

EE-Grad: Exploration and Exploitation for Cost-Efficient Mini-Batch SGD*

Mehmet A. Donmez[†]

Maxim Raginsky[‡]

Andrew C. Singer[§]

May 22, 2017

Abstract

We present a generic framework for trading off fidelity and cost in computing stochastic gradients when the costs of acquiring stochastic gradients of different quality are not known a priori. We consider a mini-batch oracle that distributes a limited query budget over a number of stochastic gradients and aggregates them to estimate the true gradient. Since the optimal mini-batch size depends on the unknown cost-fidelity function, we propose an algorithm, *EE-Grad*, that sequentially explores the performance of mini-batch oracles and exploits the accumulated knowledge to estimate the one achieving the best performance in terms of cost-efficiency. We provide performance guarantees for EE-Grad with respect to the optimal mini-batch oracle, and illustrate these results in the case of strongly convex objectives. We also provide a simple numerical example that corroborates our theoretical findings.

1 Introduction

Stochastic gradient methods are widely used to solve large-scale optimization problems in machine learning. Given a differentiable objective function $F : \mathbb{R}^d \rightarrow \mathbb{R}$ with a gradient ∇F , a stochastic gradient descent (SGD) algorithm chooses an initial iterate $\mathbf{w}_1 \in \mathbb{R}^d$, and, on each iteration $k = 1, \dots, K$, it uses a noisy gradient $\mathbf{G}(\mathbf{w}_k)$ instead of $\nabla F(\mathbf{w}_k)$ to set the next iterate as $\mathbf{w}_{k+1} = \mathbf{w}_k - \eta_k \mathbf{G}(\mathbf{w}_k)$, where $\eta_k > 0$ is a step size. The overall performance of stochastic gradient methods is controlled by the noise in $\mathbf{G}(\mathbf{w}_k)$ with respect to $\nabla F(\mathbf{w}_k)$ [1]. Often, noisy gradients with large variances lead to slower convergence and degraded performance [2].

Mini-batch stochastic gradient methods, as well as their distributed or parallelized variants, have been proposed to tackle some of these issues [3, 4]. Recently, federated learning [5] has been proposed as a decentralized optimization framework, where SGD runs on a large dataset distributed across a number of devices performing local model updates and sending them to a centralized server that aggregates them, under privacy and communication constraints. In typical resource- and budget-constrained applications, as the mini-batch size increases, the cost available to be allocated to each single stochastic gradient in the mini-batch decreases, so that its quality degrades, i.e., its noise variance increases. A common approach is to focus on the tradeoff between the rate of convergence and the computational complexity of stochastic gradient methods, where the dependence of the noise variance on the cost allocated to stochastic gradients is often omitted.

*This work was supported in part by Systems on Nanoscale Information fabriCs (SONIC), one of the six SRC STARnet Centers, sponsored by MARCO and DARPA, and in part by the Center for Science of Information (CSoI), an NSF Science and Technology Center, under grant agreement CCF-0939370.

[†]donmez2@illinois.edu

[‡]maxim@illinois.edu

[§]acsinger@illinois.edu

In this paper, we propose an alternative framework and consider the tradeoff between *fidelity* and *cost* of computing a stochastic gradient. In particular, we model a noisy gradient as an unbiased estimate of the true gradient, where the noise variance depends on the incurred cost, and this dependence is formalized through a *cost-fidelity* function. We focus on mini-batch oracles, where each mini-batch oracle distributes a limited budget across a mini-batch of stochastic gradients and aggregates them to form a final gradient estimate. We assume that the aggregation operation also incurs a cost from the budget, as does each of the noisy gradients in the mini-batch. The optimal mini-batch size in minimizing the noise variance depends on the underlying cost-fidelity function.

We focus on determining the optimal mini-batch oracle in terms of the cost-fidelity tradeoff when the cost-fidelity function is unknown. In particular, we propose and analyze EE-Grad: an algorithm that, on each iteration, performs sequential trials over different mini-batch oracles to *explore* the performance of each mini-batch oracle with high precision and *exploit* the current knowledge to focus on the one that seems to provide the best performance, i.e., the smallest noise variance. We demonstrate that the proposed algorithm performs almost as well as the optimal mini-batch oracle on each iteration in expectation. We apply this result to the case of strongly convex objectives, and prove performance guarantees in terms of the rate of convergence. We finally provide a numerical example to illustrate our theoretical results.

2 Cost-Fidelity Tradeoff and Mini-Batch Stochastic Gradient Oracles

Suppose that, on each iteration, a stochastic gradient $\mathbf{g}(\mathbf{w}, \theta)$ and the gradient $\nabla F(\mathbf{w})$ are related as

$$\mathbf{g}(\mathbf{w}, \theta) = \nabla F(\mathbf{w}) + \mathbf{U}(\mathbf{w}, \theta) \quad (1)$$

where $\mathbf{U}(\mathbf{w}, \theta)$ is a zero-mean perturbation with a positive definite and diagonal covariance matrix $\theta^{-1}\mathbf{M}(\mathbf{w})$ for $\theta > 0$. That is,

$$\mathbb{E}_{\mathbf{w}}[\mathbf{U}(\mathbf{w}, \theta)] = 0, \quad \mathbb{E}_{\mathbf{w}}[\mathbf{U}(\mathbf{w}, \theta)\mathbf{U}(\mathbf{w}, \theta)^T] = \theta^{-1}\mathbf{M}(\mathbf{w}),$$

where $\mathbb{E}_{\mathbf{w}}[\cdot]$ is the conditional expectation given \mathbf{w} . Here, θ is the *fidelity* of the stochastic gradient $\mathbf{g}(\mathbf{w}, \theta)$. We assume that i th element of $\mathbf{U}(\mathbf{w}, \theta)$ is sub-Gaussian with the parameter $\theta^{-1}\mathbf{M}(\mathbf{w})_{i,i}$, i.e.,

$$\mathbb{E}_{\mathbf{w}}[e^{\lambda \mathbf{U}(\mathbf{w}, \theta)_i}] \leq e^{\lambda^2 \mathbf{M}(\mathbf{w})_{i,i}/2\theta}, \quad \forall \lambda \in \mathbb{R}, \quad (2)$$

for $i \in [d]^1$. A mini-batch stochastic gradient is computed by averaging n independent noisy gradients $\mathbf{g}_i(\mathbf{w}, \theta) = \nabla F(\mathbf{w}) + \mathbf{U}_i(\mathbf{w}, \theta)$, $i \in [n]$, each with fidelity θ :

$$\mathbf{G}(\mathbf{w}) = \frac{1}{n} \sum_{i=1}^n \mathbf{g}_i(\mathbf{w}, \theta), \quad (3)$$

which has the covariance matrix $\mathbf{M}(\mathbf{w})/n\theta$, and satisfies $\mathbb{E}_{\mathbf{w}}[\|\nabla F(\mathbf{w}) - \mathbf{G}(\mathbf{w})\|_2^2] = S(\mathbf{w})/n\theta$, where $S(\mathbf{w}) = \text{Tr}(\mathbf{M}(\mathbf{w}))$ is the trace of the covariance matrix.

A stochastic gradient $\mathbf{g}(\mathbf{w}, \theta)$ with fidelity $\theta > 0$ incurs a cost $C(\theta)$, which is a strictly increasing function of θ with $\lim_{\theta \rightarrow 0} C(\theta) = c_{\min} \geq 0$. We assume that the cost function $C(\theta)$ is unknown. There is also an

¹For any positive integer N , $[N] \triangleq \{1, \dots, N\}$.

aggregation cost $D(n)$ to perform the averaging operation, where $D(n)$ is increasing with $D(1) = 0$. Hence, given a budget $B > 0$, the maximum feasible mini-batch size is

$$N = \max \left\{ n \in \mathbb{Z}_+ \mid B > nc_{\min} + D(n) \right\}.$$

Here, we define, for each $n \in [N]$, a mini-batch oracle $\text{MBO}(n, B, \mathbf{w})$ that computes a mini-batch stochastic gradient $\mathbf{G}(\mathbf{w}, n)$ as in (3) using the fidelity

$$\theta_n \triangleq C^{-1} \left(\frac{B - D(n)}{n} \right).$$

That is, each individual stochastic gradient in the mini-batch is allocated $(B - D(n))/n$ in cost. Therefore, the covariance matrix of $\mathbf{G}(\mathbf{w}, n)$ is $\sigma_n^2 \mathbf{M}(\mathbf{w})$, where

$$\sigma_n^2 \triangleq \frac{1}{n\theta_n}$$

is unknown, since the cost function $C(\theta)$ is assumed unknown. Note that, given $\nabla F(\mathbf{w})$, the concentration of $\mathbf{G}(\mathbf{w}, n)$ around $\nabla F(\mathbf{w})$ is completely governed by σ_n^2 for each $n \in [N]$. The optimal mini-batch size in terms of the cost-fidelity tradeoff is given by

$$n_* \triangleq \underset{n=1, \dots, N}{\operatorname{argmin}} \sigma_n^2,$$

and $\sigma_*^2 \triangleq \sigma_{n_*}^2$. In particular, we define the suboptimality gap of each mini-batch oracle $\text{MBO}(n, B, \mathbf{w})$

$$\Delta_n \triangleq \sigma_n^2 - \sigma_*^2 \geq 0.$$

Since the cost function is unknown, the optimal mini-batch size n_* and σ_*^2 , and hence the optimal mini-batch oracle $\text{MBO}(n_*, B, \mathbf{w})$, are unknown. In the next section, we propose an algorithm that attempts to *learn* the optimal mini-batch oracle over sequential trials in the sense that its noise variance is almost as small as the optimal mini-batch oracle on each iteration.

3 The EE-Grad Algorithm

In this section, we present EE-Grad: an algorithm that, on each iteration of the SGD, aggregates stochastic gradients computed over sequential trials, where at each trial it estimates the optimal mini-batch size and uses the available per-round budget to query the corresponding mini-batch oracle. EE-Grad constructs a high confidence bound on the variance estimate of each mini-batch oracle by exploiting the sub-Gaussian assumption on the noisy gradients. We demonstrate that, in expectation, the algorithm performs almost as well as the optimal mini-batch oracle at each iteration.

On each SGD iteration, EE-Grad performs the following T -round procedure. On round $t = 1, \dots, T$, it picks a mini-batch size $n_t \in [N]$ based on a strategy introduced later in this section, and uses the per-round budget B to query the mini-batch oracle $\text{MBO}(n_t, B, \mathbf{w})$. The oracle returns $\mathbf{G}_t(\mathbf{w}) = \mathbf{G}_t(\mathbf{w}, n_t)$, an unbiased estimate of $\nabla F(\mathbf{w})$, with covariance matrix $\sigma_{n_t}^2 \mathbf{M}(\mathbf{w})$. After T rounds, the algorithm outputs the stochastic gradient $\mathbf{G}(\mathbf{w}) = (1/T) \sum_{t=1}^T \mathbf{G}_t(\mathbf{w})$.

We denote the number of rounds the algorithm picks $\text{MBO}(n, B, \mathbf{w})$ up to round t as $\gamma_t(n)$, index its outputs as $\mathbf{G}_1(\mathbf{w}, n), \dots, \mathbf{G}_{\gamma_t(n)}(\mathbf{w}, n)$, and write its sample mean and sample covariance matrix as

$$\mathbf{m}_t(n) = \frac{1}{\gamma_t(n)} \sum_{i=1}^{\gamma_t(n)} \mathbf{G}_i(\mathbf{w}, n), \quad \mathbf{Cov}_t(n) = \frac{1}{\gamma_t(n) - 1} \sum_{i=1}^{\gamma_t(n)} (\mathbf{G}_i(\mathbf{w}, n) - \mathbf{m}_t(n))(\mathbf{G}_i(\mathbf{w}, n) - \mathbf{m}_t(n))^T,$$

respectively, for $n \in [N]$. The algorithm computes the trace of the sample covariance matrix, denoted by $V_t(n) = \text{Tr}(\mathbf{Cov}_t(n))$ for each $n \in [N]$. Note that $\mathbb{E}_{\mathbf{w}}[V_t(n)] = \sigma_n^2 S(\mathbf{w})$, which implies that for each $\text{MBO}(n, B, \mathbf{w})$, the trace of its sample covariance matrix is an unbiased estimate of $\sigma_n^2 S(\mathbf{w})$.

We emphasize that this framework is similar to the stochastic multi-armed bandit setup that involves an exploration/exploitation tradeoff when picking different arms over sequential trials [6]. In particular, algorithms that exploit the available knowledge on the current best arm and explore the other arms to estimate the actual best arm with higher precision have been shown to yield satisfactory performance [6, 7]. We adopt a similar approach here, and propose an algorithm that simultaneously performs exploration and exploitation. More precisely, EE-Grad first initializes by picking each mini-batch oracle exactly twice, so that $\gamma_t(n) = 2$ for each $n \in [N]$ at trial $t = 2N$, and then picks the mini-batch oracle at trial $t = 2N + 1, \dots, T$ according to

$$n_t \in \underset{n=1, \dots, N}{\text{argmin}} \left[V_t(n) - f\left(\frac{\alpha \ln(t)}{\gamma_t(n) - 1}\right) \right], \quad (4)$$

for some $\alpha > 2$, where

$$f(x) \triangleq \beta P \sqrt{\frac{xd}{c}} \max\left(1, \sqrt{\frac{x}{cd}}\right), \quad (5)$$

and $c > 0$ is a universal constant that comes from the use of Hanson-Wright inequality, as detailed in the proof of Theorem 1. Here we assume that β and P are known constants such that $\sigma_n^2 \leq \beta$ for each $n \in [N]$, and $S(\mathbf{w}) \leq P$. This algorithm constructs an upper confidence bound on the trace of the sample covariance matrix of each mini-batch oracle, and picks the one with the best estimate. The overall scheme, presented below as Algorithm 1, will be analyzed using techniques similar to the ones used in UCB strategies [7, 8, 9], as explained in the proof of Theorem 1.

Algorithm 1 EE-Grad

Input: $N, T > 1$, per-round budget B .

Initialization: Pick each mini-batch oracle twice in the first $2N$ rounds.

for $t = 2N + 1 : T$ **do**

 Compute $V_t(n)$ for each $n \in [N]$, and pick a mini-batch size n_t based on (4).

 Distribute the budget B to $\text{MBO}(n_t, B, \mathbf{w})$, which reveals $\mathbf{G}_t(\mathbf{w}, n)$, and set $\mathbf{G}_t(\mathbf{w}) = \mathbf{G}_t(\mathbf{w}, n)$.

end for

 Compute the final gradient estimate as $\mathbf{G}(\mathbf{w}) = (1/T) \sum_{t=1}^T \mathbf{G}_t(\mathbf{w})$,

4 EE-Grad Performance Guarantees

In this section, we investigate the performance of EE-Grad. In particular, we prove an upper bound on its noise variance, and compare it to the noise variance achieved by the optimal mini-batch oracle:

Theorem 1. *On each iteration, the stochastic gradient computed by EE-Grad satisfies*

$$\mathbb{E}_{\mathbf{w}} \left[\|\mathbf{G}(\mathbf{w}) - \nabla F(\mathbf{w})\|_2^2 \right] \leq Z_T(\mathbf{w}) S(\mathbf{w}),$$

where

$$Z_T(\mathbf{w}) = \frac{\sigma_*^2}{T} + \left(\frac{\ln T}{T^2} \right) C_1(\mathbf{w}) + \left(\frac{1}{T^2} \right) C_2,$$

and

$$C_1(\mathbf{w}) \triangleq \sum_{n: \Delta_n > 0} \frac{\alpha \Delta_n}{\phi(\Delta_n S(\mathbf{w})/2)}, \quad C_2 \triangleq \left(\sum_{n=1}^N \Delta_n \right) \frac{2(\alpha-1)}{\alpha-2}, \quad \phi(\varepsilon) \triangleq \frac{c\varepsilon}{\beta P} \min\left(1, \frac{\varepsilon/d}{\beta P}\right).$$

Also, the stochastic gradient $\mathbf{G}^*(\mathbf{w})$ computed by the optimal mini-batch oracle satisfies

$$\mathbb{E}_{\mathbf{w}} \left[\|\mathbf{G}^*(\mathbf{w}) - \nabla F(\mathbf{w})\|_2^2 \right] = \frac{\sigma_*^2}{T} S(\mathbf{w}).$$

Proof. We prove this theorem in several steps. We first analyze the difference between the noise variance of the stochastic gradient generated by EE-Grad and that of the optimal mini-batch oracle. We next show that this quantity is related to the pseudo-regret term that appears in stochastic multi-armed bandit problems, where UCB-type strategies are used to achieve upper bounds on the pseudo-regret by leveraging concentration inequalities. We present a similar formulation to analyze the behavior of the proposed algorithm with respect to the optimal mini-batch oracle. To prove the upper bound, we first demonstrate that the trace of the sample covariance matrix for each mini-batch oracle, which is used to pick a oracle on each trial in (4), can be written as a quadratic form of independent sub-Gaussian random variables. We combine this observation with the Hanson-Wright inequality [10] to prove a high probability tail bound on the estimate of the optimal mini-batch size. This result also is the derivation of the rule in (4). Based on these results, we prove a pseudo-regret bound and connect this bound to the noise variance achieved by EE-Grad.

Note that, on each iteration, the stochastic gradient of the optimal mini-batch oracle after T rounds is

$$\mathbf{G}^*(\mathbf{w}) \triangleq \frac{1}{T} \sum_{t=1}^T \mathbf{G}_t(\mathbf{w}, n_*),$$

where $\mathbf{G}_1(\mathbf{w}, n_*), \dots, \mathbf{G}_T(\mathbf{w}, n_*)$ are independent. We observe that

$$\begin{aligned} & \mathbb{E}_{\mathbf{w}} \left[\|\mathbf{G}(\mathbf{w}) - \nabla F(\mathbf{w})\|_2^2 \right] - \mathbb{E}_{\mathbf{w}} \left[\|\mathbf{G}^*(\mathbf{w}) - \nabla F(\mathbf{w})\|_2^2 \right] \\ &= \frac{1}{T^2} \left(\sum_{t=1}^T \mathbb{E}_{\mathbf{w}} \left[\|\mathbf{G}_t(\mathbf{w}) - \nabla F(\mathbf{w})\|_2^2 \right] - \sum_{t=1}^T \mathbb{E}_{\mathbf{w}} \left[\|\mathbf{G}_t(\mathbf{w}, n_*) - \nabla F(\mathbf{w})\|_2^2 \right] \right) \\ &= \frac{1}{T^2} \left(\sum_{t=1}^T \mathbb{E}_{\mathbf{w}} \left[\|\mathbf{G}_t(\mathbf{w}) - \nabla F(\mathbf{w})\|_2^2 \right] - T \sigma_*^2 S(\mathbf{w}) \right), \end{aligned} \tag{6}$$

where in (6) we used $\mathbb{E}_{\mathbf{w}} \left[\|\mathbf{G}_t(\mathbf{w}, n_*) - \nabla F(\mathbf{w})\|_2^2 \right] = \sigma_*^2 S(\mathbf{w})$ for each $t \in [T]$. We next observe that

$$\begin{aligned} \mathbb{E}_{\mathbf{w}} \left[\|\mathbf{G}_t(\mathbf{w}) - \nabla F(\mathbf{w})\|_2^2 \right] &= \mathbb{E}_{\mathbf{w}} \left[\mathbb{E}_{\mathbf{w}} \left[\|\mathbf{G}_t(\mathbf{w}) - \nabla F(\mathbf{w})\|_2^2 \mid n_t \right] \right] \\ &= \mathbb{E}_{\mathbf{w}} \left[\sigma_{n_t}^2 \right] S(\mathbf{w}), \end{aligned} \tag{7}$$

where in (7) the expectation is with respect to the randomness in n_t . In particular, we can write

$$\mathbb{E}_{\mathbf{w}}[\sigma_{n_t}^2] = \sum_{n=1}^N \sigma_n^2 \Pr(n_t = n) \quad (8)$$

for each $t \in [T]$. If we substitute (8) into (7) and use the result in (6), then we obtain

$$\begin{aligned} & \mathbb{E}_{\mathbf{w}} \left[\|\mathbf{G}(\mathbf{w}) - \nabla F(\mathbf{w})\|_2^2 \right] - \mathbb{E}_{\mathbf{w}} \left[\|\mathbf{G}^*(\mathbf{w}) - \nabla F(\mathbf{w})\|_2^2 \right] \\ &= \frac{1}{T^2} \left(\sum_{n=1}^N \sigma_n^2 \sum_{t=1}^T \Pr(n_t = n) - T \sigma_*^2 \right) S(\mathbf{w}), \\ &= \frac{1}{T^2} \left(\sum_{n=1}^N \sigma_n^2 \sum_{t=1}^T \mathbb{E}_{\mathbf{w}}[\mathbb{1}\{n_t = n\}] - T \sigma_*^2 \right) S(\mathbf{w}) \end{aligned} \quad (9)$$

$$= \frac{1}{T^2} \left(\sum_{n=1}^N \sigma_n^2 \mathbb{E}_{\mathbf{w}}[\gamma_T(n)] - \sigma_*^2 \sum_{n=1}^N \mathbb{E}_{\mathbf{w}}[\gamma_T(n)] \right) S(\mathbf{w}) \quad (10)$$

$$= \frac{1}{T^2} \mathbb{E}_{\mathbf{w}} \left[\sum_{n=1}^N \Delta_n \gamma_T(n) \right] S(\mathbf{w}), \quad (11)$$

where in (9) we used $\Pr(n_t = n) = \mathbb{E}_{\mathbf{w}}[\mathbb{1}\{n_t = n\}]$, in (10) we used $\gamma_T(n) = \sum_{t=1}^T \mathbb{1}\{n_t = n\}$ and $\sum_{n=1}^N \gamma_T(n) = T$, and in (11) we used $\Delta_n = \sigma_n^2 - \sigma_*^2$. We note that the term $\mathbb{E}_{\mathbf{w}}[\sum_{n=1}^N \Delta_n \gamma_T(n)] S(\mathbf{w})$ is similar to the pseudo-regret term that appears in stochastic multi-armed bandit problems, where there are N arms with unknown reward distributions [6]. We derive the strategy in (4) based on similar arguments, where we leverage a novel application of the Hanson-Wright inequality to the trace of the sample covariance matrix of each mini-batch oracle to prove concentration inequalities.

To prove an upper bound on (11), we first show in Lemma 1 that $V_t(n)$ can be written as a quadratic form of sub-Gaussian random variables as

$$V_t(n) = \mathbf{s}_{t,n}^T \mathbf{A}_{t,n} \mathbf{s}_{t,n}, \quad n \in [N],$$

where $\mathbf{s}_{t,n} \triangleq \left(\mathbf{G}_1(\mathbf{w}, n)^T, \dots, \mathbf{G}_{\gamma_t(n)}(\mathbf{w}, n)^T \right)^T$, and

$$\mathbf{A}_{t,n} = \frac{1}{\gamma_t(n) - 1} \left(\mathbf{I} - \frac{1}{\gamma_t(n)} \mathbf{E} \right),$$

$\mathbf{I} \in \mathbb{R}^{d\gamma_t(n) \times d\gamma_t(n)}$ is an identity matrix, and $\mathbf{E} \in \mathbb{R}^{d\gamma_t(n) \times d\gamma_t(n)}$ is a block matrix with $d \times d$ identity blocks. We next apply the Hanson-Wright inequality [10, 11] to $V_t(n)$ for each $n \in [N]$ to obtain high confidence bounds. This inequality provides a tail probability bound for an arbitrary quadratic function of independent sub-Gaussian random variables. We present it in the appendix for completeness. Moreover, Lemma 3 shows that the tail probability of the trace of the sample covariance matrix of each mini-batch oracle satisfies, for any $\varepsilon > 0$,

$$\Pr(V_t(n) - \sigma_n^2 S(\mathbf{w}) > \varepsilon) \leq \exp(-(\gamma_t(n) - 1)\phi(\varepsilon)), \quad (12)$$

where

$$\phi(\varepsilon) \triangleq \frac{c\varepsilon}{\beta P} \min\left(1, \frac{\varepsilon/d}{\beta P}\right),$$

for each $n \in [N]$. We observe that $\phi = f^{-1}$, where f is defined in (5).

Note that (12) is equivalent to stating that, for any $\delta \in (0, 1)$,

$$V_t(n) - f\left(\frac{1}{\gamma_t(n) - 1} \ln\left(\frac{1}{\delta}\right)\right) \leq \sigma_n^2 S(\mathbf{w}) \quad (13)$$

with probability at least $1 - \delta$. Using this result, we propose the UCB-type strategy in (4) to pick the mini-batch oracle on round t . In particular, we show in Lemma 4 that, for any $\alpha > 2$, we have

$$\mathbb{E}_{\mathbf{w}} \left[\sum_{n=1}^N \Delta_n \gamma_T(n) \right] S(\mathbf{w}) \leq (C_1(\mathbf{w}) \ln(T) + C_2) S(\mathbf{w}), \quad (14)$$

where

$$C_1(\mathbf{w}) \triangleq \sum_{n: \Delta_n > 0} \frac{\alpha \Delta_n}{\phi(\Delta_n S(\mathbf{w})/2)}, \quad C_2 \triangleq \left(\sum_{n=1}^N \Delta_n \right) \frac{2(\alpha - 1)}{\alpha - 2}.$$

Finally, if we use (14) in (11), then we obtain

$$\mathbb{E}_{\mathbf{w}} \left[\|\mathbf{G}(\mathbf{w}) - \nabla F(\mathbf{w})\|_2^2 \right] - \mathbb{E}_{\mathbf{w}} \left[\|\mathbf{G}^*(\mathbf{w}) - \nabla F(\mathbf{w})\|_2^2 \right] \leq \frac{1}{T^2} (C_1(\mathbf{w}) \ln(T) + C_2) S(\mathbf{w}), \quad (15)$$

where substituting $\mathbb{E}_{\mathbf{w}} \left[\|\mathbf{G}^*(\mathbf{w}) - \nabla F(\mathbf{w})\|_2^2 \right] = \sigma_*^2 S(\mathbf{w})/T$ in (15) yields the desired result. \square

5 SGD Performance Under Strongly Convex Objectives

In this section, we investigate the performance of EE-Grad with strongly convex objective functions with Lipschitz continuous gradients. That is, we assume that the gradient ∇F is Lipschitz continuous with Lipschitz constant $L > 0$, i.e.,

$$\|\nabla F(\mathbf{w}) - \nabla F(\bar{\mathbf{w}})\|_2 \leq L \|\mathbf{w} - \bar{\mathbf{w}}\|_2, \quad \forall \mathbf{w}, \bar{\mathbf{w}} \in \mathbb{R}^d,$$

and there exists $m > 0$ such that

$$F(\bar{\mathbf{w}}) \geq F(\mathbf{w}) + \nabla F(\mathbf{w})^T (\bar{\mathbf{w}} - \mathbf{w}) + \frac{1}{2} m \|\bar{\mathbf{w}} - \mathbf{w}\|_2^2, \quad \forall \bar{\mathbf{w}}, \mathbf{w} \in \mathbb{R}^d.$$

Let $\mathbf{w}_* = \arg\min_{\mathbf{w} \in \mathbb{R}^d} F(\mathbf{w})$ be the global minimizer. We first describe the *optimal* mini-batch SGD algorithm that uses the optimal mini-batch oracle on each iteration. We next compare its performance to EE-Grad in terms of the rate of convergence to the global solution \mathbf{w}_* . Note that the cost function $C(\theta)$, and hence the optimal mini-batch size, is allowed to vary across iterations of the SGD algorithm. We use the subscript k , which denotes the SGD iteration, for the quantities introduced in Section 2 and Section 3 to emphasize the iteration dependence whenever necessary.

On each iteration $k = 1, \dots, K$, the optimal mini-batch SGD algorithm that knows the optimal mini-batch oracle $\text{MBO}(n_{*,k}, B, \mathbf{w}_k^o)$ distributes the per-round budget B_k to it producing $\mathbf{G}_t^o(\mathbf{w}_k^o) = \mathbf{G}_t^o(\mathbf{w}_k^o, n_{*,k})$ on each trial $t = 1, \dots, T$. After T trials, it computes its final stochastic gradient as $\mathbf{G}^o(\mathbf{w}_k^o) = (1/T) \sum_{t=1}^T \mathbf{G}_t^o(\mathbf{w}_k^o)$, and sets the next iterate as $\mathbf{w}_{k+1}^o = \mathbf{w}_k^o - \eta \mathbf{G}^o(\mathbf{w}_k^o)$. We observe that \mathbf{w}_k and \mathbf{w}_k^o may be different over iterations, so the true gradients $\nabla F(\mathbf{w}_k)$ and $\nabla F(\mathbf{w}_k^o)$ also may differ. Also, note that $\mathbf{G}^o(\mathbf{w}_k^o)$ satisfies

$$\mathbb{E}_{\mathbf{w}} \left[\|\mathbf{G}^o(\mathbf{w}_k^o) - \nabla F(\mathbf{w}_k^o)\|_2^2 \right] = \frac{\sigma_*^2 S(\mathbf{w}_k^o)}{T},$$

where $S(\mathbf{w}_k^o) \triangleq \text{Tr}(\mathbf{M}(\mathbf{w}_k^o))$ for each $k \in [K]$. In this section, we focus on the case where

$$\mathbf{M}(\mathbf{w}) \triangleq \text{diag}\left(\nabla F(\mathbf{w})_1^2, \dots, \nabla F(\mathbf{w})_d^2\right)$$

for any $\mathbf{w} \in \mathbb{R}^d$, which implies that $S(\mathbf{w}) = \|\nabla F(\mathbf{w})\|_2^2$.

We define the expected gaps of EE-Grad and of the optimal mini-batch SGD algorithm with respect to the global minimizer \mathbf{w}_* on each iteration k as

$$J_{k,\eta} \triangleq \mathbb{E}[F(\mathbf{w}_k)] - F(\mathbf{w}_*), \quad J_{k,\eta}^o \triangleq \mathbb{E}[F(\mathbf{w}_k^o)] - F(\mathbf{w}_*), \quad (16)$$

respectively. The next theorem shows how these expected gaps evolve over iterations.

Theorem 2. *Suppose that the step size η_k is sufficiently small so that it satisfies*

$$0 < \eta_k < \frac{2}{L(1 + Z_T(\mathbf{w}_k))}. \quad (17)$$

Then, on each iteration k , the expected gap of the optimal mini-batch SGD algorithm satisfies

$$J_{k+1,\eta}^o \leq \tau_k^o(\eta_k) J_{k,\eta}^o,$$

where

$$0 < \tau_k^o(\eta_k) \triangleq mL\eta_k^2(1 + \sigma_{*,k}^2/T) - 2m\eta_k + 1 < 1.$$

Moreover, the expected gap of the EE-Grad Algorithm on iteration k satisfies

$$J_{k+1,\eta} \leq \tau_k(\eta_k) J_{k,\eta},$$

where

$$0 < \tau_k(\eta_k) \triangleq \tau_k^o(\eta_k) + mL\eta_k^2 O_{T,k} < 1,$$

and $O_{T,k} \triangleq Z_T(\mathbf{w}_k) - \sigma_{*,k}^2/T = C_{1,k}(\mathbf{w}) \ln T/T^2 + C_{2,k}/T^2 > 0$, where $O_{T,k} \rightarrow 0$ as $T \rightarrow \infty$.

Proof. First note that since ∇F is Lipschitz continuous with Lipschitz constant $L > 0$, it satisfies [1]

$$F(\mathbf{w}) \leq F(\bar{\mathbf{w}}) + \nabla F(\bar{\mathbf{w}})^T (\mathbf{w} - \bar{\mathbf{w}}) + \frac{1}{2}L\|\mathbf{w} - \bar{\mathbf{w}}\|_2^2, \quad \forall \mathbf{w}, \bar{\mathbf{w}} \in \mathbb{R}^d,$$

which implies that on each iteration k , we have

$$F(\mathbf{w}_{k+1}) - F(\mathbf{w}_k) \leq -\eta_k \nabla F(\mathbf{w}_k)^T \mathbf{G}(\mathbf{w}_k) + \frac{1}{2}L\eta_k^2 \|\mathbf{G}(\mathbf{w}_k)\|_2^2. \quad (18)$$

By taking conditional expectations of both sides and rearranging the terms, we obtain

$$\mathbb{E}_k[F(\mathbf{w}_{k+1})] - F(\mathbf{w}_k) \leq -\eta_k S(\mathbf{w}_k) \left(1 - \frac{1}{2}\eta_k L(1 + Z_T(\mathbf{w}_k))\right). \quad (19)$$

Performing the same steps on the optimal mini-batch SGD algorithm yields

$$\mathbb{E}_k[F(\mathbf{w}_{k+1}^o)] - F(\mathbf{w}_k^o) \leq -\eta_k S(\mathbf{w}_k^o) \left(1 - \frac{1}{2}\eta_k L(1 + \sigma_{*,k}^2/T)\right). \quad (20)$$

Since F is assumed to be m -strongly convex, the optimality gap for any $\mathbf{w} \in \mathbb{R}^d$ satisfies [1]

$$F(\mathbf{w}) - F(\mathbf{w}_*) \leq \frac{1}{2m} \|\nabla F(\mathbf{w})\|_2^2. \quad (21)$$

The assumption in (17) guarantees that $1 - \eta_k L(1 + Z_T(\mathbf{w}_k))/2 > 0$. Thus, using (21) in (19), subtracting $F(\mathbf{w}_*)$ on both sides, and rearranging terms give

$$\begin{aligned} \mathbb{E}_k[F(\mathbf{w}_{k+1})] - F(\mathbf{w}_*) &\leq F(\mathbf{w}_k) - F(\mathbf{w}_*) - \eta_k S(\mathbf{w}_k) \left(1 - \frac{1}{2} \eta_k L(1 + Z_T(\mathbf{w}_k))\right) \\ &\leq F(\mathbf{w}_k) - F(\mathbf{w}_*) - 2m\eta_k (F(\mathbf{w}_k) - F(\mathbf{w}_*)) \left(1 - \frac{1}{2} \eta_k L(1 + Z_T(\mathbf{w}_k))\right) \\ &= (mL\eta_k^2(1 + Z_T(\mathbf{w}_k)) - 2m\eta_k + 1)(F(\mathbf{w}_k) - F(\mathbf{w}_*)), \\ &= \tau_k(\eta_k)(F(\mathbf{w}_k) - F(\mathbf{w}_*)). \end{aligned}$$

Here if we take expectations of both sides and note the definition in (16), then we obtain $J_{k+1,\eta} \leq \tau_k(\eta_k)J_{k,\eta}$. Similar steps for the optimal mini-batch SGD algorithm imply $J_{k+1,\eta}^o \leq \tau_k^o(\eta_k)J_{k,\eta}^o$, where $\tau_k(\eta_k) = \tau_k^o(\eta_k) + mL\eta_k^2 O_{T,k}$, so that $\tau_k(\eta_k) - \tau_k^o(\eta_k) \rightarrow 0$ as $T \rightarrow \infty$. \square

Here, we note that $\tau_k(\eta_k)$ is a quadratic function of η_k , minimized at $\eta_k = 1/(1 + Z_T(\mathbf{w}_k))$, and $\tau_k(\eta_k) < 1$ for all η_k satisfying (17). Similarly, $\tau_k^o(\eta_k)$ is a quadratic function of η_k , minimized at $\eta_k = 1/(1 + \sigma_{*,k}^2/T)$, and $\tau_k^o(\eta_k) < 1$ for all η_k satisfying (17). Also, we observe that

$$\tau_k(\eta_k) = \tau_k^o(\eta_k) + mL\eta_k^2 O_{T,k} > \tau_k^o(\eta_k)$$

for all $\eta_k > 0$, i.e., $\tau_k(\eta_k)$ is uniformly larger than $\tau_k^o(\eta_k)$, which implies that the optimal mini-batch SGD algorithm enjoys faster convergence rate than the proposed algorithm. However, the gap between them is proportional to $O_{T,k}$ for any given step size $\eta_k > 0$, which is the gap between EE-Grad and the optimal mini-batch SGD algorithm, as shown in Theorem 1. Finally, we note that this gap diminishes as the number of trials T increases, at the expense of larger total incurred cost. In the next section, we illustrate our theoretical results with numerical examples.

6 Numerical Results

In this section, we present a numerical example based on synthetic data to illustrate our main results. We consider $d = 2$ dimensional case, where the objective function and its gradient are $F(\mathbf{w}) = \mathbf{w}^T \mathbf{w}/2$ and $\nabla F(\mathbf{w}) = \mathbf{w}$, respectively, where $F(\mathbf{w}_*) = 0$ with $\mathbf{w}_* = (0, 0)^T$.

We assume that $\mathbf{M}(\mathbf{w}) = \text{diag}(w_1^2, w_2^2)$, and each stochastic gradient $\mathbf{g}(\mathbf{w}, \theta)$ with fidelity $\theta > 0$ has uncorrelated Gaussian components with the parameters w_1^2/θ and w_2^2/θ , respectively. We next assume that the unknown parameters of the mini-batch oracles are given by $\sigma_1^2 = 50, \sigma_2^2 = 26, \sigma_3^2 = 16.7$, and run the EE-Grad algorithm and the mini-batch oracles with a randomly generated initial iterate for $T = 50$ trials and $K = 5$ iterations by using the constant step size $\eta = 0.85$, where we obtain expected results over 2000 independent realizations. We plot the resulting expected gaps achieved by EE-Grad and the mini-batch oracles in Fig. 1a. We repeat the same procedure for $T = 200$ and $T = 3000$ and plot the results in Fig. 1b and Fig. 1c, where we note that σ_i^2 are scaled accordingly, so that the results over different T s are comparable.

We observe that for this numerical example, the expected gap achieved by the EE-Grad algorithm is close to that of the optimal mini-batch oracle, where the performance difference between them shrinks with increasing T at the expense of increased total cost, as we proved in Theorem 2.




Figure 1: Expected gaps achieved by the EE-Grad Algorithm and the mini-batch oracles for different values of T over $K = 5$ iterations.

7 Discussion

We presented a new framework to analyze the tradeoff between *fidelity* and *cost* of computing a stochastic gradient, where we modeled a noisy gradient as an unbiased estimate of the true gradient such that the noise variance depends on the cost incurred to compute it. We investigated mini-batch oracles that distribute a limited budget to a mini-batch of stochastic gradients and averages them to estimate the true gradient, where the averaging operation is also assumed to be costly. In this framework, the optimal mini-batch size in minimizing the noise variance depends on the underlying cost-fidelity function, which is assumed to be unknown.

We proposed the EE-Grad algorithm that performs sequential trials over different mini-batch oracles to explore the performance of each mini-batch oracle with high precision and exploit the current knowledge to allocate the budget to the one that seems to provide the best performance. We demonstrated that the proposed algorithm performs almost as well as the optimal mini-batch oracle on each iteration in expectation. We next applied this result to the strongly convex objectives with Lipschitz continuous gradients, and provided a performance guarantee on the rate of convergence with respect to the optimal mini-batch oracle. We finally illustrated our theoretical results through numerical experiments on synthetic data.

References

- [1] Léon Bottou, Frank E. Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. arXiv preprint 1606.04838, 2016.
- [2] Chong Wang, Xi Chen, Alex Smola, and Eric P. Xing. Variance reduction for stochastic gradient optimization. In *Proc. 26th Annu. Conf. Neural Inf. Process. Syst. (NIPS)*, pages 181–189, December 2013.
- [3] Mu Li, Tong Zhang, Yuqiang Chen, and Alexander J. Smola. Efficient mini-batch training for stochastic optimization. In *Proc. 20th ACM SIGKDD Int. Conf. Knowl. Discov. Data Min.*, pages 661–670, 2014.
- [4] Martin A. Zinkevich, Markus Weimer, Alex Smola, and Lihong Li. Parallelized stochastic gradient descent. In *Proc. 23rd Annu. Conf. Neural Inf. Process. Syst. (NIPS)*, pages 2595–2603, December 2010.
- [5] H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. arXiv preprint 1602.05629, 2016.
- [6] Sebastien Bubeck and Nicolo Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Found. Trends Mach. Learn.*, 5(1):1–122, December 2012.
- [7] Peter Auer. Using confidence bounds for exploitation-exploration tradeoffs. *J. Mach. Learn. Res.*, 3(1):397–422, November 2002.
- [8] Herbert Robbins. Some aspects of the sequential design of experiments. *Bull. Am. Math. Soc.*, 58(5):527–535, 1952.
- [9] Peter Auer and Ronald Ortner. UCB revisited: Improved regret bounds for the stochastic multi-armed bandit problem. *Peri. Math. Hung.*, 61(1):55–65, September 2010.
- [10] D. L. Hanson and F. T. Wright. A bound on tail probabilities for quadratic forms in independent random variables. *Ann. Math. Stat.*, 41:1079–1083, 1971.
- [11] Mark Rudelson and Roman Vershynin. Hanson-Wright inequality and sub-Gaussian concentration. *Electron. Commun. Probab.*, 18(82):1–9, 2013.

A Trace of the Sample Covariance Matrix as a Quadratic Form

Lemma 1. *On each round t , the trace of the sample covariance matrix $\mathbf{Cov}_t(n)$ can be written as*

$$V_t(n) = \mathbf{s}_{t,n}^T \mathbf{A}_{t,n} \mathbf{s}_{t,n}, \quad n = 1, \dots, N,$$

where $\mathbf{s}_{t,n} = \left(\mathbf{G}_1(\mathbf{w}, n)^T, \dots, \mathbf{G}_{\gamma_t(n)}(\mathbf{w}, n)^T \right)^T$, $\mathbf{A}_{t,n} = (\gamma_t(n) - 1)^{-1} \left(\mathbf{I} - \gamma_t(n)^{-1} \mathbf{E} \right)$, $\mathbf{I} \in \mathbb{R}^{d\gamma_t(n) \times d\gamma_t(n)}$ is an identity matrix, and $\mathbf{E} \in \mathbb{R}^{d\gamma_t(n) \times d\gamma_t(n)}$ is a block matrix with $d \times d$ identity blocks.

Proof. Note that

$$V_t(n) = \text{Tr}(\mathbf{Cov}_t(n)) = \frac{1}{\gamma_t(n) - 1} \left(\sum_{i=1}^{\gamma_t(n)} \mathbf{G}_i(\mathbf{w}, n)^T \mathbf{G}_i(\mathbf{w}, n) - \gamma_t(n) \mathbf{m}_t(n)^T \mathbf{m}_t(n) \right),$$

where

$$\mathbf{m}_t(n)^T \mathbf{m}_t(n) = \frac{1}{\gamma_t(n)^2} \sum_{i=1}^{\gamma_t(n)} \sum_{j=1}^{\gamma_t(n)} \mathbf{G}_i(\mathbf{w}, n)^T \mathbf{G}_j(\mathbf{w}, n) = \frac{1}{\gamma_t(n)^2} \mathbf{s}_{t,n}^T \mathbf{E} \mathbf{s}_{t,n}.$$

Noting $\sum_{i=1}^{\gamma_t(n)} \mathbf{G}_i(\mathbf{w}, n)^T \mathbf{G}_i(\mathbf{w}, n) = \mathbf{s}_{t,n}^T \mathbf{s}_{t,n}$, we conclude that

$$V_t(n) = \frac{1}{\gamma_t(n) - 1} \left(\mathbf{s}_{t,n}^T \mathbf{s}_{t,n} - \frac{1}{\gamma_t(n)} \mathbf{s}_{t,n}^T \mathbf{E} \mathbf{s}_{t,n} \right) = \mathbf{s}_{t,n}^T \mathbf{A}_{t,n} \mathbf{s}_{t,n}.$$

□

B Hanson-Wright Inequality

Lemma 2. *Let $\mathbf{W} = [W_1, \dots, W_m]^T \in \mathbb{R}^m$, $m > 1$, where W_i are zero-mean sub-Gaussian with a parameter $\sigma^2 > 0$. Then, given an arbitrary matrix $\mathbf{A} \in \mathbb{R}^{m \times m}$, we have, for any $\varepsilon > 0$,*

$$\Pr(\mathbf{W}^T \mathbf{A} \mathbf{W} - \mathbb{E}_{\mathbf{W}}[\mathbf{W}^T \mathbf{A} \mathbf{W}] > \varepsilon) \leq \exp\left(-c \min\left(\frac{\varepsilon^2}{\sigma^4 \|\mathbf{A}\|_F^2}, \frac{\varepsilon}{\sigma^2 \|\mathbf{A}\|}\right)\right),$$

where $\|\mathbf{A}\|_F$ and $\|\mathbf{A}\|$ are Frobenius and operator norms of \mathbf{A} , and $c > 0$ is an absolute constant.

C Concentration Result on the Trace of the Sample Covariance Matrices

Lemma 3. *Suppose that $\gamma_t(n) > 1$. Then the tail probability of $V_t(n)$ satisfies, for any $\varepsilon > 0$,*

$$\Pr(V_t(n) - \sigma_n^2 S(\mathbf{w}) > \varepsilon) \leq \exp(-(\gamma_t(n) - 1)\phi(\varepsilon)),$$

where

$$\phi(\varepsilon) \triangleq \frac{c\varepsilon}{\beta P} \min\left(1, \frac{\varepsilon/d}{\beta P}\right),$$

for $n = 1, \dots, N$, where $c > 0$ is an absolute constant.

Proof. Note that $\mathbf{I} - (1/\gamma_t(n))\mathbf{E}$ is a $d\gamma_t(n) \times d\gamma_t(n)$ block matrix with $d \times d$ blocks, where the diagonal and non-diagonal matrices are given by $\frac{\gamma_t(n)-1}{\gamma_t(n)}\mathbf{I}$ and $-\frac{1}{\gamma_t(n)}\mathbf{I}$, respectively, and $\|\mathbf{I}\|_F^2 = d$. This implies

$$\begin{aligned}\|\mathbf{A}_{t,n}\|_F^2 &= \frac{1}{(\gamma_t(n)-1)^2} \left(\gamma_t(n) \left(\frac{\gamma_t(n)-1}{\gamma_t(n)} \right)^2 \|\mathbf{I}\|_F^2 + (\gamma_t(n)-1)\gamma_t(n) \frac{1}{\gamma_t(n)^2} \|\mathbf{I}\|_F^2 \right) \\ &= \frac{d}{\gamma_t(n)-1}.\end{aligned}$$

Next suppose that $\mathbf{v} = (\mathbf{v}_1^T, \dots, \mathbf{v}_{\gamma_t(n)}^T)^T \in \mathbb{R}^{d\gamma_t(n)}$ such that $\mathbf{v}_i \in \mathbb{R}^d$ and $\|\mathbf{v}\|_2 = 1$. Then we write

$$\begin{aligned}\|\mathbf{A}_{t,n}\mathbf{v}\|_2^2 &= \frac{1}{(\gamma_t(n)-1)^2} \left(\|\mathbf{v}\|_2^2 + \frac{1}{\gamma_t(n)^2} \|\mathbf{E}\mathbf{v}\|_2^2 - \frac{2}{\gamma_t(n)} \mathbf{v}^T \mathbf{E}\mathbf{v} \right) \\ &= \frac{1}{(\gamma_t(n)-1)^2} \left(1 - \frac{1}{\gamma_t(n)} \left\| \sum_{i=1}^{\gamma_t(n)} \mathbf{v}_i \right\|_2^2 \right) \leq \frac{1}{(\gamma_t(n)-1)^2},\end{aligned}$$

where equality is achieved by $\mathbf{v} = (\mathbf{v}_1^T, \dots, \mathbf{v}_{\gamma_t(n)}^T)^T$ such that $\mathbf{v}_1 = (\frac{1}{\sqrt{2}}, 0, \dots, 0)$, $\mathbf{v}_2 = -\mathbf{v}_1$, and $\mathbf{v}_i = (0, \dots, 0)$ for $i = 3, \dots, \gamma_t(n)$. This yields

$$\|\mathbf{A}_{t,n}\| = \sup_{\|\mathbf{v}\|_2=1} \|\mathbf{A}_{t,n}\mathbf{v}\|_2 = (\gamma_t(n)-1)^{-1}.$$

We finally note that the trace of the sample covariance matrix can be written as

$$\begin{aligned}V_t(n) &= \frac{1}{\gamma_t(n)-1} \sum_{i=1}^{\gamma_t(n)} (\mathbf{G}_i(\mathbf{w}, n) - \mathbf{m}_t(n))^T (\mathbf{G}_i(\mathbf{w}, n) - \mathbf{m}_t(n)) \\ &= \frac{1}{\gamma_t(n)-1} \sum_{i=1}^{\gamma_t(n)} (\mathbf{Q}_i(n) - \mathbf{q}_t(n))^T (\mathbf{Q}_i(n) - \mathbf{q}_t(n)),\end{aligned}$$

where $\mathbf{Q}_i(n) \triangleq \mathbf{G}_i(\mathbf{w}, n) - \nabla F(\mathbf{w})$ for $i \in [\gamma_t(n)]$, and $\mathbf{q}_t(n) = (1/\gamma_t(n)) \sum_{i=1}^{\gamma_t(n)} \mathbf{Q}_i(n)$. This implies the same expression holds for the mean-removed versions of $\mathbf{G}_i(\mathbf{w}, n)$ s. Hence, we can assume that $\mathbb{E}_{\mathbf{w}}[\mathbf{G}_i(\mathbf{w}, n)] = 0$. We apply Lemma 2 to $V_t(n)$ by using Lemma 1 to get, for any $\varepsilon > 0$,

$$\Pr(V_t(n) - \sigma_n^2 S(\mathbf{w}) > \varepsilon) \leq \exp(-(\gamma_t(n) - 1) \phi_n(\varepsilon)),$$

where

$$\phi_n(\varepsilon) \triangleq \frac{c\varepsilon}{\sigma_n^2 S(\mathbf{w})} \min\left(1, \frac{\varepsilon/d}{\sigma_n^2 S(\mathbf{w})}\right),$$

which is strictly increasing in ε , for $n \in [N]$, where $c > 0$ is an absolute constant. Finally, we note $\phi_n(\varepsilon) \geq \phi(\varepsilon)$, since we assumed $\max_{n=1, \dots, N} \sigma_n^2 \leq \beta$, and $S(\mathbf{w}) \leq P$. This concludes the proof. \square

D Pseudo-Regret Bound

Lemma 4. For any $\alpha > 2$, the pseudo-regret term in (11) satisfies, for any T ,

$$\mathbb{E}_{\mathbf{w}} \left[\sum_{n=1}^N \Delta_n \gamma_T(n) \right] S(\mathbf{w}) \leq (C_1(\mathbf{w}) \ln(T) + C_2) S(\mathbf{w}),$$

where

$$C_1(\mathbf{w}) \triangleq \sum_{n: \Delta_n > 0} \frac{\alpha \Delta_n}{\phi(\Delta_n S(\mathbf{w})/2)}, \quad C_2 \triangleq \left(\sum_{n=1}^N \Delta_n \right) \frac{2(\alpha - 1)}{\alpha - 2}. \quad (\text{D.1})$$

Proof. We follow along similar steps to the proof of Theorem 2.1 in [6]. Suppose that $n_t = n$, and consider the events

$$\begin{aligned} E_{t,1} &\triangleq \left\{ V_t(n_*) - f\left(\frac{\alpha \ln(t)}{\gamma_t(n_*) - 1}\right) \geq \sigma_*^2 S(\mathbf{w}) \right\}, \\ E_{t,2} &\triangleq \left\{ V_t(n) < \sigma_n^2 S(\mathbf{w}) - f\left(\frac{\alpha \ln(t)}{\gamma_t(n) - 1}\right) \right\}, \\ E_{t,3} &\triangleq \left\{ \gamma_t(n) < 1 + \frac{\alpha \ln(T)}{\phi(\Delta_n S(\mathbf{w})/2)} \right\}. \end{aligned}$$

We claim that $E_{t,1} \cup E_{t,2} \cup E_{t,3}$ must occur. Assume, by contradiction, that $E_{t,i}$ are all false. We obtain

$$\begin{aligned} V_t(n_*) - f\left(\frac{\alpha \ln(t)}{\gamma_t(n_*) - 1}\right) &< \sigma_*^2 S(\mathbf{w}) = \sigma_n^2 S(\mathbf{w}) - \Delta_n S(\mathbf{w}) \\ &\leq V_t(n) + f\left(\frac{\alpha \ln(t)}{\gamma_t(n) - 1}\right) - \Delta_n S(\mathbf{w}). \end{aligned} \quad (\text{D.2})$$

By assumption $E_{t,3}$ is false, and we have $\gamma_t(n) - 1 \geq \alpha \ln(T)/\phi(\Delta_n S(\mathbf{w})/2)$, which is equivalent to

$$\Delta_n S(\mathbf{w}) \geq 2f\left(\frac{\alpha \ln(T)}{\gamma_t(n) - 1}\right), \quad (\text{D.3})$$

If we use (D.3) in (D.2), then we obtain the following result, which contradicts the rule in (4):

$$V_t(n_*) - f\left(\frac{\alpha \ln(t)}{\gamma_t(n_*) - 1}\right) < V_t(n) - f\left(\frac{\alpha \ln(t)}{\gamma_t(n) - 1}\right).$$

For all n such that $\Delta_n > 0$, we define

$$M_n \triangleq \left\lceil \frac{\alpha \ln(T)}{\phi(\Delta_n S(\mathbf{w})/2)} \right\rceil.$$

We next upper bound $\mathbb{E}_{\mathbf{w}}[\gamma_T(n)]$ as

$$\begin{aligned} \mathbb{E}_{\mathbf{w}}[\gamma_T(n)] &= \mathbb{E}_{\mathbf{w}} \left[\sum_{t=1}^T \mathbb{1}\{n_t = n \text{ and } \gamma_t(n) < M_n\} \right] + \mathbb{E}_{\mathbf{w}} \left[\sum_{t=1}^T \mathbb{1}\{n_t = n \text{ and } \gamma_t(n) \geq M_n\} \right] \\ &\leq M_n + \mathbb{E}_{\mathbf{w}} \left[\sum_{t=M_n+1}^T \mathbb{1}\{n_t = n \text{ and } \gamma_t(n) \geq M_n\} \right]. \end{aligned} \quad (\text{D.4})$$

In (D.4), we observe that $\gamma_t(n) \geq M_n$ is equivalent to $E_{t,3}$ being false, which is further equivalent to $E_{t,1} \cup E_{t,2}$ being true, i.e., $E_{t,1}$ or $E_{t,2}$ must occur. Therefore we can further upper bound (D.4) as

$$\begin{aligned}\mathbb{E}_{\mathbf{w}}[\gamma_T(n)] &\leq M_n + \mathbb{E}_{\mathbf{w}} \left[\sum_{t=M_n+1}^T \mathbb{1}\{E_{t,1} \text{ or } E_{t,2} \text{ is true}\} \right] \\ &= M_n + \sum_{t=M_n+1}^T \Pr(E_{t,1} \cup E_{t,2} \text{ is true}) \\ &\leq M_n + \sum_{t=M_n+1}^T \Pr(E_{t,1}) + \sum_{t=M_n+1}^T \Pr(E_{t,2}).\end{aligned}\tag{D.5}$$

where we used the union bound. We upper bound $\Pr(E_{t,1})$ for each $t = M_n + 1, \dots, T$. Note that

$$\Pr(E_{t,1} = 1) = \Pr\left(V_t(n_*) - f\left(\frac{\alpha \ln(t)}{\gamma_t(n_*) - 1}\right) \geq \sigma_*^2 S(\mathbf{w})\right),\tag{D.6}$$

where $\gamma_t(n_*)$ can take values in $\{2, \dots, t\}$. Hence we apply the union bound in (D.6), which yields

$$\Pr(E_{t,1} = 1) \leq \sum_{s=1}^t \Pr\left(V'_s(n_*) - f\left(\frac{\alpha \ln(t)}{s}\right) \geq \sigma_*^2 S(\mathbf{w})\right) \leq \sum_{s=1}^t \frac{1}{t^\alpha} = t^{1-\alpha},\tag{D.7}$$

where (D.7) follows from (13). Here, $V'_s(n_*)$ is the trace of a sample covariance matrix given $s + 2$ independent random vectors with sub-Gaussian components with the parameter $\sigma_*^2 S(\mathbf{w})$. Hence we obtain

$$\sum_{t=M_n+1}^T \Pr(E_{t,1} = 1) \leq \sum_{t=M_n+1}^T t^{1-\alpha} \leq \sum_{t=1}^{\infty} t^{1-\alpha} \leq 1 + \int_1^{\infty} t^{1-\alpha} dt = \frac{\alpha - 1}{\alpha - 2}.\tag{D.8}$$

The same upper bound holds for $\Pr(E_{t,2} = 1)$ so that

$$\sum_{t=M_n+1}^T \Pr(E_{t,2} = 1) \leq \frac{\alpha - 1}{\alpha - 2}.$$

By incorporating these upper bounds into (D.5) we obtain

$$\mathbb{E}_{\mathbf{w}}[\gamma_T(n)] \leq M_n + 2(\alpha - 1)/(\alpha - 2).$$

Finally we use this result to get

$$\mathbb{E}_{\mathbf{w}} \left[\sum_{n=1}^N \Delta_n \gamma_T(n) \right] S(\mathbf{w}) \leq (C_1(\mathbf{w}) \ln(T) + C_2) S(\mathbf{w}),$$

where $C_1(\mathbf{w})$ and C_2 are defined in (D.1). □