and also for *fixed step size*, e.g., (Benveniste, Priouret, and Métivier, 1990), Chapter 4. All of these results, however, provide their guarantees with the requirement that the stochastic approximation iterate converges to the optimum. For decreasing step size, this is not an overly burdensome assumption, since with mild assumptions it can be shown directly. As far as we know, however, it is not clear if this holds in the fixed step size regime. To side-step this issue, (Benveniste, Priouret, and Métivier, 1990) provides results only when the (constant) step-size approaches 0 (see Section 4.4 and 4.6, and in particular Theorem 7 in (Benveniste, Priouret, and Métivier, 1990)). Similarly, while (Kushner and Yin, 2003) has asymptotic results on the average of consecutive stochastic approximation iterates with constant step size, it assumes convergence of iterates (assumption A1.7 in Ch. 10)—an assumption we are unable to justify in even simple settings.

Beyond the critical difference in the assumptions, the majority of the “classical” subject matter seeks to prove asymptotic results about different flavors of SGD, but does not properly consider its use for inference. Key exceptions are the recent work in (Toulis and Airoldi, 2014) and (Chen et al., 2016), which follow up on (Polyak and Juditsky, 1992). Both

of these rely on decreasing step size, for reasons mentioned above. The work in (Chen et al., 2016) uses SGD with decreasing step size for estimating an M -estimate’s covariance. Work in (Toulis and Airoldi, 2014) studies implicit SGD with decreasing step size and proves results similar to (Polyak and Juditsky, 1992), however it does not use SGD to compute confidence intervals.

Overall, to the best of our knowledge, there are no prior results establishing asymptotic normality for SGD with fixed step size for general M -estimation problems (that do not rely on overly restrictive assumptions, as discussed).

5 Experiments

5.1 Synthetic data

The coverage probability is defined as $\frac{1}{p} \sum_{i=1}^p \mathbb{P}[\theta_i^* \in \hat{C}_i]$ where $\theta^* = \operatorname{argmin}_{\theta} \mathbb{E}[f(\theta, X)] \in \mathbb{R}^p$, and \hat{C}_i is the estimated confidence interval for the i^{th} coordinate. The average confidence interval width is defined as $\frac{1}{p} \sum_{i=1}^p (\hat{C}_i^u - \hat{C}_i^l)$ where $[\hat{C}_i^l, \hat{C}_i^u]$ is the estimated confidence interval for the i^{th} coordinate. In our experiments, coverage probability and average confidence interval width are estimated through simulation. We use the empirical quantile of our SGD inference procedure and bootstrap to compute the 95% confidence intervals for each coordinate of the parameter. For results given as a pair (α, β) , it usually indicates (coverage probability, confidence interval length).

Univariate models In Figure 2, we compare our SGD inference procedure with (i) Bootstrap and (ii) normal approximation with inverse Fisher information in univariate models. We observe that our method and Bootstrap have similar statistical properties. Figure 5 in the appendix shows Q-Q plots of samples from our SGD inference procedure.

Normal distribution mean estimation: Figure 2a compares 500 samples from SGD inference procedure and Bootstrap versus the distribution $\mathcal{N}(0, 1/n)$, using $n = 20$ i.i.d. samples from $\mathcal{N}(0, 1)$. We used mini batch SGD described in Sec. A. For the parameters, we used $\eta = 0.8$, $t = 5$, $d = 10$, $b = 20$, and mini batch size of 2. Our SGD inference procedure gives (0.916, 0.806), Bootstrap gives (0.926, 0.841), and normal approximation gives (0.922, 0.851).

Exponential distribution parameter estimation: Figure 2b compares 500 samples from inference procedure and Bootstrap, using $n = 100$ samples from an exponential distribution with PDF $\lambda e^{-\lambda x}$ where $\lambda = 1$. We used SGD for MLE with mini batch sampled with replacement. For the parameters, we used $\eta = 0.1$, $t = 100$, $d = 5$, $b = 100$, and mini batch size of 5. Our SGD inference procedure gives (0.922, 0.364), Bootstrap gives (0.942, 0.392), and normal approximation gives (0.922, 0.393).

Poisson distribution parameter estimation: Figure 2c compares 500 samples from inference procedure and Bootstrap, using $n = 100$ samples from a Poisson distribution with PDF $\lambda^x e^{-\lambda}$ where $\lambda = 1$. We used SGD for MLE with mini batch sampled with replacement. For the parameters, we used $\eta = 0.1$, $t = 100$, $d = 5$, $b = 100$, and mini batch




Figure 2: Estimation in univariate models.

η	$t = 100$	$t = 500$	$t = 2500$	η	$t = 100$	$t = 500$	$t = 2500$
0.1	(0.957, 4.41)	(0.955, 4.51)	(0.960, 4.53)	0.1	(0.949, 4.74)	(0.962, 4.91)	(0.963, 4.94)
0.02	(0.869, 3.30)	(0.923, 3.77)	(0.918, 3.87)	0.02	(0.845, 3.37)	(0.916, 4.01)	(0.927, 4.17)
0.004	(0.634, 2.01)	(0.862, 3.20)	(0.916, 3.70)	0.004	(0.616, 2.00)	(0.832, 3.30)	(0.897, 3.93)

(a) Bootstrap (0.941, 4.14), normal approximation (0.928, 3.87) (b) Bootstrap (0.938, 4.47), normal approximation (0.925, 4.18)

Table 1: Linear regression. *Left*: Experiment 1, *Right*: Experiment 2.

η	$t = 100$	$t = 500$	$t = 2500$	η	$t = 100$	$t = 500$	$t = 2500$
0.1	(0.872, 0.204)	(0.937, 0.249)	(0.939, 0.258)	0.1	(0.859, 0.206)	(0.931, 0.255)	(0.947, 0.266)
0.02	(0.610, 0.112)	(0.871, 0.196)	(0.926, 0.237)	0.02	(0.600, 0.112)	(0.847, 0.197)	(0.931, 0.244)
0.004	(0.312, 0.051)	(0.596, 0.111)	(0.86, 0.194)	0.004	(0.302, 0.051)	(0.583, 0.111)	(0.851, 0.195)

(a) Bootstrap (0.932, 0.253), normal approximation (0.957, 0.264) (b) Bootstrap (0.932, 0.245), normal approximation (0.954, 0.256)

Table 2: Logistic regression. *Left*: Experiment 1, *Right*: Experiment 2.

size of 5. Our SGD inference procedure gives (0.942, 0.364), Bootstrap gives (0.946, 0.386), and normal approximation gives (0.960, 0.393).

Multivariate models In these experiments, we set $d = 100$, used mini-batch size of 4, and used 200 SGD samples. In all cases, we compared with Bootstrap using 200 replicates. We computed the coverage probabilities using 500 simulations. Also, we denote $1_p = [1 \ 1 \ \dots \ 1]^\top \in \mathbb{R}^p$. Additional simulations comparing covariance matrix computed with different methods are given in Sec. D.1.

Linear regression: *Experiment 1:* Results for the case where $X \sim \mathcal{N}(0, I) \in \mathbb{R}^{10}$, $Y = w^{*T}X + \epsilon$, $w^* = 1_p/\sqrt{p}$, and $\epsilon \sim \mathcal{N}(0, \sigma^2 = 10^2)$ with $n = 100$ samples is given in Table 1a. Bootstrap gives (0.941, 4.14), and confidence intervals computed using the error covariance and normal approximation gives (0.928, 3.87). *Experiment 2:* Results for the case where $X \sim \mathcal{N}(0, \Sigma) \in \mathbb{R}^{10}$, $\Sigma_{ij} = 0.3^{|i-j|}$, $Y = w^{*T}X + \epsilon$, $w^* = 1_p/\sqrt{p}$, and $\epsilon \sim \mathcal{N}(0, \sigma^2 = 10^2)$ with $n = 100$ samples is given in Table 1b. Bootstrap gives (0.938, 4.47), and confidence intervals computed using the error covariance and normal approximation gives (0.925, 4.18).

Logistic regression: Here we show results for logistic regression trained using vanilla SGD with mini batch sampled with replacement. Results for modified SGD (Sec. 3.3) are given in Sec. D.1. *Experiment 1:* Results for the case where $\mathbb{P}[Y = +1] = \mathbb{P}[Y = -1] = 1/2$, $X | Y \sim \mathcal{N}(0.01Y1_p/\sqrt{p}, I) \in \mathbb{R}^{10}$ with $n = 1000$ samples is given in Table 2a. Bootstrap gives (0.932, 0.245), and confidence intervals computed using inverse Fisher matrix as the error covariance and normal approximation gives (0.954, 0.256). *Experiment 2:* Results for the case where $\mathbb{P}[Y = +1] = \mathbb{P}[Y = -1] = 1/2$, $X | Y \sim \mathcal{N}(0.01Y1_p/\sqrt{p}, \Sigma) \in \mathbb{R}^{10}$, $\Sigma_{ij} = 0.2^{|i-j|}$ with $n = 1000$ samples is given in Table 2b. Bootstrap gives (0.932, 0.253), and confidence intervals computed using inverse Fisher matrix as the error covariance and normal approximation gives (0.957, 0.264).

5.2 Real data

Here, we compare covariance matrices computed using our SGD inference procedure, bootstrap, and inverse Fisher information matrix on the LIBSVM Splice data set, and we observe that they have similar statistical properties.

Splice data set The Splice data set³ contains 60 distinct features with 1000 data samples. This is a classification problem between two classes of splice junctions in a DNA sequence. We use a logistic regression model trained using vanilla SGD.

In Figure 3, we compare the covariance matrix computed using our SGD inference procedure and bootstrap $n = 1000$ samples. We used 10000 samples from both bootstrap and our SGD inference procedure with $t = 500$, $d = 100$, $\eta = 0.2$, and mini batch size of 6.




Figure 3: Splice data set




Figure 4: MNIST

MNIST Here, we train a binary logistic regression classifier to classify 0/1 using a noisy MNIST data set, and demonstrate that adversarial examples produced by gradient attack (Goodfellow, Shlens, and Szegedy, 2015) (perturbing an image in the direction of loss function’s gradient with respect to data) can be detected using prediction intervals. We flatten each 28×28 image into a 784 dimensional vector, and train a linear classifier using pixel values as features. To add noise to each image, where each original pixel is either 0 or 1, we randomly changed 70% pixels to random numbers uniformly on $[0, 0.9]$. Next we train the classifier on the noisy MNIST data set, and generate adversarial examples using this noisy MNIST data set. Figure 4 shows each image’s logit value ($\log \frac{\mathbb{P}[1|\text{image}]}{\mathbb{P}[0|\text{image}]}$) and its 95% confidence interval (CI) computed using quantiles from our SGD inference procedure.

³<https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/binary.html>

5.3 Discussion

In our experiments, we observed that using a larger step size η produces accurate results with significantly accelerated convergence time. This might imply that the η term in Theorem 1’s bound is an artifact of our analysis. Indeed, although Theorem 1 only applies to SGD with fixed step size, where $\eta t \rightarrow \infty$ and $\eta^2 t \rightarrow 0$ imply that the step size should be smaller when the number of consecutive iterates used for the average is larger, our experiments suggest that we can use a (data dependent) constant step size η and only require $\eta t \rightarrow \infty$.

In the experiments, our SGD inference procedure uses $(t + d) \cdot S \cdot p$ operations to produce a sample, and Newton method uses $n \cdot (\text{matrix inversion complexity} = \Omega(p^2)) \cdot (\text{number of Newton iterations } t)$ operations to produce a sample. The experiments therefore suggest that our SGD inference procedure produces results similar to Bootstrap while using far fewer operations.

6 Acknowledgments

This work was partially supported by NSF Grants 1609279, 1704778 and 1764037, and also by the USDOT through the Data-Supported Transportation Operations and Planning (D-STOP) Tier 1 University Transportation Center. A.K. is supported by the IBM Goldstine fellowship. We thank Xi Chen, Philipp Krähenbühl, Matthijs Snel, and Tom Spangenberg for insightful discussions.

References

- Bach, F. 2010. Self-concordant analysis for logistic regression. *Electronic Journal of Statistics* 4:384–414.
- Bach, F. 2014. Adaptivity of averaged stochastic gradient descent to local strong convexity for logistic regression. *Journal of Machine Learning Research* 15(1):595–627.
- Baldi, P.; Sadowski, P.; and Whiteson, D. 2014. Searching for exotic particles in high-energy physics with deep learning. *Nature Communications* 5.
- Benveniste, A.; Priouret, P.; and Métivier, M. 1990. *Adaptive Algorithms and Stochastic Approximations*. New York, NY, USA: Springer-Verlag New York, Inc.
- Bubeck, S.; Eldan, R.; and Lehec, J. 2015. Finite-time analysis of projected langevin monte carlo. In *Advances in Neural Information Processing Systems*, 1243–1251.
- Bubeck, S. 2015. Convex optimization: Algorithms and complexity. *Found. Trends Mach. Learn.* 8(3-4):231–357.
- Chen, X.; Lee, J.; Tong, X.; and Zhang, Y. 2016. Statistical inference for model parameters in stochastic gradient descent. *arXiv preprint arXiv:1610.08637*.
- Efron, B., and Tibshirani, R. 1994. *An introduction to the bootstrap*. CRC press.
- Efron, B. 1982. *The jackknife, the bootstrap and other resampling plans*. SIAM.
- Goodfellow, I.; Shlens, J.; and Szegedy, C. 2015. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations*.
- Johnson, R., and Zhang, T. 2013. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in Neural Information Processing Systems*, 315–323.

- Kushner, H., and Huang, H. 1981. Asymptotic properties of stochastic approximations with constant coefficients. *SIAM Journal on Control and Optimization* 19(1):87–105.
- Kushner, H., and Yin, G. 2003. *Stochastic Approximation and Recursive Algorithms and Applications*. Stochastic Modelling and Applied Probability. Springer New York.
- Ljung, L.; Pflug, G. C.; and Walk, H. 2012. *Stochastic approximation and optimization of random systems*, volume 17. Birkhäuser.
- Mandt, S.; Hoffman, M.; and Blei, D. 2016. A Variational Analysis of Stochastic Gradient Algorithms. In *Proceedings of The 33rd International Conference on Machine Learning*, 354–363.
- Mandt, S.; Hoffman, M. D.; and Blei, D. M. 2017. Stochastic Gradient Descent as Approximate Bayesian Inference. *arXiv preprint arXiv:1704.04289*.
- Opitz, D., and Maclin, R. 1999. Popular ensemble methods: An empirical study. *Journal of Artificial Intelligence Research* 11:169–198.
- Pflug, G. 1986. Stochastic minimization with constant step-size: asymptotic laws. *SIAM Journal on Control and Optimization* 24(4):655–666.
- Politis, D.; Romano, J.; and Wolf, M. 2012. *Subsampling*. Springer Series in Statistics. Springer New York.
- Polyak, B., and Juditsky, A. 1992. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization* 30(4):838–855.
- Robbins, H., and Monro, S. 1951. A stochastic approximation method. *The Annals of Mathematical Statistics* 400–407.
- Shao, J., and Tu, D. 2012. *The jackknife and bootstrap*. Springer Science & Business Media.
- Toulis, P., and Airolidi, E. M. 2014. Asymptotic and finite-sample properties of estimators based on stochastic gradients. *arXiv preprint arXiv:1408.2923*.
- van der Vaart, A. 2000. *Asymptotic Statistics*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.
- Wasserman, L. 2013. *All of statistics: a concise course in statistical inference*. Springer Science & Business Media.
- Welling, M., and Teh, Y. 2011. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the 28th International Conference on Machine Learning*, 681–688.

A Exact analysis of mean estimation

In this section, we give an exact analysis of our method in the least squares, mean estimation problem. For n i.i.d. samples X_1, X_2, \dots, X_n , the mean is estimated by solving the following optimization problem

$$\hat{\theta} = \operatorname{argmax}_{\theta \in \mathbb{R}^p} \frac{1}{n} \sum_{i=1}^n \frac{1}{2} \|X_i - \theta\|_2^2 = \frac{1}{n} \sum_{i=1}^n X_i.$$

In the case of mini-batch SGD, we sample $S = O(1)$ indexes uniformly randomly with replacement from $[n]$; denote that index set as I_t . For convenience, we write $Y_t = \frac{1}{S} \sum_{i \in I_t} X_i$. Then, in the t^{th} mini batch SGD step, the update step is

$$\theta_{t+1} = \theta_t - \eta(\theta_t - Y_t) = (1 - \eta)\theta_t + \eta Y_t, \quad (12)$$

which is the same as the exponential moving average. And we have

$$\sqrt{t}\hat{\theta}_t = -\frac{1}{\eta\sqrt{t}}(\theta_{t+1} - \theta_1) + \frac{1}{\sqrt{t}} \sum_{i=1}^n Y_i. \quad (13)$$

Assume that $\|\theta_1 - \hat{\theta}\|_2^2 = O(\eta)$, then from Chebyshev's inequality $-\frac{1}{\eta\sqrt{t}}(\theta_{t+1} - \theta_1) \rightarrow 0$ almost surely when $t\eta \rightarrow \infty$. By the central limit theorem, $\frac{1}{\sqrt{t}} \sum_{i=1}^n Y_i$ converges weakly to $\mathcal{N}(\hat{\theta}, \frac{1}{S}\hat{\Sigma})$ with $\hat{\Sigma} = \frac{1}{n} \sum_{i=1}^n (X_i - \hat{\theta})(X_i - \hat{\theta})^\top$. From (12), we have $\|\operatorname{Cov}(\theta_a, \theta_b)\|_2 = O(\eta(1 - \eta)^{|a-b|})$ uniformly for all a, b , where the constant is data dependent. Thus, for our SGD inference procedure, we have $\|\operatorname{Cov}(\theta^{(i)}, \theta^{(j)})\|_2 = O(\eta(1 - \eta)^{d+|i-j|})$. Our SGD inference procedure does not generate samples that are independent conditioned on the data, whereas replicates are independent conditioned on the data in bootstrap, but this suggests that our SGD inference procedure can produce ‘‘almost independent’’ samples if we discard sufficient number of SGD iterates in each segment.

When estimating a mean using our SGD inference procedure where each mini batch is S elements sampled with replacement, we set $K_s = S$ in (6).

B Linear Regression

In linear regression, the empirical risk function satisfies:

$$f(\theta) = \frac{1}{n} \sum_{i=1}^n \frac{1}{2} (\theta^\top x_i - y_i)^2,$$

where y_i denotes the observations of the linear model and x_i are the regressors. To find an estimate to θ^* , one can use SGD with stochastic gradient give by:

$$g_s[\theta_t] = \frac{1}{S} \sum_{i \in I_t} \nabla f_i(\theta_t),$$

where I_t are S indices uniformly sampled from $[n]$ with replacement.

Next, we state a special case of Theorem 1. Because the Taylor remainder $\nabla f(\theta) - H(\theta - \hat{\theta}) = 0$, linear regression has a stronger result than general M -estimation problems.

Corollary 2. Assume that $\|\theta_1 - \hat{\theta}\|_2^2 = O(\eta)$, we have

$$\left\| t\mathbb{E}[(\bar{\theta}_t - \hat{\theta})(\bar{\theta}_t - \hat{\theta})^\top] - H^{-1}GH^{-1} \right\|_2 \lesssim \sqrt{\eta} + \frac{1}{\sqrt{t\eta}},$$

where $H = \frac{1}{n} \sum_{i=1}^n x_i x_i^\top$ and $G = \frac{1}{S} \frac{1}{n} \sum_{i=1}^n (x_i^\top \hat{\theta} - y_i)^2 x_i x_i^\top$.

We assume that $S = O(1)$ is bounded, and quantities other than t and η are data dependent constants.

As with our main theorem, in the appendix we provide explicit data-dependent expressions for the constants in the result. Because in linear regression the estimate's covariance is $\frac{1}{n} (\frac{1}{n} \sum_{i=1}^n x_i x_i^\top)^{-1} (\frac{1}{n} (x_i^\top \hat{\theta} - y_i)(x_i^\top \hat{\theta} - y_i)^\top) (\frac{1}{n} \sum_{i=1}^n x_i x_i^\top)^{-1}$, we set the scaling factor $K_s = S$ in (6) for statistical inference.

C Proofs

C.1 Proof of Theorem 1

We first assume that $\theta_1 = \hat{\theta}$. For ease of notation, we denote

$$\Delta_t = \theta_t - \hat{\theta}, \quad (14)$$

and, without loss of generality, we assume that $\hat{\theta} = 0$. The stochastic gradient descent recursion satisfies:

$$\begin{aligned} \theta_{t+1} &= \theta_t - \eta \cdot g_s(\theta_t) \\ &= \theta_t - \eta \cdot (g_s(\theta_t) - \nabla f(\theta_t) + \nabla f(\theta_t)) \\ &= \theta_t - \eta \cdot \nabla f(\theta_t) - \eta \cdot e_t, \end{aligned}$$

where $e_t = g_s(\theta_t) - \nabla f(\theta_t)$. Note that e_1, e_2, \dots is a martingale difference sequence. We use

$$g_i = \nabla f_i(\hat{\theta}) \quad \text{and} \quad H_i = \nabla^2 f_i(\hat{\theta}) \quad (15)$$

to denote the gradient component at index i , and the Hessian component at index i , at optimum $\hat{\theta}$, respectively. Note that $\sum g_i = 0$ and $\frac{1}{n} \sum H_i = H$.

For each f_i , its Taylor expansion around $\hat{\theta}$ is

$$f_i(\theta) = f_i(\hat{\theta}) + g_i^\top (\theta - \hat{\theta}) + \frac{1}{2} (\theta - \hat{\theta})^\top H_i (\theta - \hat{\theta}) + R_i(\theta, \hat{\theta}), \quad (16)$$

where $R_i(\theta, \hat{\theta})$ is the remainder term. For convenience, we write $R = \frac{1}{n} \sum R_i$.

For the proof, we require the following lemmata. The following lemma states that $\mathbb{E}[\|\Delta_t\|_2^2] = O(\eta)$ as $t \rightarrow \infty$ and $\eta \rightarrow 0$.

Lemma 1. *For data dependent, positive constants α, A, B according to assumptions (F_1) and (G_2) in Theorem 1, and given assumption (G_1) , we have*

$$\mathbb{E} [\|\Delta_t\|_2^2] \leq (1 - 2\alpha\eta + A\eta^2)^{t-1} \|\Delta_1\|_2^2 + \frac{B\eta}{2\alpha - A\eta}, \quad (17)$$

under the assumption $\eta < \frac{2\alpha}{A}$.

Proof. As already stated, we assume without loss of generality that $\hat{\theta} = 0$. This further implies that: $g_s(\theta_t) = g_s(\theta_t - \hat{\theta}) = g_s(\Delta_t)$, and

$$\Delta_{t+1} = \Delta_t - \eta \cdot g_s(\Delta_t).$$

Given the above and assuming expectation $\mathbb{E}[\cdot]$ w.r.t. the selection of a sample from $\{X_i\}_{i=1}^n$, we have:

$$\begin{aligned} \mathbb{E} [\|\Delta_{t+1}\|_2^2 \mid \Delta_t] &= \mathbb{E} [\|\Delta_t - \eta g_s(\Delta_t)\|_2^2 \mid \Delta_t] \\ &= \mathbb{E} [\|\Delta_t\|_2^2 \mid \Delta_t] + \eta^2 \cdot \mathbb{E} [\|g_s(\Delta_t)\|_2^2 \mid \Delta_t] - 2\eta \cdot \mathbb{E} [g_s(\Delta_t)^\top \Delta_t \mid \Delta_t] \\ &= \|\Delta_t\|_2^2 + \eta^2 \cdot \mathbb{E} [\|g_s(\Delta_t)\|_2^2 \mid \Delta_t] - 2\eta \cdot \nabla f(\Delta_t)^\top \Delta_t \\ &\stackrel{(i)}{\leq} \|\Delta_t\|_2^2 + \eta^2 \cdot (A \cdot \|\Delta_t\|_2^2 + B) - 2\eta \cdot \alpha \|\Delta_t\|_2^2 \\ &= (1 - 2\alpha\eta + A\eta^2) \|\Delta_t\|_2^2 + \eta^2 B. \end{aligned} \quad (18)$$

where (i) is due to assumptions (F_1) and (G_2) of Theorem 1. Taking expectations for every step $t = 1, \dots$ over the whole history, we obtain the recursion:

$$\begin{aligned} \mathbb{E} [\|\Delta_{t+1}\|_2^2] &\leq (1 - 2\alpha\eta + A\eta^2)^{t-1} \|\Delta_1\|_2^2 + \eta^2 B \cdot \sum_{i=0}^{t-1} (1 - 2\alpha\eta + A\eta^2)^i \\ &= (1 - 2\alpha\eta + A\eta^2)^{t-1} \|\Delta_1\|_2^2 + \eta^2 B \cdot \frac{1 - (1 - 2\alpha\eta + A\eta^2)^t}{2\alpha\eta - A\eta^2} \\ &\leq (1 - 2\alpha\eta + A\eta^2)^{t-1} \|\Delta_1\|_2^2 + \frac{\eta B}{2\alpha - A\eta}. \end{aligned}$$

□

The following lemma states that $\mathbb{E} [\|\Delta_t\|_2^4] = O(\eta^2)$ as $t \rightarrow \infty$ and $\eta \rightarrow 0$.

Lemma 2. For data dependent, positive constants α, A, B, C, D according to assumptions $(F_1), (G_1), (G_2)$ in Theorem 1, we have:

$$\begin{aligned} \mathbb{E}[\|\Delta_t\|_2^4] &\leq (1 - 4\alpha\eta + A(6\eta^2 + 2\eta^3) + B(3\eta + \eta^2) + C(2\eta^3 + \eta^4))^{t-1} \|\Delta_1\|_2^4 \\ &\quad + \frac{B(3\eta^2 + \eta^3) + D(2\eta^2 + \eta^3)}{4\alpha - A(6\eta + 2\eta^2) - B(3 + \eta) - C(2\eta^2 + \eta^3)}. \end{aligned} \quad (19)$$

Proof. Given Δ_t , we have the following sets of (in)equalities:

$$\begin{aligned} &\mathbb{E}[\|\Delta_{t+1}\|_2^4 \mid \Delta_t] \\ &= \mathbb{E}[\|\Delta_t - \eta g_s(\Delta_t)\|_2^4 \mid \Delta_t] \\ &= \mathbb{E}[(\|\Delta_t\|_2^2 - 2\eta \cdot g_s(\Delta_t)^\top \Delta_t + \eta^2 \|g_s(\Delta_t)\|_2^2)^2 \mid \Delta_t] \\ &= \mathbb{E}[\|\Delta_t\|_2^4 + 4\eta^2 (g_s(\Delta_t)^\top \Delta_t)^2 + \eta^4 \|g_s(\Delta_t)\|_2^4 - 4\eta \cdot g_s(\Delta_t)^\top \Delta_t \|\Delta_t\|_2^2 \\ &\quad + 2\eta^2 \cdot \|g_s(\Delta_t)\|_2^2 \|\Delta_t\|_2^2 - 4\eta^3 \cdot g_s(\Delta_t)^\top \Delta_t \|g_s(\Delta_t)\|_2^2 \mid \Delta_t] \\ &\stackrel{(i)}{\leq} \mathbb{E}[\|\Delta_t\|_2^4 + 4\eta^2 \cdot \|g_s(\Delta_t)\|_2^2 \cdot \|\Delta_t\|_2^2 + \eta^4 \|g_s(\Delta_t)\|_2^4 - 4\eta \cdot g_s(\Delta_t)^\top \Delta_t \|\Delta_t\|_2^2 \\ &\quad + 2\eta^2 \cdot \|g_s(\Delta_t)\|_2^2 \cdot \|\Delta_t\|_2^2 + 2\eta^3 \cdot (\|g_s(\Delta_t)\|_2^2 + \|\Delta_t\|_2^2) \cdot \|g_s(\Delta_t)\|_2^2 \mid \Delta_t] \\ &\stackrel{(ii)}{\leq} \mathbb{E}[\|\Delta_t\|_2^4 + (2\eta^3 + \eta^4) \|g_s(\Delta_t)\|_2^4 + (6\eta^2 + 2\eta^3) \|g_s(\Delta_t)\|_2^2 \|\Delta_t\|_2^2 \mid \Delta_t] - 4\alpha\eta \|\Delta_t\|_2^4 \\ &\stackrel{(iii)}{\leq} (1 - 4\alpha\eta) \|\Delta_t\|_2^4 + (6\eta^2 + 2\eta^3)(A\|\Delta_t\|_2^2 + B)\|\Delta_t\|_2^2 + (2\eta^3 + \eta^4)(C\|\Delta_t\|_2^4 + D) \\ &\quad = (1 - 4\alpha\eta + A(6\eta^2 + 2\eta^3) + C(2\eta^3 + \eta^4)) \|\Delta_t\|_2^4 + B(6\eta^2 + 2\eta^3) \|\Delta_t\|_2^2 + D(2\eta^3 + \eta^4) \\ &\stackrel{(iv)}{\leq} (1 - 4\alpha\eta + A(6\eta^2 + 2\eta^3) + C(2\eta^3 + \eta^4)) \cdot \|\Delta_t\|_2^4 + B(3\eta + \eta^2)(\eta^2 + \|\Delta_t\|_2^4) + D(2\eta^3 + \eta^4) \\ &\quad = (1 - 4\alpha\eta + A(6\eta^2 + 2\eta^3) + B(3\eta + \eta^2) + C(2\eta^3 + \eta^4)) \cdot \|\Delta_t\|_2^4 + B\eta^2(3\eta + \eta^2) + D(2\eta^3 + \eta^4), \end{aligned} \quad (20)$$

where (i) is due to $(g_s(\Delta_t)^\top \Delta_t)^2 \leq \|g_s(\Delta_t)\|_2^2 \cdot \|\Delta_t\|_2^2$ and $-2g_s(\Delta_t)^\top \Delta_t \leq \|g_s(\Delta_t)\|_2^2 + \|\Delta_t\|_2^2$, (ii) is due to assumptions (G_1) and (F_1) in Theorem 1, (iii) is due to assumptions (G_2) and (G_3) in Theorem 1, and (iv) is due to $2\eta\|\Delta_t\|_2^2 \leq \eta^2 + \|\Delta_t\|_2^4$. Similar to the proof of the previous lemma, applying the above rule recursively and w.r.t. the whole history of estimates, we obtain:

$$\begin{aligned} \mathbb{E}[\|\Delta_{t+1}\|_2^4] &\leq (1 - 4\alpha\eta + A(6\eta^2 + 2\eta^3) + B(3\eta + \eta^2) + C(2\eta^3 + \eta^4))^{t-1} \|\Delta_1\|_2^4 \\ &\quad + (B\eta^2(3\eta + \eta^2) + D(2\eta^3 + \eta^4)) \cdot \sum_{i=0}^{t-1} (1 - 4\alpha\eta + A(6\eta^2 + 2\eta^3) + B(3\eta + \eta^2) + C(2\eta^3 + \eta^4))^i \\ &\leq (1 - 4\alpha\eta + A(6\eta^2 + 2\eta^3) + B(3\eta + \eta^2) + C(2\eta^3 + \eta^4))^{t-1} \|\Delta_1\|_2^4 \\ &\quad + \frac{B\eta^2(3\eta + \eta^2) + D(2\eta^3 + \eta^4)}{4\alpha\eta - A(6\eta^2 + 2\eta^3) - B(3\eta + \eta^2) - C(2\eta^3 + \eta^4)}, \end{aligned}$$

which is the target inequality, after simple transformations. \square

We know that:

$$\Delta_t = \Delta_{t-1} - \eta g_s(\Delta_{t-1})$$

Using the Taylor expansion formula around the point Δ_{t-1} and using the assumption that $\hat{\theta} = 0$, we have:

$$f(\Delta_{t-1}) = f(\hat{\theta}) + \nabla f(\hat{\theta})^\top \Delta_{t-1} + \frac{1}{2} \Delta_{t-1}^\top H \Delta_{t-1} + R(\Delta_{t-1})$$

Taking further the gradient w.r.t. Δ_{t-1} in the above expression, we have:

$$\nabla f(\Delta_{t-1}) = H \Delta_{t-1} + \nabla R(\Delta_{t-1})$$

Using the identity $g_s(\Delta_{t-1}) = \nabla f(\Delta_{t-1}) + e_{t-1}$, our SGD recursion can be re-written as:

$$\Delta_t = (I - \eta H) \Delta_{t-1} - \eta (\nabla R(\Delta_{t-1}) + e_{t-1}) = (I - \eta H)^{t-1} \Delta_1 - \eta \sum_{i=1}^{t-1} (I - \eta H)^{t-1-i} (e_i + \nabla R(\Delta_i)). \quad (21)$$

For $t \geq 2$ and since: $\bar{\theta} = \bar{\theta} - \hat{\theta} = \bar{\Delta}_t = \frac{1}{t} \sum_{i=1}^t (\theta_i - \hat{\theta}) = \frac{1}{t} \sum_{i=1}^t \Delta_i$, we get:

$$\begin{aligned} t(\bar{\theta} - \hat{\theta}) &= \sum_{i=1}^t \Delta_i = \sum_{i=1}^t (I - \eta H)^{i-1} \Delta_1 - \eta \sum_{j=1}^t \sum_{i=1}^j (I - \eta H)^{j-1-i} (e_i + \nabla R(\Delta_i)) \\ &\stackrel{(i)}{=} (I - (I - \eta H)^t) \frac{H^{-1}}{\eta} \Delta_1 - \eta \sum_{j=1}^t \sum_{i=1}^j (I - \eta H)^{j-1-i} (e_i + \nabla R(\Delta_i)). \end{aligned} \quad (22)$$

where (i) holds due to the assumption that the eigenvalues of $I - \eta H$ satisfy $|\lambda_i(I - \eta H)| < 1$, and thus, the geometric series of matrices: $\sum_{k=0}^{n-1} T^k = (I - T)^{-1}(I - T^n)$, is utilized above. In our case, $T = (I - \eta H)$.

For the latter term in (22), using a variant of Abel's sum formula, we have:

$$\eta \sum_{j=1}^t \sum_{i=1}^j (I - \eta H)^{j-1-i} (e_i + \nabla R(\Delta_i)) = \eta \sum_{j=0}^{t-1} \sum_{i=1}^j (I - \eta H)^{j-i} (e_i + \nabla R(\Delta_i)) \quad (23)$$

$$\begin{aligned} &= \eta \sum_{i=1}^{t-1} \left(\sum_{j=0}^{t-i-1} (I - \eta H)^j \right) (e_i + \nabla R(\Delta_i)) \\ &= \sum_{i=1}^{t-1} (I - (I - \eta H)^{t-i}) H^{-1} (e_i + \nabla R(\Delta_i)) \\ &= H^{-1} \sum_{i=1}^{t-1} e_i + H^{-1} \sum_{i=1}^{t-1} \nabla R(\Delta_i) - H^{-1} \sum_{i=1}^{t-1} (I - \eta H)^{t-i} (e_i + \nabla R(\Delta_i)) \\ &\stackrel{(i)}{=} H^{-1} \sum_{i=1}^{t-1} e_i + H^{-1} \sum_{i=1}^{t-1} \nabla R(\Delta_i) + \frac{H^{-1}}{\eta} (I - \eta H) (\Delta_t - (I - \eta H)^{t-1} \Delta_1), \end{aligned} \quad (24)$$

where (i) follows from the fact $\sum_{i=1}^{t-1} (I - \eta H)^{t-i} (e_i + \nabla R(\Delta_i)) = (I - \eta H) \frac{1}{\eta} (\Delta_t - (I - \eta H)^{t-1} \Delta_1)$, based on the expression (21).

The above combined lead to:

$$\sqrt{t} \bar{\Delta}_t = \underbrace{\frac{1}{\sqrt{t}} (I - (I - \eta H)^t) \frac{H^{-1}}{\eta} \Delta_1}_{\varphi_1} - \underbrace{\frac{1}{\sqrt{t}} H^{-1} \sum_{i=1}^{t-1} e_i}_{\varphi_2} - \underbrace{\frac{1}{\sqrt{t}} H^{-1} \sum_{i=1}^{t-1} \nabla R(\Delta_i)}_{\varphi_3} - \underbrace{\frac{1}{\sqrt{t}} \frac{H^{-1}}{\eta} (I - \eta H) (\Delta_t - (I - \eta H)^{t-1} \Delta_1)}_{\varphi_4}. \quad (25)$$

For the main result of the theorem, we are interested in the following quantity:

$$\left\| t \mathbb{E}[(\bar{\theta}_t - \hat{\theta})(\bar{\theta}_t - \hat{\theta})^\top] - H^{-1} G H^{-1} \right\|_2 = \left\| t \mathbb{E}[\bar{\Delta}_t \bar{\Delta}_t^\top] - H^{-1} G H^{-1} \right\|_2$$

Using the φ_i notation, we have $\mathbb{E}[t \bar{\Delta}_t \bar{\Delta}_t^\top] = \mathbb{E}[(\varphi_1 + \varphi_2 + \varphi_3 + \varphi_4)(\varphi_1 + \varphi_2 + \varphi_3 + \varphi_4)^\top]$. Thus, we need to bound:

$$\begin{aligned} &\left\| t \mathbb{E}[(\bar{\theta}_t - \hat{\theta})(\bar{\theta}_t - \hat{\theta})^\top] - H^{-1} G H^{-1} \right\|_2 = \left\| \mathbb{E}[(\varphi_1 + \varphi_2 + \varphi_3 + \varphi_4)(\varphi_1 + \varphi_2 + \varphi_3 + \varphi_4)^\top] - H^{-1} G H^{-1} \right\|_2 \\ &= \left\| \mathbb{E}[\varphi_2 \varphi_2^\top] - H^{-1} G H^{-1} + \mathbb{E}[\varphi_2(\varphi_1 + \varphi_4 + \varphi_3)^\top] + (\varphi_1 + \varphi_4 + \varphi_3) \varphi_2^\top + (\varphi_1 + \varphi_4 + \varphi_3)(\varphi_1 + \varphi_4 + \varphi_3)^\top \right\|_2 \\ &\leq \left\| \mathbb{E}[\varphi_2 \varphi_2^\top] - H^{-1} G H^{-1} \right\|_2 + \left\| \mathbb{E}[\varphi_2(\varphi_1 + \varphi_4 + \varphi_3)^\top] \right\|_2 + \left\| \mathbb{E}[(\varphi_1 + \varphi_4 + \varphi_3) \varphi_2^\top] \right\|_2 \\ &\quad + \left\| \mathbb{E}[(\varphi_1 + \varphi_4 + \varphi_3)(\varphi_1 + \varphi_4 + \varphi_3)^\top] \right\|_2 \\ &\stackrel{(i)}{\lesssim} \left\| \mathbb{E}[\varphi_2 \varphi_2^\top] - H^{-1} G H^{-1} \right\|_2 + \sqrt{\mathbb{E}[\|\varphi_2\|_2^2] (\mathbb{E}[\|\varphi_1\|_2^2] + \mathbb{E}[\|\varphi_4\|_2^2] + \mathbb{E}[\|\varphi_3\|_2^2]) + \mathbb{E}[\|\varphi_1\|_2^2] + \mathbb{E}[\|\varphi_4\|_2^2] + \mathbb{E}[\|\varphi_3\|_2^2]} \end{aligned} \quad (26)$$

where (i) is due to the successive use of the AM-GM rule:

$$\|ab^\top\|_2 \leq \sqrt{\mathbb{E}[\|a\|_2^2] \mathbb{E}[\|b\|_2^2]} \leq \frac{1}{2} \mathbb{E}[\|a\|_2^2] + \mathbb{E}[\|b\|_2^2]. \quad (27)$$

for two p -dimensional random vectors a and b . Indeed, for any fixed unit vector u we have $\|\mathbb{E}[ab^\top]u\|_2 = \|\mathbb{E}[a(b^\top u)]\|_2 \leq \mathbb{E}[\|a\|_2 \|b^\top u\|_2] \leq \mathbb{E}[\|a\|_2 \|b\|_2] \leq \sqrt{\mathbb{E}[\|a\|_2^2] \mathbb{E}[\|b\|_2^2]}$. We used the fact $\|\mathbb{E}[x]\|_2 \leq \mathbb{E}[\|x\|_2]$ because $\|x\|_2$ is convex. Here also, the \lesssim hides any constants appearing from applying successively the above rule.

Therefore, to proceed bounding the quantity of interest, we need to bound the terms $\mathbb{E}[\|\varphi_i\|_2^2]$. In the statement of the theorem we have $\Delta_1 = 0$ —however similar bounds will hold if $\|\Delta_1\|_2^2 = O(\eta)$; thus, for each of the above φ_i terms we have the following.

$$\varphi_1 := \frac{1}{\sqrt{t}}(I - (I - \eta H)^t) \frac{H^{-1}}{\eta} \Delta_1 = 0, \quad (\text{due to } \Delta_1 = 0) \quad (28)$$

$$\begin{aligned} \mathbb{E}[\|\varphi_4\|_2^2] &:= \mathbb{E} \left[\left\| -\frac{1}{\sqrt{t}} \frac{H^{-1}}{\eta} (I - \eta H)(\Delta_t - (I - \eta H)^{t-1} \Delta_1) \right\|_2^2 \right] \\ &\leq \mathbb{E} \left[\|H^{-1}\|_2^2 \cdot \|I - \eta H\|_2^2 \cdot \frac{1}{\eta^2 t} \|\Delta_t\|_2^2 \right] \stackrel{(i)}{\leq} \frac{1 - \eta \lambda_U}{\lambda_L} \cdot \frac{1}{\eta^2 t} \cdot \mathbb{E}[\|\Delta_t\|_2^2] \\ &\stackrel{(ii)}{\leq} \frac{1 - \eta \lambda_U}{\lambda_L} \frac{1}{\eta^2 t} \left((1 - 2\alpha\eta + A\eta^2)^{t-1} \|\Delta_1\|_2^2 + \frac{B\eta}{2\alpha - A\eta} \right) \\ &= \frac{1 - \eta \lambda_U}{\lambda_L} \frac{B}{t\eta(2\alpha - A\eta)} \\ &= O\left(\frac{1}{t\eta}\right) \end{aligned} \quad (29)$$

where (i) is due to Assumption (F₄), (ii) is due to Lemma 1, and we used in several places the fact that $\Delta_1 = 0$.

$$\begin{aligned} \mathbb{E}[\|\varphi_3\|_2^2] &:= \mathbb{E} \left[\left\| -\frac{1}{\sqrt{t}} H^{-1} \sum_{i=1}^{t-1} \nabla R(\Delta_i) \right\|_2^2 \right] \leq \mathbb{E} \left[\frac{1}{t} \cdot \|H^{-1}\|_2^2 \cdot \left\| \sum_{i=1}^{t-1} \nabla R(\Delta_i) \right\|_2^2 \right] \stackrel{(i)}{\leq} \mathbb{E} \left[\frac{1}{\lambda_L} \frac{1}{t} \left(\sum_{i=1}^{t-1} \|\nabla R(\Delta_i)\|_2 \right)^2 \right] \\ &\stackrel{(ii)}{\leq} \mathbb{E} \left[\frac{E^2}{\lambda_L \cdot t} \left(\sum_{i=1}^{t-1} \|\Delta_i\|_2^2 \right)^2 \right] \stackrel{(iii)}{\leq} \frac{E^2}{\lambda_L \cdot t} (t-1) \cdot \mathbb{E} \left[\sum_{i=1}^{t-1} \|\Delta_i\|_2^4 \right] \\ &\leq \frac{E^2}{\lambda_L} (t-1) \sum_{i=1}^{t-1} \left((1 - 4\alpha\eta + A(6\eta^2 + 2\eta^3) + C(2\eta^3 + \eta^4))^{t-1} \|\Delta_1\|_2^4 + \frac{B(3\eta^2 + \eta^3) + D(2\eta^2 + \eta^3)}{4\alpha - A(6\eta + 2\eta^2) - C(2\eta^2 + \eta^3)} \right) \\ &\stackrel{(iv)}{=} \frac{E^2}{\lambda_L} \frac{(t-1)^2}{t} \frac{B(3\eta^2 + \eta^3) + D(2\eta^2 + \eta^3)}{4\alpha - A(6\eta + 2\eta^2) - C(2\eta^2 + \eta^3)} \stackrel{(v)}{=} O(t\eta^2). \end{aligned} \quad (30)$$

where (i) is due to Assumption (F₄) and due to $|\sum_i \chi_i|^2 \leq \sum_i |\chi_i|^2$, (ii) is due to Assumption (F₃) on bounded remainder, (iii) is due to the inequality $(\sum_{i=1}^n \chi_i^2)^2 \leq n \cdot \sum_{i=1}^n \chi_i^4$, (iv) is due to $\Delta_1 = 0$, (v) is due to η being an small constant compared to α and thus $\frac{B(3\eta^2 + \eta^3) + D(2\eta^2 + \eta^3)}{4\alpha - A(6\eta + 2\eta^2) - C(2\eta^2 + \eta^3)} = \frac{O(\eta^2)}{O(1)}$.

$$\begin{aligned} \mathbb{E}[\|\varphi_2\|_2^2] &:= \mathbb{E} \left[\left\| -\frac{1}{\sqrt{t}} H^{-1} \sum_{i=1}^{t-1} e_i \right\|_2^2 \right] \stackrel{(i)}{=} \frac{1}{t} \sum_{i=1}^{t-1} \mathbb{E}[\|H^{-1} e_i\|_2^2] \stackrel{(ii)}{\leq} \frac{\lambda_U}{t} \sum_{i=1}^{t-1} \mathbb{E}[\|e_i\|_2^2] \\ &= \frac{\lambda_U}{t} \sum_{i=1}^{t-1} \mathbb{E}[\|g_s(\Delta_i) - \nabla f(\Delta_i)\|_2^2] \leq \frac{2\lambda_U}{t} \left(\sum_{i=1}^{t-1} \mathbb{E}[\|g_s(\Delta_i)\|_2^2] + \sum_{i=1}^{t-1} \mathbb{E}[\|\nabla f(\Delta_i)\|_2^2] \right) \\ &\stackrel{(iii)}{\leq} \frac{2\lambda_U}{t} \left((t-1)B + (A + L^2) \sum_{i=1}^{t-1} \mathbb{E}[\|\Delta_i\|_2^2] \right) \\ &\stackrel{(iv)}{\leq} \frac{2\lambda_U}{t} \left((t-1)B + (A + L^2) \sum_{i=1}^{t-1} \left((1 - 2\alpha\eta + A\eta^2)^{t-1} \|\Delta_1\|_2^2 + \frac{B\eta}{2\alpha - A\eta} \right) \right) \\ &= \frac{2\lambda_U(t-1)}{t} \left(B + (A + L^2) \frac{B\eta}{2\alpha - A\eta} \right) = O(1), \end{aligned} \quad (31)$$

where (i) is due to $\mathbb{E}[(H^{-1}e_i)^\top H^{-1}e_j] = 0$ for $i \neq j$, (ii) is due to Assumption (F_4) , (iii) is due to Assumptions (F_2) and (G_2) , (iv) is due to Lemma 1.

Finally, for the term $\mathbb{E}[\varphi_2 \varphi_2^\top]$, we have

$$\mathbb{E}[\varphi_2 \varphi_2^\top] = \mathbb{E} \left[\left(-\frac{1}{\sqrt{t}} H^{-1} \sum_{i=1}^{t-1} e_i \right) \left(-\frac{1}{\sqrt{t}} H^{-1} \sum_{i=1}^{t-1} e_i \right)^\top \right] = \frac{1}{t} H^{-1} \left(\sum_{i=1}^{t-1} \mathbb{E}[e_i e_i^\top] \right) H^{-1}. \quad (32)$$

and thus:

$$\begin{aligned} \|\mathbb{E}[\varphi_2 \varphi_2^\top] - H^{-1} G H^{-1}\|_2 &= \left\| \frac{1}{t} H^{-1} \left(\sum_{i=1}^{t-1} \mathbb{E}[e_i e_i^\top] \right) H^{-1} - H^{-1} G H^{-1} \right\|_2 \\ &= \left\| \frac{1}{t} H^{-1} \left(\sum_{i=1}^{t-1} \mathbb{E}[e_i e_i^\top] - G + G \right) H^{-1} - H^{-1} G H^{-1} \right\|_2 \\ &= \left\| \frac{1}{t} H^{-1} \left(\sum_{i=1}^{t-1} \mathbb{E}[e_i e_i^\top] - G \right) H^{-1} - \frac{t-1}{t} \cdot H^{-1} G H^{-1} \right\|_2 \\ &\leq \frac{1}{t} H^{-1} \left(\sum_{i=1}^{t-1} \|\mathbb{E}[e_i e_i^\top] - G\|_2 \right) H^{-1} + \frac{t-1}{t} \|H^{-1} G H^{-1}\|_2 \end{aligned}$$

For each term $\|\mathbb{E}[e_i e_i^\top] - G\|_2, \forall i$, we have

$$\begin{aligned} \|\mathbb{E}[e_i e_i^\top] - G\|_2 &= \left\| \mathbb{E}[g_s(\Delta_i) g_s(\Delta_i)^\top] - \mathbb{E}[\nabla f(\Delta_i) \nabla f(\Delta_i)^\top] - G \right\|_2 \\ &= \left\| \mathbb{E}[(g_s(\Delta_i) - \nabla f(\Delta_i))(g_s(\Delta_i) - \nabla f(\Delta_i))^\top] - G \right\|_2 \\ &= \left\| \mathbb{E}[g_s(\Delta_i) g_s(\Delta_i)^\top] - \mathbb{E}[g_s(\Delta_i) \nabla f(\Delta_i)^\top] - \mathbb{E}[\nabla f(\Delta_i) g_s(\Delta_i)^\top] + \mathbb{E}[\nabla f(\Delta_i) \nabla f(\Delta_i)^\top] - G \right\|_2 \\ &\stackrel{(i)}{=} \left\| \mathbb{E}[g_s(\Delta_i) g_s(\Delta_i)^\top] - 2\mathbb{E}[\nabla f(\Delta_i) \nabla f(\Delta_i)^\top] + \mathbb{E}[\nabla f(\Delta_i) \nabla f(\Delta_i)^\top] - G \right\|_2 \\ &\stackrel{(ii)}{\leq} \mathbb{E}[\|\nabla f(\Delta_i)\|_2^2] + \mathbb{E}[A_1 \|\Delta_i\|_2 + A_2 \|\Delta_i\|_2^2 + A_3 \|\Delta_i\|_2^3 + A_4 \|\Delta_i\|_2^4] \\ &\stackrel{(iii)}{\leq} L^2 \mathbb{E}[\|\Delta_i\|_2^2] + A_1 \sqrt{\mathbb{E}[\|\Delta_i\|_2^2]} + A_2 \mathbb{E}[\|\Delta_i\|_2^2] + \frac{A_3}{2} \mathbb{E}[\|\Delta_i\|_2^2 + \|\Delta_i\|_2^4] + A_4 \mathbb{E}[\|\Delta_i\|_2^4] \\ &= A_1 \sqrt{\mathbb{E}[\|\Delta_i\|_2^2]} + (L^2 + A_2 + \frac{A_3}{2}) \mathbb{E}[\|\Delta_i\|_2^2] + (\frac{A_3}{2} + A_4) \mathbb{E}[\|\Delta_i\|_2^4] \\ &\stackrel{(iv)}{\leq} A_1 \sqrt{(1 - 2\alpha\eta + A\eta^2)^{t-1} \|\Delta_1\|_2^2 + \frac{B\eta}{2\alpha - A\eta}} + \left(L^2 + A_2 + \frac{A_3}{2} \right) \left((1 - 2\alpha\eta + A\eta^2)^{t-1} \|\Delta_1\|_2^2 + \frac{B\eta}{2\alpha - A\eta} \right) \\ &\quad + \left(\frac{A_3}{2} + A_4 \right) \left((1 - 4\alpha\eta + A(6\eta^2 + 2\eta^3) + C(2\eta^3 + \eta^4))^{t-1} \|\Delta_1\|_2^4 + \frac{B(3\eta^2 + \eta^3) + D(2\eta^2 + \eta^3)}{4\alpha - A(6\eta + 2\eta^2) - C(2\eta^2 + \eta^3)} \right) \\ &= A_1 \sqrt{\frac{B\eta}{2\alpha - A\eta}} + \left(L^2 + A_2 + \frac{A_3}{2} \right) \frac{B\eta}{2\alpha - A\eta} + \left(\frac{A_3}{2} + A_4 \right) \frac{B(3\eta^2 + \eta^3) + D(2\eta^2 + \eta^3)}{4\alpha - A(6\eta + 2\eta^2) - C(2\eta^2 + \eta^3)}. \quad (33) \end{aligned}$$

where (i) is due to Assumption (G_1) , (ii) is due to Cauchy-Schwartz inequality and Assumption (G_4) , (iii) is due to Assumption (F_2) , (iv) is due to Lemmas 1-2.

Then, we have:

$$\begin{aligned} &\left\| \mathbb{E}[\varphi_2 \varphi_2^\top] - H^{-1} G H^{-1} \right\|_2 \\ &\stackrel{(i)}{\leq} \frac{t-1}{t} \|H^{-1} G H^{-1}\|_2 + \frac{t-1}{\lambda_L^2 t} \left(A_1 \sqrt{\frac{B\eta}{2\alpha - A\eta}} + (L^2 + A_2 + \frac{A_3}{2}) \frac{B\eta}{2\alpha - A\eta} + (\frac{A_3}{2} + A_4) \frac{B(3\eta^2 + \eta^3) + D(2\eta^2 + \eta^3)}{4\alpha - A(6\eta + 2\eta^2) - C(2\eta^2 + \eta^3)} \right) \\ &= O(\sqrt{\eta}). \quad (34) \end{aligned}$$

where (i) is due to Assumption, and (ii) is after removing constants and observing that the dominant term in the second part is $O(\sqrt{\eta})$. Combining all the above in (26), we obtain:

$$\left\| t\mathbb{E}[(\bar{\theta}_t - \hat{\theta})(\bar{\theta}_t - \hat{\theta})^\top] - H^{-1} G H^{-1} \right\|_2 \lesssim \sqrt{\eta} + \sqrt{\frac{1}{t\eta} + t\eta^2}.$$

■

C.2 Proof of Corollary 2

Proof of Corollary 2. Here we use the same notations as the proof of Theorem 1. Because linear regression satisfies $\nabla f(\theta) - H(\theta - \hat{\theta}) = 0$, we do not have to consider the Taylor remainder term in our analysis. And we do not need 4-th order bound for SGD. Due to the fact that the quadratic function is strongly convex, we have $\Delta^\top \nabla f(\Delta + \hat{\theta}) \geq \lambda_L \|\Delta\|_2^2$.

By sampling with replacement, we have

$$\begin{aligned}
\mathbb{E}[\|g_s(\theta_t)\|_2^2 \mid \theta_t] &= \|\nabla f(\theta_t)\|_2^2 + \mathbb{E}[\|e_t\|_2^2 \mid \theta_t] \\
&= \|\nabla f(\theta_t)\|_2^2 + \frac{1}{S} \left(\frac{1}{n} \sum \|\nabla f_i(\theta_t)\|_2^2 - \|\nabla f(\theta_t)\|_2^2 \right) \\
&\leq L^2 \left(1 - \frac{1}{S}\right) \|\Delta_t\|_2^2 + \frac{1}{S} \frac{1}{n} \sum \|x_i(x_i^\top \theta_t - y_i)\|_2^2 \\
&= L^2 \left(1 - \frac{1}{S}\right) \|\Delta_t\|_2^2 + \frac{1}{S} \frac{1}{n} \sum \|x_i x_i^\top \Delta_t + x_i x_i^\top \hat{\theta} - y_i x_i\|_2^2 \\
&\leq L^2 \left(1 - \frac{1}{S}\right) \|\Delta_t\|_2^2 + 2 \frac{1}{S} \frac{1}{n} \sum (\|x_i x_i^\top \Delta_t\|_2^2 + \|x_i x_i^\top \hat{\theta} - y_i x_i\|_2^2) \\
&\leq \left(L^2 \left(1 - \frac{1}{S}\right) + 2 \frac{1}{S} \frac{1}{n} \sum \|x_i\|_2^4 \right) \|\Delta_t\|_2^2 + 2 \frac{1}{S} \frac{1}{n} \sum \|x_i x_i^\top \hat{\theta} - y_i x_i\|_2^2.
\end{aligned} \tag{35}$$

We also have

$$\begin{aligned}
\|\mathbb{E}[g_s(\theta)g_s(\theta)^\top \mid \theta] - G\|_2 &= \left\| \frac{1}{S} \frac{1}{n} \sum \nabla f_i(\theta) f_i(\theta)^\top - \nabla f(\theta) \nabla f(\theta)^\top - G \right\|_2 \\
&\leq \|\nabla f(\theta)\|_2^2 + \frac{1}{S} \left\| \frac{1}{n} \sum \nabla f_i(\theta) f_i(\theta)^\top - G \right\|_2 \\
&\leq \|\nabla f(\theta)\|_2^2 + \frac{1}{S} \left\| \frac{1}{n} \sum (g_i + H_i \Delta)(g_i + H_i \Delta)^\top - G \right\|_2 \\
&\leq \|\nabla f(\theta)\|_2^2 + \frac{1}{S} \left\| \frac{1}{n} \sum H_i \Delta g_i^\top + g_i \Delta^\top H_i + H_i \Delta \Delta^\top H_i \right\|_2 \\
&\leq \|\nabla f(\theta)\|_2^2 + \frac{1}{S} \left(\frac{2}{n} \|H_i\|_2 \|g_i\|_2 \right) \|\Delta\|_2 + \frac{1}{S} \left(\frac{1}{n} \sum \|H_i\|_2^2 \right) \|\Delta\|_2^2 \\
&\leq \frac{1}{S} \left(\frac{2}{n} \|H_i\|_2 \|g_i\|_2 \right) \|\Delta\|_2 + \left(L^2 + \frac{1}{S} \frac{1}{n} \sum \|H_i\|_2^2 \right) \|\Delta\|_2^2,
\end{aligned} \tag{36}$$

where $g_i = x_i(x_i^\top \hat{\theta} - y_i)$ and $H_i = x_i x_i^\top$.

Following Theorem 1's proof, we have

$$\left\| t \mathbb{E}[(\bar{\theta}_t - \hat{\theta})(\bar{\theta}_t - \hat{\theta})^\top] - H^{-1} G H^{-1} \right\|_2 \lesssim \sqrt{\eta} + \frac{1}{\sqrt{t\eta}}. \tag{37}$$

□

C.3 Proof of Corollary 1

Proof of Corollary 1. Here we use the same notations as the proof of Theorem 1. Because $\nabla^2 f(\theta) = \nabla k(\theta) \nabla k(\theta)^\top + (k(\theta) + c) \nabla^2 k(\theta)$, $f(\theta)$ is convex. The following lemma shows that $\nabla f(\theta) = (k(\theta) + c) \nabla k(\theta)$ is Lipschitz.

Lemma 3.

$$\|\nabla f(\theta)\|_2 \leq L \|\Delta\|_2 \tag{38}$$

for some data dependent constant L .

Proof. First, because

$$\nabla k(\theta) = \frac{1}{n} \sum -\frac{-y_i x_i}{1 + \exp(y_i \theta^\top x_i)}, \tag{39}$$

we have

$$\|\nabla k(\theta)\|_2 \leq \frac{1}{n} \sum \|x_i\|_2. \tag{40}$$

Also, we have

$$\|\nabla^2 k(\theta)\|_2 = \left\| \frac{1}{n} \sum \frac{\exp(y_i \theta^\top x_i)}{(1 + \exp(y_i \theta^\top x_i))^2} x_i x_i^\top \right\|_2 \leq \frac{1}{4n} \sum \|x_i\|_2^2, \tag{41}$$

which implies

$$\|\nabla k(\theta)\|_2 \leq \frac{1}{4n} \sum \|x_i\|_2^2 \|\Delta\|_2. \quad (42)$$

Further:

$$\begin{aligned} k(\theta) &= \frac{1}{n} \sum \log(1 + \exp(-y_i \Delta^\top x_i - y_i \hat{\theta}^\top x_i)) \\ &\leq \frac{1}{n} \sum \log(1 + \exp(\|x_i\|_2 \|\Delta\|_2 - y_i \hat{\theta}^\top x_i)) \\ &\stackrel{(i)}{\leq} \frac{1}{n} \sum (\log(1 + \exp(-y_i \hat{\theta}^\top x_i)) + \|x_i\|_2 \|\Delta\|_2) \end{aligned} \quad (43)$$

where step (i) follows from $\log(1 + \exp(a + b)) \leq \log(1 + e^b) + |a|$. Thus, we have

$$\begin{aligned} \|\nabla f(\theta)\|_2 &= \|(k(\theta) + c)\nabla k(\theta)\|_2 \leq k(\theta) \|\nabla k(\theta)\|_2 + c \|\nabla k(\theta)\|_2 \\ &\leq \left(c + \frac{1}{n} \sum \log(1 + \exp(-y_i \hat{\theta}^\top x_i))\right) \|\nabla k(\theta)\|_2 + \left(\frac{1}{n} \sum \|x_i\|_2\right)^2 \|\Delta\|_2, \end{aligned} \quad (44)$$

and we can conclude that $\|\nabla f(\theta)\|_2 \leq L \|\Delta\|_2$ for some data dependent constant L . \square

Next, we show that $f(\theta)$ has a bounded Taylor remainder.

Lemma 4.

$$\|\nabla f(\theta) - H(\theta - \hat{\theta})\|_2 \leq E \|\theta - \hat{\theta}\|_2^2, \quad (45)$$

for some data dependent constant E .

Proof. Because $\nabla f(\theta) = (k(\theta) + c)\nabla k(\theta)$, we know that $\|\nabla f(\theta)\|_2 = O(\|\Delta\|_2)$ when $\|\Delta\|_2 = \Omega(1)$ where the constants are data dependent. Because $f(\theta)$ is infinitely differentiable, by the Taylor expansion we know that $\|\nabla f(\theta) - H(\theta - \hat{\theta})\|_2 = O(\|\theta - \hat{\theta}\|_2^2)$ when $\|\Delta\|_2 = O(1)$ where the constants are data dependent. Combining the above, we can conclude $\|\nabla f(\theta) - H(\theta - \hat{\theta})\|_2 \leq E \|\theta - \hat{\theta}\|_2^2$ for some data dependent constant E . \square

In the following lemma, we will show that $\nabla f(\theta)^\top (\theta - \hat{\theta}) \geq \alpha \|\theta - \hat{\theta}\|_2^2$ for some data dependent constant α .

Lemma 5.

$$\nabla f(\theta)^\top (\theta - \hat{\theta}) \geq \alpha \|\theta - \hat{\theta}\|_2^2, \quad (46)$$

for some data dependent constant α .

Proof. We know that

$$\nabla f(\theta)^\top \Delta = (k(\theta) + c)\nabla k(\theta)^\top \Delta. \quad (47)$$

First, notice that locally (when $\|\Delta\|_2 = O(\frac{\lambda_L}{E})$) we have

$$\nabla k(\theta)^\top \Delta \gtrsim \Delta^\top H \Delta \gtrsim \lambda_L \|\Delta\|_2^2, \quad (48)$$

because of the optimality condition. This lower bounds $\nabla f(\theta)^\top (\theta - \hat{\theta})$ when $\|\Delta\|_2 = O(\frac{\lambda_L}{E})$. Next we will lower bound it when $\|\Delta\|_2 = \Omega(\frac{\lambda_L}{E})$.

Consider the function for $t \in [0, \infty)$, we have

$$\begin{aligned} g(t) &= \nabla f(\hat{\theta} + ut)^\top ut \\ &= (k(\hat{\theta} + ut) + c)\nabla k(\hat{\theta} + ut)^\top ut \\ &= k(\hat{\theta} + ut)\nabla k(\hat{\theta} + ut)^\top ut + c\nabla k(\hat{\theta} + ut)^\top ut, \end{aligned} \quad (49)$$

where $u = \frac{\Delta}{\|\Delta\|_2}$. Because $k(\theta)$ is convex, $\nabla k(\hat{\theta} + ut)^\top u$ is an increasing function in t , thus we have $\nabla k(\hat{\theta} + ut)^\top u = \Omega(\frac{\lambda_L^2}{E})$ when $t = \Omega(\frac{\lambda_L}{E})$. And we can deduce $\nabla k(\hat{\theta} + ut)^\top ut = \Omega(\frac{\lambda_L^2}{E} t)$ when $t = \Omega(\frac{\lambda_L}{E})$.

Similarly, because $k(\theta)$ is convex, $k(\hat{\theta} + ut)$ is an increasing function in t . Its derivative $\nabla k(\hat{\theta} + ut)^\top u = \Omega(\frac{\lambda_L^2}{E})$ when $t = \Omega(\frac{\lambda_L}{E})$. So we have $k(\hat{\theta} + ut) = \Omega(\frac{\lambda_L^2}{E} t)$ when $t = \Omega(\frac{\lambda_L}{E})$.

Thus, we have

$$k(\hat{\theta} + ut)\nabla k(\hat{\theta} + ut)^\top ut = \Omega\left(\frac{\lambda_L^4}{E^2}t^2\right), \quad (50)$$

when $t = \Omega(\frac{E}{\lambda_L})$.

And we can conclude that $\nabla f(\theta)^\top (\theta - \hat{\theta}) \geq \alpha \|\theta - \hat{\theta}\|_2^2$ for some data dependent constant $\alpha = \Omega(\min\{\lambda_L, \frac{\lambda_L^4}{E^2}\})$. \square

Next, we will prove properties about $g_s = \Psi_s \Upsilon_s$.

$$\mathbb{E}[\|\Upsilon\|_2^2 \mid \theta] = \frac{1}{S_\Upsilon} \left(\frac{1}{n} \sum \|\nabla k_i(\theta)\|_2^2 - \|\nabla k(\theta)\|_2^2 \right) + \|\nabla k(\theta)\|_2^2 \lesssim \frac{1}{n} \|x_i\|_2^2 \quad (51)$$

$$\begin{aligned} \mathbb{E}[\Psi_s^2] &\stackrel{(i)}{\leq} \frac{1}{n} \sum (c + k_i(\theta))^2 \\ &= \frac{1}{n} \sum \left(c + \log(1 + \exp(-y_i \hat{\theta}^\top x_i - y_i \Delta x_i)) \right)^2 \\ &\stackrel{(ii)}{\lesssim} \frac{1}{n} \sum \|x_i\|_2^2 \|\Delta\|_2^2 + \frac{1}{n} \sum (c + \log(1 + \exp(-y_i \hat{\theta}^\top x_i)))^2, \end{aligned} \quad (52)$$

where (i) follows from $\mathbb{E} \left[\left(\frac{\sum_{j=1}^S X_j}{S} \right)^2 \right] \leq \mathbb{E} \left[\frac{\sum_{j=1}^S X_j^2}{S} \right]$ and (ii) follows from $\log(1 + \exp(a + b)) \leq \log(1 + e^b) + |a|$.

Thus, we have

$$\mathbb{E}[\|g_s\|_2^2(\theta) \mid \theta] = \mathbb{E}[\Psi^2 \mid \theta] \cdot \mathbb{E}[\|\Upsilon\|_2^2 \mid \theta] \lesssim A \|\Delta\|_2^2 + B \quad (53)$$

for some data dependent constants A and B .

For the fourth-moment quantities, we have:

$$\begin{aligned} \mathbb{E}[\|\Upsilon\|_2^4 \mid \theta] &= \mathbb{E} \left[\left\| \frac{1}{S_\Upsilon} \sum_{i \in I_t^\Upsilon} \nabla \log(1 + \exp(-y_i \theta^\top x_i)) \right\|_2^4 \right] \\ &\leq \mathbb{E} \left[\left(\frac{1}{S_\Upsilon} \sum_{i \in I_t^\Upsilon} \|\nabla \log(1 + \exp(-y_i \theta^\top x_i))\|_2 \right)^4 \right] \\ &\leq \mathbb{E} \left[\left(\frac{1}{S_\Upsilon} \sum_{i \in I_t^\Upsilon} \|x_i\|_2 \right)^4 \right] \leq \frac{1}{n} \sum \|x_i\|_2^4. \end{aligned} \quad (54)$$

$$\begin{aligned} \mathbb{E}[\Psi_s^4] &\stackrel{(i)}{\leq} \frac{1}{n} \sum (c + k_i(\theta))^4 = \frac{1}{n} \sum \left(c + \log(1 + \exp(-y_i \hat{\theta}^\top x_i - y_i \Delta x_i)) \right)^4 \\ &\stackrel{(ii)}{\lesssim} \frac{1}{n} \sum \|x_i\|_2^4 \|\Delta\|_2^4 + \frac{1}{n} \sum \left(c + \log(1 + \exp(-y_i \hat{\theta}^\top x_i)) \right)^4, \end{aligned} \quad (55)$$

‘where (i) follows from $\mathbb{E} \left[\left(\frac{\sum_{j=1}^S X_j}{S} \right)^4 \right] \leq \mathbb{E} \left[\frac{\sum_{j=1}^S X_j^4}{S} \right]$ and (ii) follows from $\log(1 + \exp(a + b)) \leq \log(1 + e^b) + |a|$.

Combining the above, we get:

$$\mathbb{E}[\|g_s\|_2^4(\theta) \mid \theta] = \mathbb{E}[\Psi^4 \mid \theta] \cdot \mathbb{E}[\|\Upsilon\|_2^4 \mid \theta] \lesssim C \|\Delta\|_2^4 + D, \quad (56)$$

for some data dependent constants C and D .

Finally, we need a bound for the quantity $\|\mathbb{E}[\nabla g_s(\theta)\nabla g_s(\theta)^\top] - G\|_2$. We observe:

$$\begin{aligned} \|\mathbb{E}[\nabla g_s(\theta)\nabla g_s(\theta)^\top] - G\|_2 &\leq \left\| \frac{K_G(\theta)}{n} \sum \nabla k_i(\theta)\nabla k_i(\theta)^\top - \frac{K_G(\hat{\theta})}{n} \sum \nabla k_i(\hat{\theta})\nabla k_i(\hat{\theta})^\top \right\|_2 \\ &\leq \left\| \frac{K_G(\theta)}{n} \sum \nabla k_i(\theta)\nabla k_i(\theta)^\top - \frac{K_G(\theta)}{n} \sum \nabla k_i(\hat{\theta})\nabla k_i(\hat{\theta})^\top \right. \\ &\quad \left. + \frac{K_G(\theta)}{n} \sum \nabla k_i(\hat{\theta})\nabla k_i(\hat{\theta})^\top - \frac{K_G(\hat{\theta})}{n} \sum \nabla k_i(\hat{\theta})\nabla k_i(\hat{\theta})^\top \right\|_2 \\ &\leq \frac{K_G(\theta)}{n} \left\| \sum \left(\nabla k_i(\theta)\nabla k_i(\theta)^\top - \nabla k_i(\hat{\theta})\nabla k_i(\hat{\theta})^\top \right) \right\|_2 \\ &\quad + |K_G(\theta) - K_G(\hat{\theta})| \cdot \left\| \frac{1}{n} \sum \nabla k_i(\hat{\theta})\nabla k_i(\hat{\theta})^\top \right\|_2. \end{aligned} \quad (57)$$

Because

$$K_G(\theta) = O(1 + \|\Delta\|_2 + \|\Delta\|_2^2), \quad (58)$$

$$\frac{1}{n} \left\| \sum (\nabla k_i(\theta)\nabla k_i(\theta)^\top - \nabla k_i(\hat{\theta})\nabla k_i(\hat{\theta})^\top) \right\|_2 = O(\|\Delta\|_2 + \|\Delta\|_2^2), \quad (59)$$

$$|K_G(\theta) - K_G(\hat{\theta})| = O(\|\Delta\|_2 + \|\Delta\|_2^2), \quad (60)$$

we may conclude that

$$\|\mathbb{E}[g_s(\theta)g_s(\theta)^\top \mid \theta] - G\|_2 \leq A_1\|\theta - \hat{\theta}\|_2 + A_2\|\theta - \hat{\theta}\|_2^2 + A_3\|\theta - \hat{\theta}\|_2^3 + A_4\|\theta - \hat{\theta}\|_2^4, \quad (61)$$

for some data dependent constants A_1, A_2, A_3 , and A_4 .

Combining above results and using Theorem 1, we have

$$\left\| t\mathbb{E}[(\bar{\theta}_t - \hat{\theta})(\bar{\theta}_t - \hat{\theta})^\top] - H^{-1}GH^{-1} \right\|_2 \lesssim \sqrt{\eta} + \sqrt{\frac{1}{t\eta} + t\eta^2}. \quad (62)$$

□

D Experiments

Here we present additional experiments on our SGD inference procedure.

D.1 Synthetic data




Figure 5: Estimation in univariate models: Q-Q plots for samples shown in Figure 2

Figure 5 shows Q-Q plots for samples shown in Figure 2.

Multivariate models Here we show Q-Q plots per coordinate for samples from our SGD inference procedure.

Q-Q plots per coordinate for samples in linear regression experiment 1 is shown in Figure 6. Q-Q plots per coordinate for samples in linear regression experiment 2 is shown in Figure 7.

Q-Q plots per coordinate for samples in logistic regression experiment 1 is shown in Figure 8. Q-Q plots per coordinate for samples in logistic regression experiment 2 is shown in Figure 9.

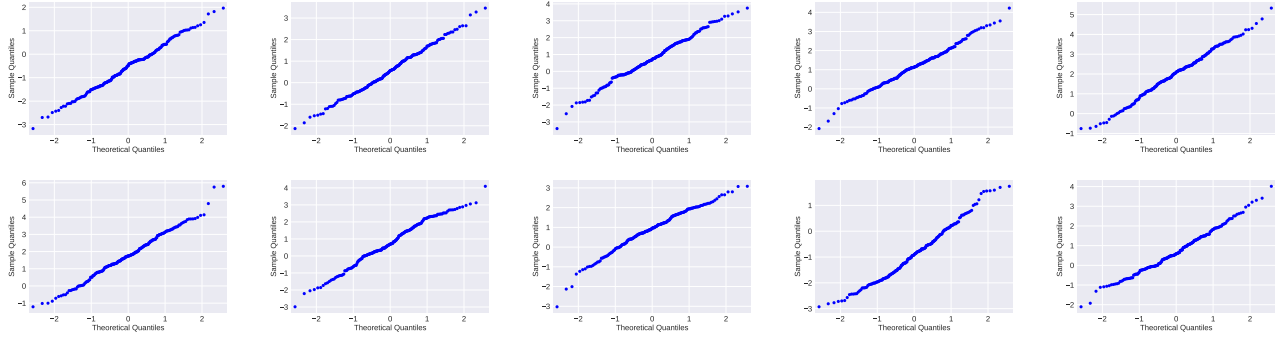


Figure 6: Linear regression experiment 1: Q-Q plots per coordinate

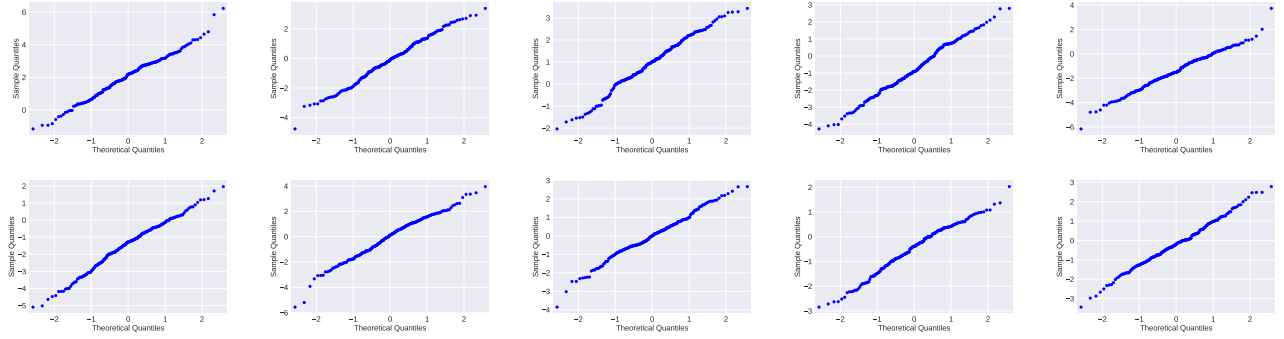


Figure 7: Linear regression experiment 2: Q-Q plots per coordinate

Additional experiments

2-Dimensional Linear Regression. Consider:

$$y = x_1 + x_2 + \epsilon, \quad \text{where } \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \sim \mathcal{N}\left(0, \begin{bmatrix} 1 & 0.8 \\ 0.8 & 1 \end{bmatrix}\right) \text{ and } \epsilon \sim \mathcal{N}(0, \sigma^2 = 10^2).$$

Each sample consists of $Y = y$ and $X = [x_1, x_2]^\top$. We use linear regression to estimate w_1, w_2 in $y = w_1x_1 + w_2x_2$. In this case, the minimizer of the population least square risk is $w_1^* = 1, w_2^* = 1$.

For 300 i.i.d. samples, we plotted 100 samples from SGD inference in Figure 10. We compare our SGD inference procedure against bootstrap in Figure 10a. Figure 10b and Figure 10c show samples from our SGD inference procedure with different parameters.

10-Dimensional Linear Regression.

Here we consider the following model

$$y = x^\top w^* + \epsilon,$$

where $w^* = \frac{1}{\sqrt{10}}[1, 1, \dots, 1]^\top \in \mathbb{R}^{10}$, $x \sim \mathcal{N}(0, \Sigma)$ with $\Sigma_{ij} = 0.8^{|i-j|}$, and $\epsilon \sim \mathcal{N}(0, \sigma^2 = 20^2)$, and use $n = 1000$ samples. We estimate the parameter using

$$\hat{w} = \underset{w}{\operatorname{argmin}} \frac{1}{n} \sum_{i=1}^n \frac{1}{2} (x_i^\top w - y_i)^2.$$

Figure 11 shows the diagonal terms of the covariance matrix computed using the sandwich estimator and our SGD inference procedure with different parameters. 100000 samples from our SGD inference procedure are used to reduce the effect of randomness.

2-Dimensional Logistic Regression.

Here we consider the following model

$$\mathbb{P}[Y = +1] = \mathbb{P}[Y = -1] = \frac{1}{2}, \quad X | Y \sim \mathcal{N}(\mu = 1.1 + 0.1Y, \sigma^2 = 1). \quad (63)$$

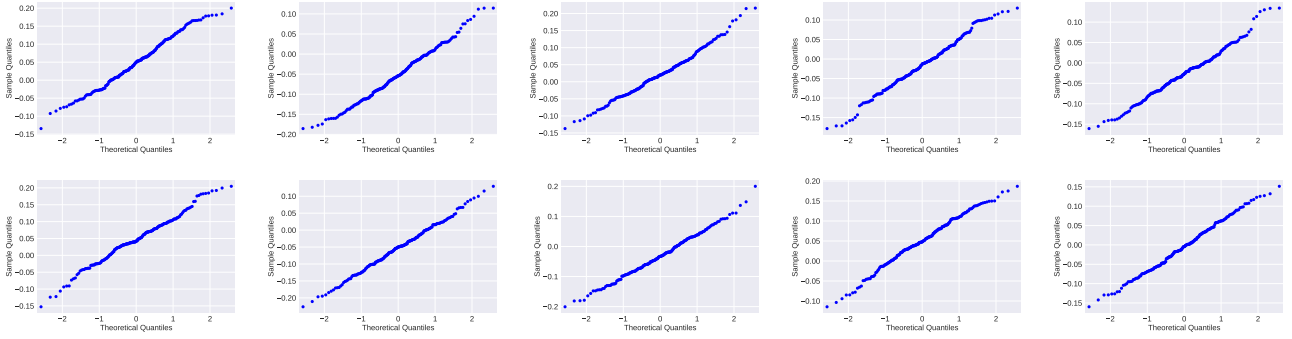


Figure 8: Logistic regression experiment 1: Q-Q plots per coordinate

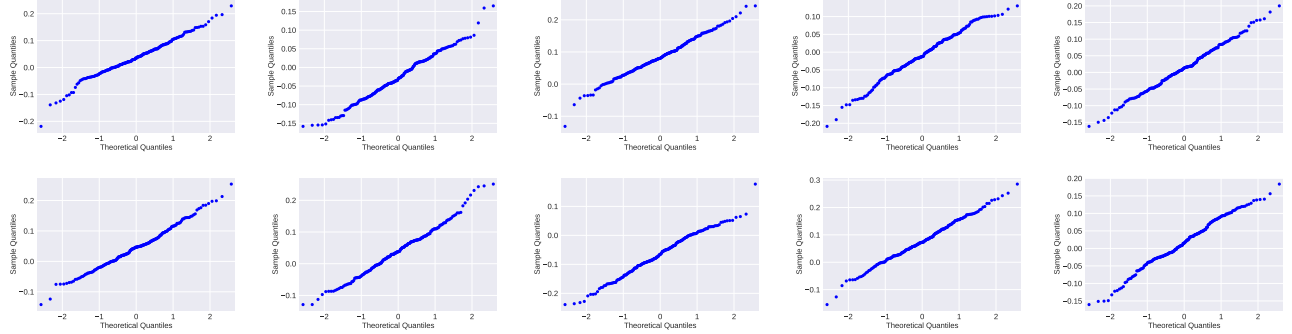


Figure 9: Logistic regression experiment 2: Q-Q plots per coordinate


We use logistic regression to estimate w, b in the classifier $\text{sign}(wx + b)$ where the minimizer of the population logistic risk is $w^* = 0.2, b^* = -0.22$.

For 100 i.i.d. samples, we plot 1000 samples from SGD in Figure 12. In our simulations, we notice that our modified SGD for logistic regression behaves similar to vanilla logistic regression. This suggests that an assumption weaker than $(\theta - \hat{\theta})^\top \nabla f(\theta) \geq \alpha \|\theta - \hat{\theta}\|_2^2$ (assumption (F_1) in Theorem 1) is sufficient for SGD analysis. Figure 12b and Figure 12d suggest that the $t\eta^2$ term in Corollary 1 is an artifact of our analysis, and can be improved.


11-Dimensional Logistic Regression.

Here we consider the following model


$$\mathbb{P}[Y = +1] = \mathbb{P}[Y = -1] = \frac{1}{2}, \quad X | Y \sim \mathcal{N}(0.01Y\mu, \Sigma),$$



(a) SGD inference vs. bootstrap



(b) $t = 800$



(c) $\eta = 0.1$

Figure 10: 2-dimensional linear regression

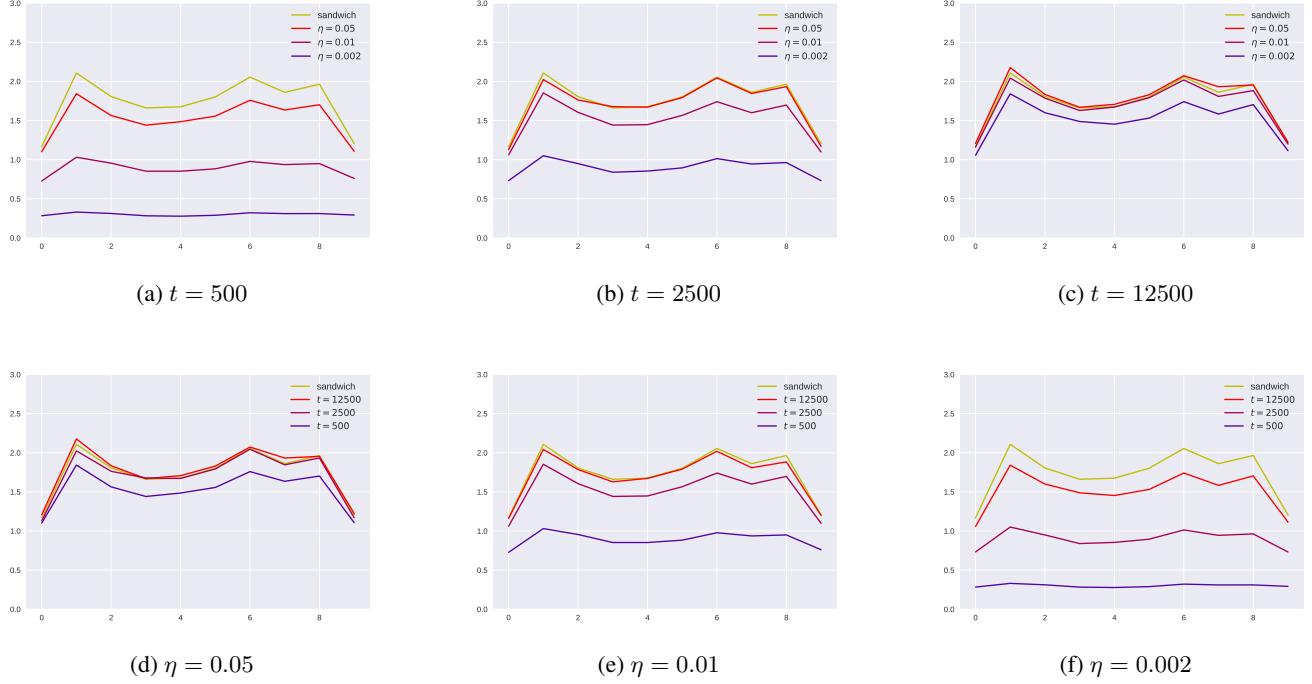


Figure 11: 11-dimensional linear regression: covariance matrix diagonal terms of SGD inference and sandwich estimator

where $\Sigma_{ii} = 1$ and when $i \neq j$ $\Sigma_{ij} = \rho^{|i-j|}$ for some $\rho \in [0, 1)$, and $\mu = \frac{1}{\sqrt{10}}[1, 1, \dots, 1]^\top \in \mathbb{R}^{10}$. We estimate a classifier $\text{sign}(w^\top x + b)$ using

$$\hat{w}, \hat{b} = \underset{w, b}{\operatorname{argmin}} \frac{1}{n} \sum_{i=1}^n \log(1 + \exp(-Y_i(w^\top X_i + b))). \quad (64)$$

Figure 13 shows results for $\rho = 0$ with $n = 80$ samples. We use $t = 100$, $d = 70$, $\eta = 0.8$, and mini batch of size 4 in vanilla SGD. Bootstrap and our SGD inference procedure each generated 2000 samples. In bootstrap, we used Newton method to perform optimization over each replicate, and 6-7 iterations were used. In figure 14, we follow the same procedure for $\rho = 0.6$ with $n = 80$ samples. Here, we use $t = 200$, $d = 70$, $\eta = 0.85$; the rest of the setting is the same.

D.2 Real data

Here, we compare covariance matrix computed using our SGD inference procedure, bootstrap, and inverse Fisher information matrix on the Higgs data set (Baldi, Sadowski, and Whiteson, 2014) and the LIBSVM Splice data set, and we observe that they have similar statistical properties.


Higgs data set The Higgs data set ⁴ (Baldi, Sadowski, and Whiteson, 2014) contains 28 distinct features with 11,000,000 data samples. This is a classification problem between two types of physical processes: one produces Higgs bosons and the other is a background process that does not. We use a logistic regression model, trained using vanilla SGD, instead of the modified SGD described in Section 3.3.

To understand different settings of sample size, we subsampled the data set with different sample size levels: $n = 200$ and $n = 50000$. We investigate the empirical performance of SGD inference on this subsampled data set. In all experiments below, the batch size of the mini batch SGD is 10.

In the case $n = 200$, the asymptotic normality for the MLE is not a good enough approximation. Hence, in this small-sample inference, we compare the SGD inference covariance matrix with the one obtained by inverse Fisher information matrix and bootstrap in Figure 15.

For our SGD inference procedure, we use $t = 100$ samples to average, and discard $d = 50$ samples. We use $R = 20$ averages from 20 segments (as in Figure 1). For bootstrap, we use 2000 replicates, which is much larger than the sample size $n = 200$.


⁴<https://archive.ics.uci.edu/ml/datasets/HIGGS>




(a) Modified SGD with $t = 1000$ and $\eta = 0.1$




(b) Modified SGD with $\eta = 0.1$



(c) Vanilla SGD with $t = 1000$ and $\eta = 0.1$



(d) Vanilla SGD with $\eta = 0.1$



(e) $t = 1000$ and $\eta = 0.1$

Figure 12: 2-dimensional logistic regression


Figure 15 shows that the covariance matrix obtained by SGD inference is comparable to the estimation given by bootstrap and inverse Fisher information.

In the case $n = 50000$, we use $t = 5000$ samples to average, and discard $d = 500$ samples. We use $R = 20$ averages from 20 segments (as in Figure 1). For this large data set, we present the estimated covariance of SGD inference procedure and inverse Fisher information (the asymptotic covariance) in Figure 16 because bootstrap is computationally prohibitive. Similar to the small sample result in Figure 15, the covariance of our SGD inference procedure is comparable to the inverse Fisher information.


In Figure 17, we compare the covariance matrix computed using our SGD inference procedure and inverse Fisher information with $n = 90000$ samples. We used 25 samples from our SGD inference procedure with $t = 5000$, $d = 1000$, $\eta = 0.2$, and mini batch size of 10.

Splice data set The Splice data set⁵ contains 60 distinct features with 1000 data samples. This is a classification problem between two classes of splice junctions in a DNA sequence. Similar to the Higgs data set, we use a logistic regression model, trained using vanilla SGD.


⁵<https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/binary.html>



(a) SGD covariance



(b) Bootstrap estimated covariance



(c) Diagonal terms

Figure 13: 11-dimensional logistic regression: $\rho = 0$

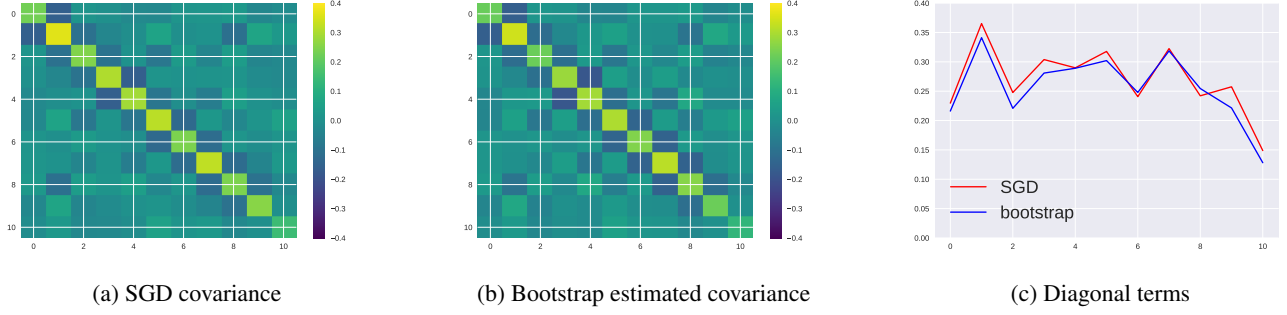


Figure 14: 11-dimensional logistic regression: $\rho = 0.6$






Figure 15: Higgs data set with $n = 200$

In Figure 18, we compare the covariance matrix computed using our SGD inference procedure and bootstrap $n = 1000$ samples. We used 10000 samples from both bootstrap and our SGD inference procedure with $t = 500$, $d = 100$, $\eta = 0.2$, and mini batch size of 6.

MNIST Here, we train a binary logistic regression classifier to classify 0/1 using perturbed MNIST data set, and demonstrate that certain adversarial examples (e.g. (Goodfellow, Shlens, and Szegedy, 2015)) can be detected using prediction confidence intervals. For each image, where each original pixel is either 0 or 1, we randomly changed 70% pixels to random numbers uniformly on $[0, 0.9]$. Figure 19 shows each image's logit value ($\log \frac{\mathbb{P}[1|\text{image}]}{\mathbb{P}[0|\text{image}]}$) and its 95% confidence interval computed using our SGD inference procedure. The adversarial perturbation used here is shown in Figure 20 (scaled for display).




(a) Inverse Fisher information




(b) SGD inference covariance


Figure 16: Higgs data set with $n = 50000$



(a) Inverse Fisher information




(b) SGD inference covariance




(c) Diagonal terms


Figure 17: Higgs data set with $n = 90000$



(a) Bootstrap




(b) SGD inference covariance




(c) Diagonal terms


Figure 18: Splice data set




(a) Original “0”: logit -72.6,
CI (-106.4, -30.0)




(b) Adversarial “0”: logit 15.3,
CI (-6.5, 26.2)




(c) Original “0”: logit -62.1,
CI (-101.6, -5.5)




(d) Adversarial “0”: logit 1.9,
CI (-4.9, 11.6)




(e) Original “0”: logit -42.9,
CI (-75.4, 5.1)



(f) Adversarial “0”: logit 4.8,
CI (-3.4, 17.9)



(g) Original “0”: logit -77.0,
CI (-110.7, -32.2)



(h) Adversarial “0”: logit 13.3,
CI (-8.0, 25.7)

Figure 19: MNIST




Figure 20: MNIST adversarial perturbation (scaled for display)