

# Beating SGD Saturation with Tail-Averaging and Minibatching

Nicole Mücke\* , Gergely Neu† and Lorenzo Rosasco‡

## Abstract

While stochastic gradient descent (SGD) is a workhorse in machine learning, the learning properties of many practically used variants are hardly known. In this paper, we consider least squares learning and contribute filling this gap focusing on the effect and interplay of multiple passes, mini-batching and averaging, and in particular tail averaging. Our results show how these different flavors of SGD can be combined to achieve optimal learning errors, hence providing practical insights.

## 1 Introduction

Stochastic gradient descent (SGD) provides a simple and yet stunningly efficient way to solve a broad range of machine learning problems. Our starting observation is that, while a number of variants including multiple passes over the data, mini-batching and averaging are commonly used, their combination and learning properties are studied only partially. The literature on convergence properties of SGD is vast, but usually only one pass over the data is considered, see, e.g., [23]. In the context of nonparametric statistical learning, which we consider here, the study of one-pass SGD was probably first considered in [35] and then further developed in a number of papers (e.g., [25, 36, 37]). Another line of work derives statistical learning results for one pass SGD with averaging from a worst-case sequential prediction analysis [18, 28, 29]. The idea of using averaging also has a long history going back to at least the works of [32] and [27], see also [34] and references therein. More recently, averaging was shown to lead to larger, possibly constant, step-sizes, see [2, 10, 11]. A different take on the role of (weighted) averaging was given in [24], highlighting a connection with ridge regression, a.k.a. Tikhonov regularization. A different flavor of averaging called *tail averaging* for one-pass SGD was considered in [19] in a parametric setting. The role of minibatching has also been considered and shown to potentially lead to linear parallelization speedups, see e.g. [7] and references therein. Very few results consider the role of multiple passes for learning. Indeed, this variant of SGD is typically analyzed for the minimization of the empirical risk, rather than the actual population risk, see for example [4]. To the best of our knowledge the first paper to analyze the learning properties of multipass SGD was [31], where a cyclic selection strategy was considered. Other results for multipass SGD were then given in [16] and [20]. Our starting point are the results in [21] where optimal results for multipass SGD were derived considering also the effect of mini-batching. Following the approach in this latter paper, multipass SGD with averaging was analyzed by [26] with no minibatching.

In this paper, we develop and improve the above results on two fronts. On the one hand, we consider for the first time the role of multiple passes, mini-batching and averaging at once. On the other hand, we further study the beneficial effect of tail averaging. Both mini-batching

---

\*Institute for Stochastics and Applications, University of Stuttgart, [nicole.muecke@mathematik.uni-stuttgart.de](mailto:nicole.muecke@mathematik.uni-stuttgart.de)

†Universitat Pompeu Fabra, Barcelona, Spain, [gergely.neu@gmail.com](mailto:gergely.neu@gmail.com)

‡LCSL, Massachusetts Institute of Technology & Istituto Italiano di Tecnologia & DIBRIS, Università degli Studi di Genova, [lrosasco@mit.edu](mailto:lrosasco@mit.edu)

and averaging are known to allow larger step-sizes. Our results show that their combination allows even more aggressive parameter choices. At the same time averaging was shown to lead to slower convergence rates in some cases. In a parametric setting, averaging prevents linear convergence rates [2, 11]. In a nonparametric setting, it prevents exploiting the possible regularity in the solution [10], a phenomenon called *saturation* [12]. In other words, uniform averaging can prevent optimal rates in a nonparametric setting. Our results provide a simple explanation to this effect, showing it has a purely deterministic nature. Further, we show that tail averaging allows to bypass this problem. These results parallel the findings of [19], showing similar beneficial effects of tail-averaging and minibatching in the finite-dimensional setting. Following [21], our analysis relies on the study of batch gradient descent and then of the discrepancy between batch gradient and SGD, with the additional twist that it also considers the role of tail-averaging. The rest of the paper is organized as follows. In Section 2, we describe the least-squares learning problem that we consider, as well as the different SGD variants we analyze. In Section 3, we collect a number of observations shedding light on the role of uniform and tail averaging. In Section 4, we present and discuss our main results. In Section 5 we illustrate our results via some numerical simulations. Proofs and technical results are deferred to the appendices.

## 2 Least Squares Learning with SGD

In this section, we introduce the problem of supervised learning with the least squares loss and then present SGD and its variants.

### 2.1 Least squares learning

We let  $(X, Y)$  be a pair of random variables with values in  $\mathcal{H} \times \mathbb{R}$ , with  $\mathcal{H}$  a real separable Hilbert space. This latter setting is known to be equivalent to nonparametric learning with kernels [31]. We focus on this setting since considering infinite dimensions allows to highlight more clearly the regularization role played by different parameters. Indeed, unlike in finite dimensions, regularization is needed to derive learning rates in this case. Throughout the paper we will suppose that the following assumption holds:

**Assumption 1.** *Assume  $\|X\| \leq \kappa$ ,  $|Y| \leq M$  almost surely, for some  $\kappa, M > 0$ .*

The problem of interest is to solve

$$\min_{w \in \mathcal{H}} \mathcal{L}(w), \quad \mathcal{L}(w) = \frac{1}{2} \mathbb{E}[(Y - \langle w, X \rangle)^2] \quad (1)$$

provided a realization  $x_1, \dots, x_n$  of  $n$  identical copies  $X_1, \dots, X_n$  of  $X$ . Defining

$$\Sigma = \mathbb{E}[X \otimes X], \quad \text{and} \quad h = \mathbb{E}[XY], \quad (2)$$

the optimality condition of problem (1) shows that a solution  $w_*$  satisfies the normal equation

$$\Sigma w_* = h. \quad (3)$$

Finally, recall that the excess risk associated with any  $w \in \mathcal{H}$  can be written as <sup>1</sup>

$$\mathcal{L}(w) - \mathcal{L}(w_*) = \left\| \Sigma^{1/2}(w - w_*) \right\|^2.$$

---

<sup>1</sup>It is a standard fact that the operator  $\Sigma$  is symmetric, positive definite and trace class (hence compact), since  $X$  is bounded. Then fractional powers of  $\Sigma$  are naturally defined using spectral calculus.

## 2.2 Learning with stochastic gradients

We now introduce various gradient iterations relevant in the following. The basic stochastic gradient iteration is given by the recursion

$$w_{t+1} = w_t - \gamma_t x_t (\langle x_t, w_t \rangle - y_t)$$

for all  $t = 0, 1 \dots$ , with  $w_0 = 0$ . For all  $w \in \mathcal{H}$  and  $t = 1, \dots, n$ ,

$$\mathbb{E}[X_t(\langle X_t, w \rangle - Y_t)] = \nabla L(w), \quad (4)$$

hence the name. While the above iteration is not ensured to decrease the objective at each step, the above procedure and its variants are commonly called Stochastic Gradient Descent (SGD). We will also use this terminology. The sequence  $(\gamma_t)_t > 0$ , is called step-size or learning rate. In its basic form, the above iteration prescribes to use each data point only once. This is the classical stochastic approximation perspective pioneered by [30].

In practice, however, a number of different variants are considered. In particular, often times, data points are visited multiple times, in which case we can write the recursion as

$$w_{t+1} = w_t - \gamma_t x_{i_t} (\langle x_{i_t}, w_t \rangle - y_{i_t}).$$

Here  $i_t = i(t)$  denotes a map specifying a strategy with which data are selected at each iteration. Popular choices include: *cyclic*, where an order over  $[n]$  is fixed a priori and data points are visited multiple times according to it; *reshuffling*, where the order of the data points is permuted after all of them have been sampled once, amounting to sampling without replacement; and finally the most common approach, which is sampling each point *with replacement* uniformly at random. This latter choice is also the one we consider in this paper. We broadly refer to this variant of SGD as *multipass-SGD*, referring to the “multiple passes” ‘over the data set as  $t$  grows larger than  $n$ .

Another variant of SGD is based on considering more than one data point at each iteration, a procedure called *mini-batching*. Given  $b \in [n]$  the mini-batch SGD recursion is given by

$$w_{t+1} = w_t + \gamma_t \frac{1}{b} \sum_{i=b(t-1)+1}^{bt} (\langle w_t, x_{j_i} \rangle - y_{j_i}) x_{j_i},$$

where  $j_1, \dots, j_{bT}$  are i.i.d. random variables, distributed according to the uniform distribution on  $[n]$ . Here the number of *passes* over the data after  $t$  iterations is  $\lceil bt/n \rceil$ . Mini-batching can be useful for at least two different reasons. The most important is that considering mini-batches is natural to make the best use of memory resources, in particular when distributed computations are available. Another advantage is that in this case more accurate gradient estimates are clearly available at each step.

Finally, one last idea is considering averaging of the iterates, rather than working with the final iterate,

$$\bar{w}_T = \frac{1}{T} \sum_{t=1}^T w_t.$$

This is a classical idea in optimization, where it is known to provide improved convergence results [2, 15, 27, 32], but it is also used when recovering stochastic results from worst case sequential prediction analysis [17, 33]. More recently, averaging was shown to lead to larger step-sizes, see [2, 10, 11]. In the following, we consider a variant of the above idea, namely *tail-averaging*, where for  $0 \leq S \leq T - 1$  we let

$$\bar{w}_{S,T} = \frac{1}{T - S} \sum_{t=S+1}^T w_t.$$

We will occasionally write  $\bar{w}_L = \bar{w}_{S,T}$ , with  $L = T - S$ . In the following, we study how the above ideas can be combined to solve problem (1) and how such combinations affect the learning properties of the obtained solutions.

### 3 An appetizer: Averaging and Gradient Descent Convergence

Averaging is known to allow larger step-sizes for SGD but also to slower convergence rates in certain settings [10]. In this section, we present calculations shedding light on these effects. In particular, we show how the slower convergence is a completely deterministic effect and how *tail* averaging can provide a remedy. In the rest of the paper, we will build on these reasonings to derive novel quantitative results in terms of learning bounds. The starting observation is that since SGD is based on stochastic estimates of the expected risk gradient (cf. equations (1), (4)) it is natural to start from the exact gradient descent to understand the role played by averaging.

For  $\gamma > 0$ ,  $w_0 = 0$ , consider the population gradient descent iteration,

$$u_t = u_{t-1} - \gamma \mathbb{E}[X(\langle X, u_{t-1} \rangle - Y)] = (I - \gamma \Sigma)u_{t-1} + \gamma h,$$

where the last equality follows from (2). Then using the normal equation (3) and a simple induction argument [12], it is easy to see that,

$$u_T = g_T(\Sigma)\Sigma w_*, \quad g_T(\Sigma) = \gamma \sum_{j=0}^{T-1} (I - \gamma \Sigma)^j. \quad (5)$$

Here,  $g_T$  is a *spectral filtering* function corresponding to a truncated matrix geometric series (the von Neumann series). For the latter to converge, we need  $\gamma$  such that  $\|I - \gamma \Sigma\| < 1$ , e.g.  $\gamma < 1/\sigma_M < 1/\kappa^2$ , with  $\sigma_M = \sigma_{\max}(\Sigma) \leq \kappa^2$ , hence recovering a classical step-size choice. The above computation provides a way to analyze gradient descent convergence. Indeed, one can easily show that

$$w_* - u_T = r_T(\Sigma)w_*, \quad r_T(\Sigma) = (I - \gamma \Sigma)^T$$

since

$$g_T(\Sigma)\Sigma = (I - (I - \gamma \Sigma)^T)w_*$$

from basic properties of the Neumann series defining  $g_T$ .

The properties of the so-called *residual operators*  $r_T(\Sigma)$  control the convergence of GD. Indeed, if  $\sigma_m = \sigma_{\min}(\Sigma) > 0$ , then

$$\left\| \Sigma^{1/2}(u_T - w_*) \right\|^2 = \left\| \Sigma^{1/2}r_T(\Sigma)w_* \right\|^2 \leq \sigma_M(1 - \gamma\sigma_m)^{2T} \|w_*\|^2 \leq \sigma_M e^{-2\sigma_m \gamma T} \|w_*\|^2,$$

from the basic inequality  $1 + z \leq e^z$ , highlighting that the population GD iteration converges exponentially fast to the risk minimizer. However, a major caveat is that assuming  $\sigma_{\min}(\Sigma) > 0$  is clearly restrictive in an infinite dimensional (nonparametric) setting, since it effectively implies that  $\Sigma$  has finite rank. In general,  $\Sigma$  will not be finite rank, but rather compact with 0 as the only accumulation point of its spectrum. In this case, it is easy to see that the slower rate

$$\left\| \Sigma^{1/2}(u_T - w_*) \right\|^2 \leq \frac{1}{\gamma T} \|w_*\|^2$$

holds without any further assumption on the spectrum, since one can show, using spectral calculus and a direct computation <sup>2</sup>, that  $s^{1/2}r_T(s) \leq 1/\gamma T$ . It is reasonable to ask whether it is possible to interpolate between the above-described slow and fast rates by making some

---

<sup>2</sup>Setting  $\frac{d}{ds}s(1 - \gamma s)^T = 0$  gives  $1 - \gamma s - s\gamma T = 0 \Rightarrow s = \frac{1}{\gamma(T+1)}$  and  $\frac{1}{\gamma(T+1)} \left(1 - \gamma \frac{1}{\gamma(T+1)}\right)^t \leq \frac{1}{\gamma t}$ .

intermediate assumption. Rather than making assumption on the spectrum of  $\Sigma$ , one can assume the optimal solution  $w_*$  to belong to a subspace of the range of  $\Sigma$ , more precisely that

$$w_* = \Sigma^r v_* \tag{6}$$

holds for some  $r \geq 0$  and  $v_* \in \mathcal{H}$ , where larger values of  $r$  correspond to making more stringent assumptions. In particular, as  $r$  goes to infinity we are essentially assuming  $w_*$  to belong to a finite dimensional space. Assumption (6) is common in the literature of inverse problems [12] and statistical learning [8, 9]. Interestingly, it is also related to so-called conditioning and Łojasiewicz conditions, known to lead to improved rates in continuous optimization, see [13] and references therein. Under assumption (6), and using again spectral calculus, it is possible to show that, for all  $r \geq 0$ ,

$$\left\| \Sigma^{1/2}(u_T - w_*) \right\|^2 = \left\| \Sigma^{1/2} r_T(\Sigma) \Sigma^r v_* \right\|^2 \lesssim \left( \frac{1}{\gamma T} \right)^{2r+1} \|v_*\|^2.$$

Thus, higher values of  $r$  result in faster convergence rates, at the price of more stringent assumptions.

### 3.1 Tail averaged gradient descent

Given the above discussion, we can derive analogous computations for (tail) averaged GD and draw some insights. Using (5), for  $S < T$ , we can write the tail-averaged gradient

$$\bar{u}_{S,T} = \frac{1}{T-S} \sum_{t=S+1}^T u_t \tag{7}$$

as

$$\bar{u}_{S,T} = G_{S,T}(\Sigma) \Sigma w_*, \quad G_{S,T}(\Sigma) = \frac{1}{T-S} \sum_{j=S+1}^T g_j(\Sigma). \tag{8}$$

As before, we can analyze convergence considering a suitable residual operator

$$w_* - u_{S,T} = R_{S,T}(\Sigma) w_*, \quad R_{S,T}(\Sigma) = I - G_{S,T}(\Sigma) \Sigma \tag{9}$$

which, in this case, can be shown to take the form,

$$R_{S,T}(\Sigma) = \frac{(I - \gamma \Sigma)^{S+1}}{\gamma(T-S)} (I - (I - \gamma \Sigma)^{T-S}) \Sigma^{-1}$$

and where with an abuse of notation we denote by  $\Sigma^{-1}$  the pseudoinverse of  $\Sigma$ . The case of uniform averaging corresponds to  $S = 0$ , in which case the residual operator simplifies to

$$R_{0,T}(\Sigma) = \frac{(I - \gamma \Sigma)}{\gamma T} (I - (I - \gamma \Sigma)^T) \Sigma^{-1}.$$

When  $\sigma_m > 0$ , the residual operators behave roughly as

$$\|R_{S,T}(\Sigma)\|^2 \approx \frac{e^{-\sigma_m \gamma(S+1)}}{\gamma(T-S)}, \quad \|R_{0,T}(\Sigma)\|^2 \approx \frac{1}{\gamma T},$$

respectively. This leads to a slower convergence rate for uniform averaging and shows instead how tail averaging with  $S \propto T$  can preserve the fast convergence of GD.

When  $\sigma_m = 0$ , taking again  $S \propto T$ , it is easy to see by spectral calculus that the residual operators behave similarly,

$$\left\| \Sigma^{1/2} R_{S,T}(\Sigma) \right\|^2 \approx \frac{1}{\gamma T}, \quad \left\| \Sigma^{1/2} R_{0,T}(\Sigma) \right\|^2 \approx \frac{1}{\gamma T},$$

leading to comparable rates. The advantage of tail averaging is again apparent if we consider Assumption (6). In this case for all  $r > 0$ , if we take  $S \propto T$

$$\left\| \Sigma^{1/2} R_{S,T}(\Sigma) \Sigma^r \right\|^2 \approx \left( \frac{1}{\gamma T} \right)^{2r+1}, \quad (10)$$

whereas with uniform averaging one can only prove

$$\left\| \Sigma^{1/2} R_{0,T}(\Sigma) \Sigma^r \right\|^2 \approx \left( \frac{1}{\gamma T} \right)^{2 \min(r, 1/2) + 1}. \quad (11)$$

One immediate observation following from the above discussion is that uniform averaging induces a so-called *saturation effect* [12], meaning that the rates do not improve after  $r$  reaches a critical point. As shown above, this effect vanishes considering tail-averaging and the convergence rate of GD is recovered. These results are critically important for our analysis and constitute the main conceptual contribution of our paper. They are proved in Appendix B, while Section A.1 highlights their critical role for SGD. To the best of our knowledge, we are the first to highlight this acceleration property of tail averaging beyond the finite-dimensional setting.

## 4 Main Results and Discussion

In this section we present and discuss our main results. We start by presenting a general bound and then use it to derive the optimal parameter settings and corresponding performance guarantees. A key quantity in our results will be the *effective dimension*

$$\mathcal{N}(1/\gamma L) = \text{Tr} \left[ \left( \Sigma + \frac{1}{\gamma L} \right)^{-1} \Sigma \right],$$

introduced in [38] to generalize results from parametric estimation problems to non-parametric kernel methods. Similarly this will be one of the main quantities in our learning bounds.

Further, in all our results we will require that the stepsize is bounded as  $\gamma \kappa^2 < 1/4$ , and that the tail length  $L = T - S$  is scaled appropriately with the total number of iterations  $T$ . More precisely, our analysis considers two different scenarios where  $S = 0$  (plain averaging) is explicitly allowed and where  $S > 0$ , i.e., where we investigate the merits of tail-averaging. To do so, we will assume  $0 \leq S \leq \frac{K-1}{K+1} T$  for some  $1 \leq K$ , and also  $T \leq (K+1)S$  for the latter case.

The following theorem presents a simplified version of our main technical result that we present in its general form in the Appendix. Here, we omit constants and lower order terms for clarity and give the first insights into the interplay between the tuning parameters, namely the step-size  $\gamma$ , tail-length  $L$ , and mini-batch size  $b$ , and the number of points  $n$ . Note that in a nonparametric setting these are the quantities controlling learning rates. The following result provides a bound for any choice of the tuning parameters, and will allow to derive optimal choices balancing the various error contributions.

**Theorem 1.** *Let  $\alpha \in (0, 1]$ ,  $1 \leq L \leq T$  and let Assumption 1 hold. Assume  $\gamma \kappa^2 < 1/4$  as well as  $n \gtrsim \gamma L \mathcal{N}(1/\gamma L)$ . Then, the excess risk of the tail-averaged SGD iterates satisfies*

$$\mathbb{E} \left[ \left\| \Sigma^{\frac{1}{2}} (\bar{w}_L - w_*) \right\|^2 \right] \lesssim \left\| \Sigma^{1/2} R_L(\Sigma) w_* \right\|^2 + \frac{\mathcal{N}(1/\gamma L)}{n} + \frac{\gamma \text{Tr}[\Sigma^\alpha]}{b(\gamma L)^{1-\alpha}}.$$

The proof of the result is given in Appendix E. We make a few comments. The first term in the bound is the approximation error, already discussed in Section 3. It is controlled by the bound in (10) and which is decreasing in  $\gamma L$ . The second term corresponds to a variance

error due to sampling and noise in the data. It depends on the effective dimension which is increasing in  $\gamma L$ . The third term is a computational error due to the randomization in SGD. Note how it depends on both  $\gamma L$  and the minibatch size  $b$ . The larger  $b$  is, the smaller this error becomes. The dependence of all three terms on  $\gamma L$  suggest already at this stage that  $(\gamma L)^{-1}$  plays the role of a regularization parameter. We derive our final bound by balancing all terms, i.e. choosing them to be of the same order. To do so we make additional assumptions. The first one is Eq. (6), enforcing the optimal solution  $w_*$  to belong to a subspace of the range of  $\Sigma$ .

**Assumption 2.** For some  $r \geq 0$  we assume  $w_* = \Sigma^r v_*$ , for some  $v_* \in \mathcal{H}$  satisfying  $\|v_*\| \leq R$ .

The larger is  $r$  the more stringent is the assumption, or, equivalently, the easier is the problem, see Section 3. A second further assumption is related to the effective dimension.

**Assumption 3.** For some  $\nu \in (0, 1]$  and  $C_\nu < \infty$  we assume  $\mathcal{N}(1/\gamma L) \leq C_\nu(\gamma L)^\nu$ .

This assumption is common in the nonparametric regression setting, see e.g [6]. Roughly speaking, it quantifies how far  $\Sigma$  is from being finite rank. Indeed, it is satisfied if the eigenvalues  $(\sigma_i)_i$  of  $\Sigma$  have a polynomial decay  $\sigma_i \sim i^{-\frac{1}{\nu}}$ . Since  $\Sigma$  is trace class, the assumption is always satisfied for  $\nu = 1$  with  $C_\nu = \kappa^2$ . Smaller values of  $\nu$  lead to faster convergence rates.

The following corollary of Theorem 1, together with Assumptions 2 and 3, derives optimal parameter settings and corresponding learning rates.

**Corollary 1.** Let all assumptions of Theorem 1 be satisfied, and suppose that Assumptions 2, 3 also hold. Further, assume either

1.  $0 \leq r \leq 1/2$ ,  $1 \leq L \leq T$  (here  $S = 0$ , i.e., full averaging is allowed) or
2.  $1/2 < r$ ,  $1 \leq L < T$  with the additional constraint that for some  $K \geq 2$

$$\frac{K+1}{K-1}S \leq T \leq (K+1)S,$$

(only tail-averaging is considered).

Then, for any  $n$  sufficiently large, the excess risk of the (tail)-averaged SGD iterate satisfies

$$\mathbb{E} \left[ \left\| \Sigma^{\frac{1}{2}} (\bar{w}_{L_n} - w_*) \right\|^2 \right] \lesssim n^{-\frac{2r+1}{2r+1+\nu}}$$

for each of the following choices:

- (a)  $b_n \simeq 1$ ,  $L_n \simeq n$ ,  $\gamma_n \simeq n^{-\frac{2r+\nu}{2r+1+\nu}}$  (one pass over data)
- (b)  $b_n \simeq n^{\frac{2r+\nu}{2r+1+\nu}}$ ,  $L_n \simeq n^{\frac{1}{2r+1+\nu}}$ ,  $\gamma_n \simeq 1$  (one pass over data)
- (c)  $b_n \simeq n$ ,  $L_n \simeq n^{\frac{1}{2r+1+\nu}}$ ,  $\gamma_n \simeq 1$  ( $n^{\frac{1}{2r+1+\nu}}$  passes over data).

The proof of Corollary 1 is given in Appendix E. It gives optimal rates [5, 6] under different assumptions and choices for the stepsize  $\gamma$ , the minibatch size  $b$  and the tail length  $L$ , considered as functions of  $n$  and the parameters  $r$  and  $\nu$  from Assumptions 2, 3. We now discuss our findings in more detail and compare them to previous related work.

**Optimality of the bound:** The above results show that different parameter choices allow to achieve the same error bound. The latter is known to be optimal in minmax sense, see e.g. [6]. As noted before, here we provide simplified statements highlighting the dependence of the bound on the number of points  $n$  and the parameters  $r$  and  $\nu$  that control the regularity of the problem. These are quantities controlling the learning rates and for which lower bounds are available. Note however, that all the constants in the Theorem are worked out and reported in detail in the Appendices.

**Regularization properties of tail-length:** We recall that for GD it is well known that  $(\gamma T)^{-1}$  serves as a regularization parameter, having a quantitatively similar effect to Tikhonov regularization with parameter  $\lambda > 0$ , see e.g. [12]. More generally, our result shows that in the case of tail averaging the quantity  $(\gamma L)^{-1}$  becomes the regularizing parameter for *both GD and SGD*.

**The benefit of tail-averaging:** For SGD with  $b = 1$  and full averaging it has been shown by [10] that a single pass over data (i.e.,  $T_n = n$ ) gives optimal rates of convergence provided that  $\gamma_n$  is chosen as in case (a) in the corollary. However the results in [10] held only in the case  $r \leq 1/2$ . Indeed, beyond this regime, there is a saturation effect which precludes optimality for higher smoothness, see the discussion in Section 3, eq. (11). Our analysis for case (a) shows that optimal rates for  $r \geq 0$  can still be achieved with the same number of passes and step-size *by using non-trivial tail averaging*. Additionally, we compare our results with those from [26]. In that paper it is shown that multi-passes are beneficial for obtaining improved rates for averaged SGD in a regime where the optimal solution  $w^*$  does not belong to  $\mathcal{H}$  (Assumption 2 does not hold in that case). In that regime, tail-averaging does not improve convergence. Our analysis focuses on the “opposite” regime where  $w^* \in \mathcal{H}$  and saturation slows down the convergence of uniformly-averaged SGD, preventing optimal rates. Here, tail-averaging is indeed beneficial and leads to improved rates.

**The benefit of multi-passes and mini-batching:** We compare our results with those in [21] where no averaging but mini-batching is considered. In particular, there it is shown that a relatively large stepsize of order  $\log(n)^{-1}$  can be chosen provided the minibatch size is set to  $n^{\frac{2r+1}{2r+1+\nu}}$  and a number of  $n^{\frac{1}{2r+1+\nu}}$  passes is considered. Comparing to these results we can see the benefits of combining minibatching with tail averaging. Indeed from (c) we see that with a comparable number of passes, we can use a larger, constant step-size already with a much smaller minibatch size. Further, comparing (b) and (c) we see that the setting of  $\gamma$  and  $L$  is the same and there is a full range of possible values for  $b_n$  between  $[n^{\frac{2r+\nu}{2r+1+\nu}}, n]$  where a constant stepsize is allowed, still ensuring optimality. As noted in [21], increasing the minibatch size beyond a critical value does not yield any benefit. Compared to [21], we show that that tail-averaging can lead to a much smaller critical minibatch size, and hence more efficient computations.

**Comparison to finite-dimensional setting:** The relationship between the step-size and batch size in finite dimensions  $\dim \mathcal{H} = d < \infty$  is derived in [19] where also tail-averaging but only one pass over the data is considered. One of the main contributions of this work is characterizing the largest stepsize that allows achieving statistically optimal rates, showing that the largest permissible stepsize grows linearly in  $b$  before hitting a certain quantity  $b_{\text{thresh}}$ . Setting  $b > b_{\text{thresh}}$  results in loss of computational and statistical efficiency: in this regime, each step of minibatch SGD is exactly as effective in decreasing the bias as a step of batch gradient descent. The critical value  $b_{\text{thresh}}$  and the corresponding largest admissible stepsize is problem dependent and does not depend on the sample size  $n$ . Notably, the statistically optimal rate of order  $\sigma^2 d/n$  is achieved for all *constant* minibatch sizes, and the particular choice of  $b$  only impacts the constants in the decay rate of the bias (which is of the lower order  $1/n^2$  anyway). That is, choosing the right minibatch size does not involve a tradeoff between statistical and optimization error. In contrast, our work shows that setting a large batch size  $b_n \simeq n^\alpha$ ,  $\alpha \in [0, 1]$  yields optimality guarantees in the infinite dimensional setting. This is due to the fact that choosing the optimal values for parameters like  $\gamma$  and  $b$  involve a tradeoff between the bias and the variance in this setting. [19] also show that tail-averaging improves the rate at which the initial bias decays if the smallest eigenvalue of the covariance matrix  $\sigma_{\min}(\Sigma)$  is lower-bounded by a constant. Their analysis of this algorithmic component






Figure 1: Illustration of the effects of tail-averaging and minibatching. (a) Excess risk as a function of  $r$  with uniform and tail averaging. (b) Excess risk as a function of stepsize  $\gamma$  and minibatch-size  $b$  for SGD with uniform averaging. (c) Excess risk as a function of stepsize  $\gamma$  and minibatch-size  $b$  for SGD with tail-averaging.

is based on observations similar to the ones we made in Section 3. Our analysis significantly extends these arguments by showing the usefulness of tail-averaging in cases when  $\sigma_{\min}$  is not necessarily lower-bounded.

## 5 Numerical Illustration

This section provides an empirical illustration to the effects characterized in the previous sections. We focus on two aspects of our results: the benefits of tail-averaging over uniform averaging as a function of the smoothness parameter  $r$ , and the impact of tail-averaging on the best choice of minibatch sizes. All experiments are conducted on synthetic data with  $d = 1,000$  dimensions, generated as follows. We set  $\Sigma$  as a diagonal matrix with entries  $\Sigma_{ii} = i^{-1/\nu}$  and choose  $w^* = \Sigma^r e$ , where  $e$  is a vector of all 1's. The covariates  $X_t$  are generated from a Gaussian distribution with covariance  $\Sigma$ , and labels are generated as  $Y_t = \langle w^*, X_t \rangle + \varepsilon_t$ , where  $\varepsilon_t$  is standard Gaussian noise. For all experiments, we choose  $\nu = 1/2$  and  $n = 10,000$ . With this choice of parameters, we have seen that increasing  $d$  beyond 100 does not yield any noticeable change in the results, indicating that setting  $d = 1,000$  is an appropriate approximation to the infinite-dimensional setting.

Our first experiment illustrates the saturation effect described in Section 3 (cf. Eqs. 10,11) by plotting the respective excess risks of uniformly-averaged and tail-averaged SGD as a function of  $r$  (Figure 1(a)). We fix  $b = 1$  and set  $\gamma = n^{-\frac{2r+\nu}{2r+1+\nu}}$  as recommended in Corollary 1. As predicted by our theoretical results, the two algorithms behave similarly for smaller values of  $r$ , but uniformly-averaged SGD noticeably starts to lag behind its tail-averaged counterpart for larger values of  $r$  exceeding  $1/2$ , eventually flattening out and showing no improvement as  $r$  increases. On the other hand, the performance of the tail-averaged version continues to improve for large values of  $r$ , confirming that this algorithm can indeed massively benefit from favorable structural properties of the data.

In our second experiment, we study the performance of both tail- and uniformly-averaged SGD as a function of the stepsize  $\gamma$  and the minibatch-size  $b$  (Figure 1(b), (c)). We fix  $r = 1/2$  and set  $T = n/b$  for all tested values of  $b$ , amounting to a single pass over the data. Again, as theory predicts, performance remains largely constant as  $\gamma \cdot b$  remains constant for both algorithms, until a critical threshold stepsize is reached. However, it is readily apparent from the figures that tail-averaging permits the use of larger minibatch sizes, therefore allowing for more efficient parallelization.

## Acknowledgments

NM is supported by the German Research Foundation under DFG Grant STE 1074/4-1. L. R. acknowledges the financial support of the AFOSR projects FA9550-17-1-0390 and BAA-AFRL-AFOSR-2016-0007 (European Office of Aerospace Research and Development), and the EU H2020-MSCA-RISE project NoMADS - DLV-777826.

## References

- [1] R. Aguech, E. Moulines, and P. Priouret. On a perturbation approach for the analysis of stochastic tracking algorithms. *SIAM J. Control and Optimization*, 39 (3):872–899, 2000.
- [2] Francis R. Bach and Eric Moulines. Non-strongly-convex smooth stochastic approximation with convergence rate  $o(1/n)$ . In *NIPS*, pages 773–781, 2013.
- [3] Frank Bauer, Sergei Pereverzev, and Lorenzo Rosasco. On regularization algorithms in learning theory. *Journal of Complexity*, 23(1):52 – 72, 2007.
- [4] Dimitri P. Bertsekas. A new class of incremental gradient methods for least squares problems. *SIAM Journal on Optimization*, 7(4):913–926, 1997.
- [5] Gilles Blanchard and Nicole Mücke. Optimal rates for regularization of statistical inverse learning problems. *Foundations of Computational Mathematics*, 18(4):971–1013, Aug 2017.
- [6] Andrea Caponnetto and E. De Vito. Optimal rates for regularized least-squares algorithm. *Foundations of Computational Mathematics*, 7(3):331–368, 2006.
- [7] Andrew Cotter, Ohad Shamir, Nati Srebro, and Karthik Sridharan. Better mini-batch algorithms via accelerated gradient methods. In J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 1647–1655. Curran Associates, Inc., 2011.
- [8] F. Cucker and S. Smale. On the mathematical foundations of learning. *Bulletin of the American mathematical society*, 39(1):1–49, 2002.
- [9] E. De Vito, A. Caponnetto, and L. Rosasco. Model selection for regularized least-squares algorithm in learning theory. *Foundations of Computational Mathematics*, 5(1):59–85, 2005.
- [10] Aymeric Dieuleveut and Francis Bach. Nonparametric stochastic approximation with large step-sizes. *Ann. Statist.*, 44(4):1363–1399, 08 2016.
- [11] Aymeric Dieuleveut, Nicolas Flammarion, and Francis Bach. Harder, better, faster, stronger convergence rates for least-squares regression. *Journal of Machine Learning Research*, 18:101:1–101:51, 2017.
- [12] Heinz Werner Engl, Martin Hanke, and Andreas Neubauer. *Regularization of inverse problems*, volume 375. Springer Science & Business Media, 1996.
- [13] G. Garrigos, L. Rosasco, and S. Villa. Convergence of the forward-backward algorithm: Beyond the worst case with the help of geometry. *arXiv:1703.09477*, 2017.
- [14] Zheng-Chu Guo, Shao-Bo Lin, and Ding-Xuan Zhou. Learning theory of distributed spectral algorithms. *Inverse Problems*, 33(7):074009, 2017.

- [15] László Györfi and Harro Walk. On the averaged stochastic approximation for linear regression. *SIAM Journal on Control and Optimization*, 34(1):31–61, 1996.
- [16] Moritz Hardt, Ben Recht, and Yoram Singer. Train faster, generalize better: Stability of stochastic gradient descent. In *ICML*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 1225–1234. JMLR.org, 2016.
- [17] Elad Hazan. Introduction to online convex optimization. *Foundations and Trends in Optimization*, 2(3-4):157–325, 2016.
- [18] Elad Hazan and Satyen Kale. Beyond the regret minimization barrier: optimal algorithms for stochastic strongly-convex optimization. *The Journal of Machine Learning Research*, 15(1):2489–2512, 2014.
- [19] Prateek Jain, Sham M. Kakade, Rahul Kidambi, Praneeth Netrapalli, and Aaron Sidford. Parallelizing stochastic gradient descent for least squares regression: Mini-batching, averaging, and model misspecification. *Journal of Machine Learning Research*, 18(223):1–42, 2018.
- [20] Junhong Lin, Raffaello Camoriano, and Lorenzo Rosasco. Generalization properties and implicit regularization for multiple passes SGM. *CoRR*, abs/1605.08375, 2016.
- [21] Junhong Lin and Lorenzo Rosasco. Optimal rates for multi-pass stochastic gradient methods. *Journal of Machine Learning Research*, 18:97:1–97:47, 2017.
- [22] Junhong Lin, Alessandro Rudi, Lorenzo Rosasco, and Volkan Cevher. Optimal rates for spectral algorithms with least-squares regression over hilbert spaces. *Applied and Computational Harmonic Analysis*, 2018.
- [23] Arkadi Nemirovski, Anatoli Juditsky, Guanghui Lan, and Alexander Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.
- [24] Gergely Neu and Lorenzo Rosasco. Iterate averaging as regularization for stochastic gradient descent. In *COLT*, volume 75 of *Proceedings of Machine Learning Research*, pages 3222–3242. PMLR, 2018.
- [25] Francesco Orabona. Simultaneous model selection and optimization through parameter-free stochastic learning. In *Advances in Neural Information Processing Systems*, pages 1116–1124, 2014.
- [26] Loucas Pillaud-Vivien, Alessandro Rudi, and Francis Bach. Statistical optimality of stochastic gradient descent on hard learning problems through multiple passes. *CoRR*, abs/1805.10074, 2018.
- [27] B. T. Polyak and A. B. Juditsky. Acceleration of stochastic approximation by averaging. *SIAM J. Control Optim.*, 30(4):838–855, jul 1992.
- [28] A. Rakhlin, O. Shamir, and K. Sridharan. Making gradient descent optimal for strongly convex stochastic optimization. *arXiv:1109.5647*, 2011.
- [29] Alexander Rakhlin, Ohad Shamir, and Karthik Sridharan. Making gradient descent optimal for strongly convex stochastic optimization. In *Proceedings of the 29th International Conference on Machine Learning (ICML)*, pages 1571–1578, 2012.
- [30] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.

- [31] Lorenzo Rosasco and Silvia Villa. Learning with incremental iterative regularization. In *NIPS*, pages 1630–1638, 2015.
- [32] David Ruppert. Efficient estimations from a slowly convergent Robbins–Monro process. Technical report, Cornell University Operations Research and Industrial Engineering, 1988.
- [33] Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2012.
- [34] o. Shamir and T. Zhang. Stochastic gradient descent for non-smooth optimization: convergence results and optimal averaging schemes. In *Proceedings of the 30th International Conference on Machine Learning*, 2013.
- [35] Steve Smale and Yuan Yao. Online learning algorithms. *Foundations of Computational Mathematics*, 6(2):145–170, 2006.
- [36] Pierre Tarres and Yuan Yao. Online learning as stochastic approximation of regularization paths: Optimality and almost-sure convergence. *IEEE Trans. Information Theory*, 60(9):5716–5735, 2014.
- [37] Yiming Ying and Massimiliano Pontil. Online gradient descent learning algorithms. *Foundations of Computational Mathematics*, 8(5):561–596, 2008.
- [38] T. Zhang. Effective dimension and generalization of kernel learning. *Advances in Neural Information Processing Systems 2003*, 2003.

## A Analysis

This section presents the key components of the proofs of our main results. Recall that the goal of the analysis is to understand the rate at which the tail-averaged SGD iterates  $\bar{w}_{S,T}$  approach the risk minimizer  $w_*$ . The main error decomposition underlying our proofs is borrowed from [21], and based on introducing two intermediate objects that will be shown to converge towards  $w_*$ , yet stay close to the SGD iterates  $w_t$ . In Section 3 we have already introduced one of these components: population GD. We will further need the *empirical* (batch) GD iteration, defined as

$$v_{t+1} = v_t - \gamma \frac{1}{n} \sum_{j=1}^n (\langle v_t, x_j \rangle_{\mathcal{H}} - y_j) x_j = (I - \gamma \widehat{\Sigma}) v_t - \gamma \hat{h}, \quad (12)$$

where we also introduced the important notations

$$\widehat{\Sigma} = \frac{1}{n} \sum_{j=1}^n x_j \otimes x_j \quad \hat{h} = \frac{1}{n} \sum_{j=1}^n x_j y_j.$$

Analogously to the tail-averaged SGD/GD we define the tail-averaged batch GD iterates

$$\bar{v}_{S,T} = \frac{1}{T-S} \sum_{t=S+1}^T v_t,$$

which will act as our proxy to  $\bar{w}_{S,T}$ . With these definitions in place, we can upper bound the excess risk of  $\bar{w}_{S,T}$  as

$$\left\| \Sigma^{1/2}(w_* - \bar{w}_{S,T}) \right\|^2 \leq 2 \left\| \Sigma^{1/2}(w_* - \bar{v}_{S,T}) \right\|^2 + 2 \left\| \Sigma^{1/2}(\bar{v}_{S,T} - \bar{w}_{S,T}) \right\|^2. \quad (13)$$

The purpose of this decomposition is to help us separate the inherent statistical errors due to using an i.i.d. sample of fixed size  $n$  (first term) and the errors introduced by the randomized algorithm (second term). Accordingly, we will refer to this latter term as the *computational variance*. In the sections below, we give bounds on both terms separately.

### A.1 Learning properties of GD with tail averaging

In this section, we discuss how to bound the first term in the decomposition of Equation (13). In analogy to the discussion in Section 3, we rewrite the empirical GD using spectral filtering functions,

$$v_{t+1} = g_{t+1}(\widehat{\Sigma}) \hat{h}, \quad g_{t+1}(\widehat{\Sigma}) = \gamma \sum_{j=0}^t (I - \gamma \widehat{\Sigma})^j \quad (14)$$

With this notation, the tail-averaged GD iterates can be written as

$$\bar{v}_{T,S} = G_{S,T}(\widehat{\Sigma}) \hat{h}, \quad G_{S,T}(\sigma) = \frac{1}{T-S} \sum_{t=S+1}^T g_t(\sigma). \quad (15)$$

Most of the analysis in this section will rely on the regularization properties of the spectral filter  $G_{S,T}(\widehat{\Sigma})$ , the corresponding residual operators

$$R_{S,T}(\widehat{\Sigma}) = 1 - \widehat{\Sigma} G_{S,T}(\widehat{\Sigma}), \quad (16)$$

and the analogous population quantities introduced in Section 3. Denoting the tail-length by  $L = T - S$ , we will occasionally use the notations  $G_L = G_{S,T}$  and  $R_L = R_{S,T}$ .

Our error bounds are derived by means of a classical error decomposition in bias and variance (see, e.g., 6, 3, 5 and 22). Recalling the definition of the averaged population GD in Equations (7) and (8), we consider the decomposition

$$\begin{aligned}\bar{v}_L - w_* &= (\bar{v}_L - \bar{u}_L) + (\bar{u}_L - w_*) = (\bar{v}_L - \bar{u}_L) + R_L(\Sigma)w_* \\ &= (G_L(\hat{\Sigma})\hat{h} - G_L(\hat{\Sigma})\hat{\Sigma}\bar{u}_L) + (G_L(\hat{\Sigma})\hat{\Sigma}\bar{u}_L - \bar{u}_L) + R_L(\Sigma)w_* \\ &= G_L(\hat{\Sigma})(\hat{h} - \hat{\Sigma}\bar{u}_L) + R_L(\hat{\Sigma})\bar{u}_L + R_L(\Sigma)w_* .\end{aligned}\tag{17}$$

We refer to

$$\mathcal{A}(L) = \|\Sigma^{1/2}R_L(\Sigma)w_*\|^2\tag{18}$$

as the *deterministic approximation error*, to

$$\hat{\mathcal{A}}(L) = \|\Sigma^{1/2}R_L(\hat{\Sigma})\bar{u}_L\|^2\tag{19}$$

as the *stochastic approximation error* and to

$$\hat{\mathcal{V}}(L) = \|\Sigma^{1/2}G_L(\hat{\Sigma})(\hat{h} - \hat{\Sigma}\bar{u}_L)\|^2\tag{20}$$

as the *sample variance*. Our analysis will crucially rely on the properties of the residual operator  $R_{S,T}$  already discussed in Section 3. Here we show that these arguments made about population GD also impacts the learning error for empirical GD in the same qualitative way. More precisely, Propositions 1 and 2 in Appendices C.1 and C.2 show that, under appropriate conditions, the (expected) approximation errors can be bounded as

$$\mathcal{A}(L) \lesssim \begin{cases} R^2(\gamma L)^{-2(r+1/2)} & \text{if } r \leq \frac{1}{2}, \\ K^2 R^2(\gamma L)^{-2(r+1/2)} & \text{else, and} \end{cases} \quad \mathbb{E}[\hat{\mathcal{A}}(L)] \lesssim \begin{cases} R^2(\gamma L)^{-2(r+1/2)} & \text{if } r \leq \frac{1}{2}, \\ K^{4(r+1)} R^2(\gamma L)^{-2(r+1/2)} & \text{else.} \end{cases}$$

Notably, proving this result for  $r > 1/2$  critically relies on setting  $S$  as a *constant* fraction of  $T$  that enables the rapid decay of  $R_{S,T}$  in  $S$ , highlighting the important role of tail averaging to obtain these results. The precise condition we require is  $S \leq \frac{K-1}{K+1}T$  and  $T \leq (K+1)T$  to hold for some constant  $K > 1$  see Corollary 1. Regarding the sample variance, Proposition 4 in Appendix C.3 shows the bound

$$\mathbb{E}[\hat{\mathcal{V}}(L)] \lesssim \mathcal{A}(L) + \frac{\gamma L(1 + \|w_*\|^2)}{n^2} + \frac{\mathcal{N}(1/\gamma L)}{n}.$$

Putting these results together, we can conclude that the excess risk of tail-averaged GD satisfies the bound

$$\mathbb{E}\left[\left\|\Sigma^{1/2}(w_* - \bar{v}_{S,T})\right\|^2\right] \lesssim R^2(\gamma L)^{-2(r+1/2)} + \frac{\gamma L(1 + \|w_*\|^2)}{n^2} + \frac{\mathcal{N}(1/\gamma L)}{n}$$

whenever  $K$  is set as  $O(1)$ . The precise bound is stated in Appendix C.4 as Theorem 2. A particularly important consequence of this result is that, under the additional Assumption 3, the excess-risk bound can be further rewritten as

$$\mathbb{E}\left[\left\|\Sigma^{1/2}(\bar{v}_L - w_*)\right\|^2\right] \leq R^2 C_K n^{-\frac{2r}{2r+1+\nu}},$$

when choosing  $\gamma_n \simeq n^{-a}$  and  $T \simeq n^{\tilde{a}}$  for some  $a, \tilde{a} > 0$  satisfying  $a - \tilde{a} = \frac{1}{2r+1+\nu}$ . Once again, these results rely on choosing  $T \simeq S$  in the case  $r > 1/2$ , whereas choosing  $S = 0$  is sufficient for the case  $r \leq 1/2$ . This result is formally stated as Corollary 2 in Appendix C.4.

In the low smoothness regime, i.e.  $0 \leq r \leq 1/2$ , the choice  $0 < S, S_n \asymp T_n$  is also possible but does not affect the rate of convergence, whereas in the high smoothness regime, i.e.  $1/2 < r$ , tail averaging is necessary to avoid saturation.

## A.2 SGD vs. GD

We now move on to analyze the difference between the tail-averaged SGD and GD iterates and provide a bound on the second term in the decomposition (13). To relate the two iterations, we introduce the notation

$$\widehat{\Sigma}_t = \frac{1}{b} \sum_{i=b(t-1)+1}^{bt} x_{j_i} \otimes x_{j_i} \quad \text{and} \quad \widehat{h}_t = \frac{1}{b} \sum_{i=b(t-1)+1}^{bt} y_{j_i} x_{j_i},$$

so that the minibatch SGD iteration can be written as  $w_{t+1} = (I - \gamma \widehat{\Sigma}_t)w_t - \gamma \widehat{h}_t$ . Thus, the difference between the two iterate sequences can be written in the recursive form

$$w_{t+1} - v_{t+1} = \left( I - \gamma \widehat{\Sigma}_{t+1} \right) (w_t - v_t) + \gamma \xi_{t+1}, \quad (21)$$

where  $\xi_{t+1} = \xi_{t+1}^{(1)} + \xi_{t+1}^{(2)}$  and

$$\xi_{t+1}^{(1)} = (\widehat{\Sigma} - \widehat{\Sigma}_{t+1})v_t, \quad \xi_{t+1}^{(2)} = \widehat{h}_{t+1} - \widehat{h}. \quad (22)$$

It is easy to see that  $\xi_{t+1}$  has zero mean when conditioned on the history  $\mathcal{F}_t$  and the dataset. Notice that the recursion above is of the form

$$\mu_{t+1} = (I - \gamma \widehat{H}_{t+1})\mu_t + \gamma \zeta_{t+1}$$

for i.i.d. self-adjoint positive operators  $\widehat{H}_t$  satisfying  $\mathbb{E}[\widehat{H}_t | \mathcal{F}_t] = H$  and  $\mathbb{E}[\zeta_t | \mathcal{F}_t] = 0$ . Such recursions have been well studied in the stochastic approximation literature, and can be analyzed by techniques proposed by [1] (and later used by [2, 10, 26], among many others). Our analysis builds on a recent result by [26] that we generalize to account for minibatching and tail-averaging. On a high level, this result states that if  $\widehat{H}_t$  and  $\zeta_t$  respectively satisfy  $\mathbb{E}[\widehat{H}_t^2 | \mathcal{F}_t] \leq \kappa^2 H$  and  $\mathbb{E}[\zeta_t \otimes \zeta_t | \mathcal{F}_t] \preceq \sigma^2 \widehat{\Sigma}$  for some  $\kappa, \sigma$ , then the tail-averaged iterate  $\bar{\mu}_{S,T} = \frac{1}{T-S} \sum_{t=S+1}^T \mu_t$  satisfies

$$\mathbb{E} \left[ \left\| H^{\frac{\alpha}{2}} \bar{\mu}_{S,T} \right\|^2 \right] \lesssim \sigma^2 \text{Tr}[H^\alpha] \gamma^{1-u+\alpha} (T-S)^{\alpha-u} \frac{S+1}{T-S}$$

for arbitrary  $\alpha \in (0, 1]$  and  $u \in [0, 1 + \alpha]$ . Appendix D is dedicated to formally proving this result, presented precisely as Proposition 5.

Our analysis crucially relies on applying the above lemma for  $H = \widehat{\Sigma}$  under an appropriately defined condition  $\mathcal{E}_1$  on the data, see (68), guaranteeing that  $\widehat{\Sigma}$  is “close enough” to its population counterpart  $\Sigma$ . A second condition  $\mathcal{E}_2$  in (69) ensures the boundedness of  $\mathbb{E}[\xi_t \otimes \xi_t | \mathcal{F}_t, \mathcal{G}_n]$ , conditioned on the data  $\mathcal{G}_n$ . We note that, ensuring the condition about  $\xi_t$  is rather challenging due to the fact that the size of  $\xi_t^{(2)}$  depends on the norm of the GD iterate  $\|v_t\|$ , which can be, in principle, unbounded. Consequently, the resulting error terms can only be controlled in a probabilistic sense. Our analysis relies on showing that there indeed exists a condition  $\mathcal{E}_1 \cap \mathcal{E}_2$  that holds with high probability and ensures the desired properties. A formal treatment of these matters is presented in Appendix E. The final result of these derivations is Proposition 6 that, under appropriate conditions on the algorithm’s parameters, bounds the deviations between the averaged GD and SGD iterates as

$$\mathbb{E} \left[ \left\| \Sigma^{\frac{1}{2}} (\bar{w}_{S,T} - \bar{v}_{S,T}) \right\|^2 \right] \lesssim \frac{\gamma^\alpha \text{Tr}[\Sigma^\alpha]}{bL^{1-\alpha}}.$$

## B Spectral Filtering properties of averaged GD

Consider the function

$$g_t(\sigma) = \gamma \sum_{k=0}^{t-1} (1 - \gamma\sigma)^k = \sigma^{-1} (1 - (1 - \gamma\sigma)^t). \quad (23)$$

defined on the spectrum  $\sigma(\Sigma) \subseteq [0, \kappa^2]$  of  $\Sigma$  and let

$$r_t(\sigma) = 1 - \sigma g_t(\sigma).$$

Then, for any  $\alpha \in [0, 1]$  [12]

$$\sup_{0 < \sigma \leq \kappa^2} |\sigma^\alpha g_t(\sigma)| \leq (\gamma t)^{1-\alpha}. \quad (24)$$

Moreover, for any  $0 \leq u$

$$\sup_{0 < \sigma \leq \kappa^2} |r_t(\sigma)| \sigma^u \leq C_u (\gamma t)^{-u}, \quad (25)$$

for some  $C_u > 0$ . In particular,  $C_0 = 1$ .

For  $0 \leq S \leq T - 1$  consider

$$G_{S,T}(\sigma) = \frac{1}{T-S} \sum_{t=S+1}^T g_t(\sigma)$$

and let

$$R_{T,S}(\sigma) = 1 - \sigma G_{S,T}(\sigma).$$

**Lemma 1** (Filter).

$$\sigma G_{S,T}(\sigma) = 1 - \frac{1}{(T-S)\gamma\sigma} (1 - \gamma\sigma)^{S+1} (1 - (1 - \gamma\sigma)^{T-S}).$$

*Proof of Lemma 1.* By (23), we have

$$\begin{aligned} \sigma G_{S,T}(\sigma) &= \frac{1}{T-S} \sum_{t=S+1}^T 1 - (1 - \gamma\sigma)^t \\ &= 1 - \frac{1}{T-S} \sum_{t=S+1}^T (1 - \gamma\sigma)^t \\ &= 1 - \frac{1}{T-S} \sum_{t=0}^T (1 - \gamma\sigma)^t + \frac{1}{T-S} \sum_{t=0}^S (1 - \gamma\sigma)^t \\ &= 1 - \frac{((1 - \gamma\sigma)^{S+1} - (1 - \gamma\sigma)^{T+1})}{(T-S)\gamma\sigma} \\ &= 1 - \frac{(1 - \gamma\sigma)^{S+1} (1 - (1 - \gamma\sigma)^{T-S})}{(T-S)\gamma\sigma}. \end{aligned}$$

□

**Lemma 2** (Properties I). *Let  $u \in [0, 1]$ . For any  $1 \leq T$ ,  $0 \leq S \leq T - 1$  we have*

$$\sup_{0 < \sigma \leq \kappa^2} |\sigma^u G_{S,T}(\sigma)| \leq C_u \gamma^{1-u} (T+S)(T-S)^{-u}.$$

*In particular, for  $1 \leq K$ , choosing  $S \leq \frac{K-1}{K+1} T$  gives*

$$\sup_{0 < \sigma \leq \kappa^2} |\sigma^u G_{S,T}(\sigma)| \leq K \gamma^{1-u} (T-S)^{1-u}.$$



*Proof.* By (16) we have

$$\begin{aligned} \sup_{0 < \sigma \leq \kappa^2} |\sigma^u G_{S,T}(\sigma)| &\leq \sup_{0 < \sigma \leq \kappa^2} \frac{1}{T-S} \sum_{t=S+1}^T |\sigma^u g_t(\sigma)| \\ &\leq \frac{\gamma^{1-u}}{T-S} \sum_{t=S+1}^T t^{1-u}. \end{aligned}$$

From Lemma 14 and Lemma 12, we find

$$\begin{aligned} \sum_{t=S+1}^T t^{1-u} &\leq \int_{S+1}^T t^{1-u} dt \\ &= \frac{1}{2-u} (T^{2-u} - (S+1)^{2-u}) \\ &\leq \frac{1}{2-u} (T^{2-u} - S^{2-u}) \\ &\leq \frac{1}{2-u} (T+S)(T-S)^{1-u}. \end{aligned}$$

This proves the first statement. The second statement follows by observing that  $T+S \leq K(T-S)$  if  $S \leq \frac{K-1}{K+1} T$ .  $\square$

**Remark 1.** A more refined bound for the case  $u = 1$  can be obtained by considering (16), which directly leads to

$$\sup_{0 < \sigma \leq \kappa^2} |\sigma G_{S,T}(\sigma)| \leq \sup_{0 < \sigma \leq \kappa^2} \frac{1}{T-S} \sum_{t=S+1}^T |\sigma g_t(\sigma)| \leq 1.$$

**Lemma 3** (Properties II). *Let  $1 \leq T$  and  $0 \leq S \leq T-1$ .*

1. *For any  $u \in [0, 1]$ , we have*

$$\sup_{0 < \sigma \leq \kappa^2} |\sigma^u R_{S,T}(\sigma)| \leq \gamma^{-u} (T-S)^{-u}.$$

2. *For any  $u > 1$  we have*

$$\sup_{0 < \sigma \leq \kappa^2} |\sigma^u R_{S,T}(\sigma)| \leq \tilde{C}_u \gamma^{-u} \frac{S+1}{T-S} \left( \frac{1}{S+1} \right)^u,$$

*for some  $\tilde{C}_u < \infty$ . In particular, if  $1 \leq K$  and  $S \leq \frac{K-1}{K+1} T$  one has*

$$\sup_{0 < \sigma \leq \kappa^2} |\sigma^u R_{S,T}(\sigma)| \leq \tilde{C}_u \gamma^{-u} K \left( \frac{1}{S+1} \right)^u.$$

*If additionally  $T \leq (K+1)S$ , one has*

$$\sup_{0 < \sigma \leq \kappa^2} |\sigma^u R_{S,T}(\sigma)| \leq 2\tilde{C}_u \gamma^{-u} K^2 \left( \frac{1}{T-S} \right)^u.$$

*Proof of Lemma 3.* By Lemma 1 we have for any  $u \in [0, 1]$

$$\begin{aligned}
|\sigma^u R_{S,T}(\sigma)| &= \frac{\sigma^u}{(T-S)\gamma\sigma} (1-\gamma\sigma)^{S+1} (1-(1-\gamma\sigma)^{T-S}) \\
&\leq \frac{\sigma^u}{(T-S)\gamma\sigma} (1-\gamma\sigma)^{S+1} ((T-S)\gamma\sigma)^{1-u} \\
&= (\gamma(T-S))^{-u} (1-\gamma\sigma)^{S+1} \\
&\leq \gamma^{-u} (T-S)^{-u},
\end{aligned}$$

where we use that

$$|1 - (1-x)^t| \leq (tx)^{1-u}$$

for any  $x \in [0, 1]$  and for any  $u \in [0, 1]$ .

For  $u > 1$  we apply (25) and Lemma 14 and obtain

$$\begin{aligned}
\sup_{0 < \sigma \leq \kappa^2} |\sigma^u R_{S,T}(\sigma)| &\leq C_u \frac{\gamma^{-u}}{T-S} \sum_{t=S+1}^T t^{-u} \\
&\leq C_u \frac{\gamma^{-u}}{T-S} \left( (S+1)^{-u} + \int_{S+1}^T t^{-u} dt \right) \\
&\leq C_u \frac{\gamma^{-u}}{T-S} \left( (S+1)^{-u} + \frac{1}{u-1} \left( \left( \frac{1}{S+1} \right)^{u-1} - \left( \frac{1}{T+1} \right)^{u-1} \right) \right) \\
&\leq \tilde{C}_u \gamma^{-u} \frac{S+1}{T-S} \left( \frac{1}{S+1} \right)^u. \tag{26}
\end{aligned}$$

Note that  $S \leq \frac{K-1}{K+1} T$  implies

$$\frac{S+1}{T-S} \leq K.$$

Finally,  $T \leq (K+1)S$  gives

$$\frac{1}{S+1} \leq \frac{K+1}{T-S} \leq \frac{2K}{T-S}.$$

□

## C Bounds Tail-Averaged Gradient Descent

Our error bounds are derived by means of a classical error decomposition in bias and variance, see e.g. [6], [3], [5] and [22]. More precisely, recalling the filter expression of the population GD,

$$u_L = G_L(\Sigma)\Sigma w_*, \tag{27}$$

we consider

$$\begin{aligned}
\bar{v}_L - w_* &= (\bar{v}_L - u_L) + (u_L - w_*) \\
&= (\bar{v}_L - u_L) + R_L(\Sigma)w_* \\
&= (G_L(\hat{\Sigma})\hat{h} - G_L(\hat{\Sigma})\hat{\Sigma}u_L) + (G_L(\hat{\Sigma})\hat{\Sigma}u_L - u_L) + R_L(\Sigma)w_* \\
&= G_L(\hat{\Sigma})(\hat{h} - \hat{\Sigma}u_L) + R_L(\hat{\Sigma})u_L + R_L(\Sigma)w_*. \tag{28}
\end{aligned}$$

We refer to

$$\mathcal{A}(L) = \|\Sigma^{1/2}R_L(\Sigma)w_*\|^2 \tag{29}$$

as the *deterministic Approximation error*, to

$$\hat{\mathcal{A}}(L) = \|\Sigma^{1/2} R_L(\hat{\Sigma}) u_L\|^2 \quad (30)$$

as the *stochastic Approximation error* and to

$$\hat{\mathcal{V}}(L) = \|\Sigma^{1/2} G_L(\hat{\Sigma})(\hat{h} - \hat{\Sigma} u_L)\|^2 \quad (31)$$

as the *Sample variance*. In what follows we successively bound each error term in Section C.1, Section C.2 and Section C.3. Finally, the total bound is given in Section C.4.

In the following we let

$$\Sigma_L = \left(\Sigma + \frac{1}{\gamma L}\right), \quad \text{and} \quad \hat{\Sigma}_L = \left(\hat{\Sigma} + \frac{1}{\gamma L}\right).$$

## C.1 Bounding the deterministic Approximation Error

**Proposition 1** (Deterministic Approximation Error). *Let  $1 \leq T$ ,  $0 \leq S \leq T - 1$ ,  $\gamma\kappa^2 < 1$  and Assumption 2 hold.*

1. *If  $0 \leq r \leq 1/2$ , we have*

$$\mathcal{A}(L) \leq R^2 (\gamma L)^{-2(r+1/2)}.$$

2. *If  $1/2 < r$  we have*

$$\mathcal{A}(L) \leq C_r R^2 \gamma^{-2(r+1/2)} \left(\frac{S+1}{L}\right)^2 \left(\frac{1}{S+1}\right)^{2(r+1/2)}.$$

*for some  $C_r < \infty$ . In particular, if  $1 \leq K$ ,  $S \leq \frac{K-1}{K+1} T$  and  $T \leq (K+1)S$ , one has*

$$\mathcal{A}(L) \leq C_r K^2 R^2 (\gamma L)^{-2(r+1/2)}.$$

*Proof of Proposition 1.* By Assumption 2 we have

$$\mathcal{A}(L) = \|\Sigma^{1/2} R_L(\Sigma) w_*\|^2 \leq R^2 \|\Sigma^{r+1/2} R_L(\Sigma)\|^2.$$

Since

$$\|\Sigma^{r+1/2} R_L(\Sigma)\| \leq \sup_{0 < \sigma \leq \kappa^2} |\sigma^{r+1/2} R_L(\sigma)|,$$

the result follows immediately by applying Lemma 3.  $\square$

## C.2 Bounding the stochastic Approximation Error

**Proposition 2** (Stochastic Approximation Error). *Let  $1 \leq T$ ,  $0 \leq S \leq T - 1$ ,  $\gamma\kappa^2 < 1$  and Assumption 2 hold. Further assume*

$$n \geq 16\kappa^2 \gamma L \max\{1, \mathcal{N}(1/(\gamma L))\}.$$

1. *If  $0 \leq r \leq 1/2$ , we have*

$$\mathbb{E}[\hat{\mathcal{A}}(L)] \leq C_r R^2 \left(\frac{T+S}{L}\right)^2 (\gamma L)^{-2(r+1/2)},$$

*for some  $C_r < \infty$ .*

2. If  $1/2 < r$  we have

$$\mathbb{E}[\hat{A}(L)] \leq C_{r,\kappa} R^2 \left( \frac{T+S}{L} \right)^2 \left( (\gamma L) \frac{(S+1)^2}{L^2} \Psi_r^2(S, T) + \frac{1}{n} \right),$$

for some  $C_{r,\kappa} < \infty$  and where  $\Psi_r$  is defined in (37). In particular, if  $1 \leq K$ ,  $S \leq \frac{K-1}{K+1} T$  and  $T \leq (K+1)S$ , one has

$$\mathbb{E}[\hat{A}(L)] \leq \tilde{C}_{r,\kappa} K^{4(r+1)} (\gamma L)^{-2(r+1/2)},$$

for some  $\tilde{C}_{r,\kappa} < \infty$ .

*Proof of Proposition 2.* We start with deriving bounds holding with high probability, bounds in expectation follow then by integration. From Lemma 10 we derive with probability at least  $1 - \delta/2$

$$\begin{aligned} \hat{A}(L) &= \|\Sigma^{1/2} R_L(\hat{\Sigma}) u_L\|^2 \\ &\leq 16 \log^2(4\delta^{-1}) \|\hat{\Sigma}_L^{1/2} R_L(\hat{\Sigma}) u_L\|^2. \end{aligned} \quad (32)$$

We separate the analysis by considering two cases.

**Case 1** ( $0 \leq r \leq 1/2$ ): Recalling the definition of  $u_L$  in (27) gives

$$\|\hat{\Sigma}_L^{1/2} R_L(\hat{\Sigma}) u_L\| \leq R \|\hat{\Sigma}_L R_L(\hat{\Sigma})\| \cdot \|\hat{\Sigma}_L^{-1/2} \Sigma_L^{1/2}\| \cdot \|\Sigma^{r+1/2} G_L(\Sigma)\|.$$

Bounding the first term is done by using Lemma 1 and Lemma 3, leading to

$$\begin{aligned} \|\hat{\Sigma}_L R_L(\hat{\Sigma})\| &\leq \sup_{0 < \sigma \leq \kappa^2} |(\sigma + 1/(\gamma L)) R_L(\sigma)| \\ &\leq \sup_{0 < \sigma \leq \kappa^2} |\sigma R_L(\sigma)| + 1/(\gamma L) \sup_{0 < \sigma \leq \kappa^2} |R_L(\sigma)| \\ &\leq 2(\gamma L)^{-1}. \end{aligned}$$

From Lemma 2 we obtain

$$\|\Sigma^{r+1/2} G_L(\Sigma)\| \leq C_r \gamma^{1/2-r} (T+S) L^{-(r+1/2)}.$$

Thus, applying Corollary 3, gives with probability at least  $1 - \delta/2$

$$\hat{A}(L) \leq 64 \cdot 16 C_r^2 R^2 \log^4(4\delta^{-1}) (\gamma L)^{-2} \gamma^{1-2r} (T+S)^2 L^{-2(r+1/2)}, \quad (33)$$

for some  $C_r < \infty$ .

**Case 2** ( $1/2 < r$ ): In this case we split (32) differently. Using Assumption 2 and Definition (27), we obtain

$$\begin{aligned} \|\hat{\Sigma}_L^{1/2} R_L(\hat{\Sigma}) u_L\| &= \|\hat{\Sigma}_L^{1/2} R_L(\hat{\Sigma}) G_L(\Sigma) \Sigma w_*\| \\ &\leq R \underbrace{\|\hat{\Sigma}_L^{1/2} R_L(\hat{\Sigma}) (\Sigma^{r+1/2} - \hat{\Sigma}^{r+1/2}) G_L(\Sigma) \Sigma^{1/2}\|}_{A_1} \\ &\quad + R \underbrace{\|\hat{\Sigma}_L^{1/2} R_L(\hat{\Sigma}) \hat{\Sigma}^{r+1/2} G_L(\Sigma) \Sigma^{1/2}\|}_{A_2}. \end{aligned} \quad (34)$$

**Bounding  $A_1$ :** For bounding  $A_1$  we apply [5], Proposition 5.5. and Proposition 5.6., to obtain with probability at least  $1 - \delta/2$

$$\begin{aligned} \|\Sigma^{r+1/2} - \hat{\Sigma}^{r+1/2}\| &\leq C_r \|\Sigma - \hat{\Sigma}\| \\ &\leq 6C_r \frac{\kappa^2}{\sqrt{n}} \log(4\delta^{-1}) . \end{aligned}$$

Furthermore, Lemma 2 gives

$$\|G_L(\Sigma)\Sigma^{1/2}\| \leq C \gamma^{1/2} (T + S)L^{-1/2} , \quad (35)$$

for some numerical constant  $C < \infty$ . Moreover, using Lemma 3 leads to

$$\begin{aligned} \|\hat{\Sigma}_L^{1/2} R_L(\hat{\Sigma})\| &\leq \sup_{0 < \sigma \leq \kappa^2} |(\sigma + (1/(\gamma L)))^{1/2} R_L(\sigma)| \\ &\leq \sup_{0 < \sigma \leq \kappa^2} |\sigma^{1/2} R_L(\sigma)| + (1/(\gamma L))^{1/2} \sup_{0 < \sigma \leq \kappa^2} |R_L(\sigma)| \\ &\leq 2(\gamma L)^{-1/2} . \end{aligned}$$

Collecting the previous steps we arrive at

$$\|A_1\| \leq C'_r \left( \frac{T + S}{L} \right) \frac{\kappa^2}{\sqrt{n}} \log(4\delta^{-1}) , \quad (36)$$

with probability at least  $1 - \delta/2$ , for some numerical constant  $C'_r < \infty$ .

**Bounding  $A_2$ :** For bounding  $A_2$  we apply Lemma 3 once more, giving

$$\begin{aligned} &\|\hat{\Sigma}_L^{1/2} R_L(\hat{\Sigma}) \hat{\Sigma}^{r+1/2}\| \\ &\leq \sup_{0 < \sigma \leq \kappa^2} |(\sigma + (1/(\gamma L)))^{1/2} R_L(\sigma) \sigma^{r+1/2}| \\ &\leq \sup_{0 < \sigma \leq \kappa^2} |\sigma^{r+1} R_L(\sigma)| + (\gamma L)^{-1/2} \sup_{0 < \sigma \leq \kappa^2} |\sigma^{r+1/2} R_L(\sigma)| \\ &\leq C'_r \gamma^{-(r+1)} \frac{S+1}{L} \left( \frac{1}{S+1} \right)^{r+1} + C''_r \gamma^{-(r+1/2)} \frac{S+1}{L} \left( \frac{1}{S+1} \right)^{r+1/2} (\gamma L)^{-1/2} \\ &\leq C'''_r \frac{S+1}{L} \Psi_r(S, T) . \end{aligned}$$

where we set

$$\Psi_r(S, T) := \gamma^{-(r+1)} \left[ \left( \frac{1}{S+1} \right)^{r+1} + L^{-1/2} \left( \frac{1}{S+1} \right)^{r+1/2} \right] . \quad (37)$$

Thus, combining with (35), we find

$$\|A_2\| \leq \tilde{C}_r (\gamma L)^{1/2} \frac{T + S}{L} \frac{S + 1}{L} \Psi_r(S, T) . \quad (38)$$

Finally, note that  $S \leq \frac{K-1}{K+1} T$  implies

$$\frac{T + S}{L} \leq K , \quad \frac{S + 1}{L} \leq K$$

and  $T \leq (K + 1)S$  gives

$$\frac{1}{S + 1} \leq \frac{K + 1}{L} \leq \frac{2K}{L} .$$

Hence,

$$\Psi_r(S, T) \leq (4K)^{r+1} (\gamma L)^{-(r+1)}. \quad (39)$$

Thus,

$$\|A_2\| \leq \tilde{C}_r (4K)^{2(r+1)} (\gamma L)^{-(r+1/2)}. \quad (40)$$

The result in this case then follows by combining (39), (34) with (36), (38) and (40) and by integration, Lemma 11.  $\square$

### C.3 Bounding the Sample Variance

For proving the bound for the sample variance we need a concentration result which we slightly generalize from [22].

**Proposition 3.** *Let  $u_L$  be defined by (27),  $\mathcal{A}(L)$  by (18) and  $\delta \in (0, 1]$ . Under Assumption 1, one has with probability at least  $1 - \delta$*

$$\begin{aligned} & \left\| \left( \Sigma + \frac{1}{\gamma L} \right)^{-1/2} \left( (\hat{\Sigma} u_L - \hat{h}) - (\Sigma u_L - h) \right) \right\| \\ & \leq c \log(2\delta^{-1}) \left( \frac{\sqrt{\gamma L} (\kappa M + \kappa^2 \|u_L\|)}{n} + \sqrt{\frac{\kappa^2 \gamma L \mathcal{A}(L) + \mathcal{N}(\frac{1}{\gamma L})}{n}} \right), \end{aligned}$$

for some numerical constant  $c < \infty$ .

**Proposition 4** (Sample Variance). *Set  $L = T - S$  and assume  $\gamma \kappa^2 < 1$  as well as*

$$n \geq 16\kappa^2 \gamma L \max\{1, \mathcal{N}(1/(\gamma L))\}.$$

Under Assumption 1 one has

$$\begin{aligned} \mathbb{E} \left[ \hat{V}(L) \right] & \leq C_{\tilde{\sigma}, M, \kappa} \left( 1 + \frac{\Delta(S, T)}{L^2} \right)^2 \\ & \quad \left( \mathcal{A}(L) + \frac{\gamma L (1 + \|w_*\|^2)}{n^2} + \frac{\mathcal{N}(1/\gamma L)}{n} \right) \end{aligned}$$

for some  $C_{\tilde{\sigma}, M, \kappa} < \infty$  and where

$$\Delta(S, T) = T(T+1) - S(S+1).$$

In particular, if  $1 \leq K$  and  $S \leq \frac{K-1}{K+1} T$  one has

$$1 + \frac{\Delta(S, T)}{L^2} \leq 1 + 2K.$$

*Proof of Proposition 4.* According to Lemma 10, we have with probability at least  $1 - \delta/2$

$$\begin{aligned} \hat{V}(L) & = \|\Sigma^{1/2} G_L(\hat{\Sigma})(\hat{h} - \hat{\Sigma} u_L)\|^2 \\ & \leq 16 \log^2(4\delta^{-1}) \|\hat{\Sigma}_L^{1/2} G_L(\hat{\Sigma})(\hat{h} - \hat{\Sigma} u_L)\|^2. \end{aligned} \quad (41)$$

We proceed by decomposing as follows:

$$\hat{\Sigma}_L^{1/2} G_L(\hat{\Sigma})(\hat{h} - \hat{\Sigma} u_L) = \hat{\Sigma}_L G_L(\hat{\Sigma}) \cdot \hat{\Sigma}_L^{-1/2} \Sigma_L^{1/2} \cdot \hat{h}_L,$$

with

$$\hat{h}_L = \Sigma_L^{1/2} \left( \hat{h} - \hat{\Sigma} u_L \right).$$

Using the filter function properties in Lemma 1 and Lemma 2 gives

$$\begin{aligned} \|\hat{\Sigma}_L G_L(\hat{\Sigma})\| &\leq \sup_{0 < \sigma \leq \kappa^2} |(\sigma + 1/(\gamma L)) G_L(\sigma)| \\ &\leq 1 + \frac{1}{L^2} (T(T+1) - S(S+1)) \\ &= 1 + \frac{\Delta(S, T)}{L^2}, \end{aligned} \quad (42)$$

with  $\Delta(S, T) = T(T+1) - S(S+1)$ . Furthermore, Corollary 3 gives

$$\|\hat{\Sigma}_L^{-1/2} \Sigma_L^{1/2}\| \leq 4 \log(8\delta^{-1}) \quad (43)$$

with probability at least  $1 - \delta/4$ .

For bounding  $\hat{h}_L$  we need to decompose once more: Since  $\Sigma w_* = h$ , we find

$$\begin{aligned} \hat{h}_L &= \Sigma_L^{1/2} \left( (\hat{h} - \hat{\Sigma} u_L) - (h - \Sigma u_L) \right) + \Sigma_L^{1/2} (h - \Sigma u_L) \\ &= \Sigma_L^{1/2} \left( (\hat{h} - \hat{\Sigma} u_L) - (h - \Sigma u_L) \right) + \Sigma_L^{1/2} \Sigma R_L(\Sigma) w_*, \end{aligned}$$

satisfying

$$\|\hat{h}_L\| \leq \|\Sigma_L^{1/2} \left( (\hat{h} - \hat{\Sigma} u_L) - (h - \Sigma u_L) \right)\| + \sqrt{\mathcal{A}(L)}.$$

Applying Proposition 3 gives

$$\|\hat{h}_L\| \leq \sqrt{\mathcal{A}(L)} + 12 \log(8\delta^{-1}) \left( \frac{\sqrt{\gamma L} (\kappa M + \kappa^2 \|u_L\|)}{n} + \sqrt{\frac{\kappa^2 \gamma L \mathcal{A}(L) + \tilde{\sigma}^2 \mathcal{N}(\frac{1}{\gamma L})}{n}} \right), \quad (44)$$

with probability at least  $1 - \delta/4$ . Collecting (42), (43) and (44) yields

$$\begin{aligned} \hat{\mathcal{V}}(L) &\leq C_{\tilde{\sigma}, M, \kappa} \log^2(8\delta^{-1}) \left( 1 + \frac{\Delta(S, T)}{L^2} \right)^2 \\ &\quad \left( \mathcal{A}(L) + \frac{\gamma L (1 + \|u_L\|^2)}{n^2} + \frac{\mathcal{N}(1/\gamma L)}{n} \right), \end{aligned} \quad (45)$$

with probability at least  $1 - \delta/4$ , for some  $C_{\tilde{\sigma}, M, \kappa} < \infty$ . Finally, Lemma 1 ensures that

$$\|u_L\| = \|\Sigma G_L(\Sigma) w_*\| \leq \|w_*\|.$$

The bound in expectation follows from Lemma 11 by integration.

For the last part we refer to the proof of Lemma 2, from which we deduce that

$$1 + \frac{\Delta(S, T)}{L^2} \leq 1 + 2K,$$

provided that  $1 \leq K$  and  $S \leq \frac{K-1}{K+1} T$ . □

## C.4 Main result on GD convergence

Proposition 1, Proposition 2 and Proposition 4 together lead our main result regarding the convergence of batch gradient descent, stated as the following theorem.

**Theorem 2.** *Let  $1 \leq T$ ,  $0 \leq S \leq T - 1$ , Assumptions 1, 2 hold. Set  $L = T - S$  and assume  $\gamma\kappa^2 < 1$  as well as*

$$n \geq 16\kappa^2 \gamma L \max\{1, \mathcal{N}(1/(\gamma L))\}. \quad (46)$$

1. *If  $0 \leq r \leq 1/2$  and  $1 \leq K$ ,  $0 \leq S \leq \frac{K-1}{K+1} T$ , we have*

$$\begin{aligned} \mathbb{E} \left[ \|\Sigma^{1/2}(\bar{v}_L - w_*)\|^2 \right] &\leq C_r R^2 K^2 (\gamma L)^{-2(r+1/2)} \\ &\quad + C_{\kappa, M, \sigma, \nu} K^2 \left( \mathcal{A}(L) + \frac{\gamma L(1 + \|w_*\|)^2}{n^2} + \frac{\gamma L \mathcal{A}(L)}{n} + \frac{\mathcal{N}(\frac{1}{\gamma L})}{n} \right), \end{aligned}$$

*for some  $C_r < \infty$  and  $C_{\kappa, M, \sigma, \nu} < \infty$ .*

2. *If  $1/2 < r$ ,  $1 < K$ ,  $0 < S \leq \frac{K-1}{K+1} T$  and  $T \leq (K+1)S$ , we have*

$$\begin{aligned} \mathbb{E} \left[ \|\Sigma^{1/2}(\bar{v}_L - w_*)\|^2 \right] &\leq C_r C_K R^2 \left[ (\gamma L)^{-2(r+1/2)} + \frac{1}{n} \right] \\ &\quad + C_{\kappa, M, \sigma, \nu} C'_K \left( \mathcal{A}(L) + \frac{\gamma L(1 + \|w_*\|)^2}{n^2} + \frac{\gamma L \mathcal{A}(L)}{n} + \frac{\mathcal{N}(\frac{1}{\gamma L})}{n} \right), \end{aligned}$$

*for some  $C_{\kappa, r} < \infty$  and  $C_{\kappa, M, \sigma, \tilde{r}} < \infty$ .*

From Theorem 2 we can immediately derive the Proof of Corollary 2.

**Corollary 2** (Rates of Convergence). *Let any assumption of Theorem 2 hold and assume additionally Assumption 3. One has for any  $n$  sufficiently large*

$$\mathbb{E} \left[ \|\Sigma^{1/2}(\bar{v}_L - w_*)\|^2 \right] \leq C n^{-\frac{2(r+1/2)}{2r+\nu}},$$

*under each of the following choices:*

1. *If  $0 \leq r \leq 1/2$ :  $S = 0$ ,  $\alpha, \beta \geq 0$  and*

$$\gamma_n \simeq n^{-\alpha} \quad T_n \simeq n^\beta \quad \text{such that} \quad \alpha - \beta = \frac{1}{2r+1+\nu}. \quad (47)$$

2. *If  $1/2 < r$ :  $0 < S$ ,  $S_n \asymp T_n$ , with  $T_n, \gamma_n$  as in (47).*

*Proof of Corollary 2.* Let  $\gamma_n \simeq n^{-a}$ ,  $L_n \simeq n^{\tilde{a}}$ , with  $a, \tilde{a} > 0$  satisfying  $a - \tilde{a} = \frac{1}{2r+1+\nu}$ . Plugging in Assumptions 2 and 3 gives in either case

$$\mathbb{E} \left[ \|\Sigma^{1/2}(\bar{v}_{L_n} - w_*)\|^2 \right] \leq C \left( (\gamma_n L_n)^{-2(r+1/2)} + \frac{\gamma_n L_n}{n^2} + \frac{(\gamma_n L_n)^{-2r}}{n} + \frac{(\gamma_n L_n)^\nu}{n} + \frac{1}{n} \right),$$

for some constant  $C < \infty$ , depending on all model parameters  $\kappa, M, \nu, r, R$  and  $\|w_*\|$ . A short calculation shows that

$$n^{-1} = o\left((\gamma_n L_n)^{-2(r+1/2)}\right), \quad \frac{\gamma_n L_n}{n^2} = o\left((\gamma_n L_n)^{-2(r+1/2)}\right)$$

and

$$\frac{(\gamma_n L_n)^{-2r}}{n} = o\left((\gamma_n L_n)^{-2(r+1/2)}\right),$$

so we can disregard the terms  $n^{-1}$ ,  $\frac{\gamma_n L_n}{n^2}$ ,  $\frac{(\gamma_n L_n)^{-2r}}{n}$  for  $n$  large enough. The choice

$$\gamma_n L_n \simeq n^{\frac{1}{2r+1+\nu}}$$

precisely balances the two remaining terms  $(\gamma_n L_n)^{-2(r+1/2)}$  and  $\frac{(\gamma_n L_n)^\nu}{n}$ . This choice also implies Assumption (46) if  $n$  is sufficiently large.  $\square$



## D A general Result

Consider the recursion

$$\mu_{t+1} = \hat{Q}_{t+1} \mu_t + \gamma \xi_{t+1}, \quad \hat{Q}_t = (I - \gamma \hat{H}_t), \quad (48)$$

with  $\mu_0 = 0$ , with  $\hat{H}_t$  linear i.i.d. random operators acting on  $\mathcal{H}$  and with  $\xi_t \in \mathcal{H}$  i.i.d. random variables, satisfying  $\mathbb{E}[\xi_t] = 0$ . For  $0 \leq S \leq T - 1$  we let

$$\bar{\mu} := \bar{\mu}_{S,T} := \frac{1}{T-S} \sum_{t=S+1}^T \mu_t. \quad (49)$$

Denote  $H = \mathbb{E}[\hat{H}_t]$ . We assume that  $\text{Tr}[H^\alpha] < \infty$  for some  $\alpha \in (0, 1]$  and

$$\mathbb{E}[\xi_t \otimes \xi_t] \preceq \sigma^2 H, \quad \mathbb{E}[\hat{H}_t^2] \preceq \kappa^2 H. \quad (50)$$

The last condition holds in particular when the  $\hat{H}_t$  are bounded a.s. by  $\kappa^2$ . We generalize Proposition 1 given in [26] (see also [10]) to more general recursions and to tail-averaging, including full averaging and mini-batching as special cases.

**Proposition 5.** *Let  $\alpha \in (0, 1]$ ,  $\gamma \kappa^2 \leq 1/4$  and  $u \in [0, 1 + \alpha]$ . Under Assumption (50), one has*

$$\mathbb{E} \left[ \|H^{\frac{u}{2}} \bar{\mu}_{S,T}\|^2 \right] \leq 16\sigma^2 \text{Tr}[H^\alpha] \gamma^{1-u+\alpha} (T-S)^{\alpha-u} \Upsilon(S, T),$$

with  $\Upsilon(S, T) = 1 + \frac{S+1}{T-S}$ . If additionally  $1 \leq K$  and  $1 \leq T$ ,  $0 \leq S \leq T - 1$  satisfy  $S \leq \frac{K-1}{K+1} T$ , we have

$$\Upsilon(S, T) \leq 1 + K.$$

The proof of this result is carried out in Section D.2. The basic idea is to derive a similar bound for the related *semi-stochastic recursion* (51), where  $\hat{H}_t$  is replaced by its expectation  $H$ , leaving the randomness in the noise variables  $\xi_t$ . This is done in Section D.1. In a second step one needs to control the difference between the full-stochastic recursion and the semi-stochastic iterates. This relies on a *perturbation argument*, summarized in Section D.2.

### D.1 Semi-Stochastic Recursion (SSR)

Let  $H$  be a positive, self-adjoint operator on some Hilbert space  $\mathcal{H}$ , satisfying  $H \preceq \kappa^2 I$ . Consider the general recursion in  $\mathcal{H}$

$$\mu_{t+1} = (1 - \gamma H) \mu_t + \gamma \xi_{t+1}, \quad (51)$$

with  $\mu_0 = 0$  and  $\gamma \kappa^2 < 1$ . We further assume that

$$\mathbb{E}[\xi_t] = 0, \quad \mathbb{E}[\xi_t \otimes \xi_t] \preceq \sigma^2 H.$$

For  $1 \leq T$  and  $0 \leq S \leq T - 1$ , we consider

$$\bar{\mu} := \bar{\mu}_{S,T} := \frac{1}{T-S} \sum_{t=S+1}^T \mu_t.$$

**Lemma 4** (SSR). *Let  $\alpha \in (0, 1]$  and assume that  $\text{Tr}(H^\alpha) < \infty$ . Let  $1 \leq T$ . For any  $u \in [0, 1 + \alpha]$  we have*

$$\mathbb{E} \left[ \|H^{u/2} \bar{\mu}\|^2 \right] \leq 4\sigma^2 \text{Tr}[H^\alpha] \gamma^{1-u+\alpha} (T-S)^{\alpha-u} \Upsilon(S, T),$$

with  $\Upsilon(S, T) = 1 + \frac{S+1}{T-S}$ . In particular, given  $1 \leq K$  and if  $0 \leq S \leq \frac{K-1}{K+1} T$  one has

$$\Upsilon(S, T) \leq 1 + K.$$

*Proof of Lemma 4.* Setting  $Q = 1 - \gamma H$ , a standard calculation combined with the fact

$$\sum_{t=S+1}^T q^t = \frac{q^{S+1}(1 - q^{T-S})}{1 - q} \quad (52)$$

shows that the averaged iterates are given by

$$\begin{aligned} \bar{\mu} &= \frac{\gamma}{T-S} \sum_{t=S+1}^T \sum_{k=0}^{t-1} Q^{t+1-k} \xi_k \\ &= \frac{\gamma}{T-S} \sum_{t=0}^S \left( \sum_{k=S-t}^{T-(t+1)} Q^k \right) \xi_t + \frac{\gamma}{T-S} \sum_{t=S+1}^{T-1} \left( \sum_{k=0}^{T-(t+1)} Q^k \right) \xi_t \\ &= \sum_{t=0}^S A_t \xi_t + \sum_{t=S+1}^{T-1} \tilde{A}_t \xi_t, \end{aligned}$$

where we set

$$A_t := \frac{\gamma}{T-S} \left( \sum_{k=S-t}^{T-(t+1)} Q^k \right), \quad \tilde{A}_t := \frac{\gamma}{T-S} \left( \sum_{k=0}^{T-(t+1)} Q^k \right).$$

Thus, since  $\mathbb{E}[\xi_t \otimes \xi_t] \preceq \sigma^2 H$ , we find

$$\begin{aligned} \mathbb{E}[\|H^{u/2} \bar{\mu}\|^2] &\leq 2 \sum_{t=0}^S \mathbb{E}[\text{Tr}[H^u A_t^2 \xi_t \otimes \xi_t]] + 2 \sum_{t=S+1}^{T-1} \mathbb{E}[\text{Tr}[H^u \tilde{A}_t^2 \xi_t \otimes \xi_t]] \\ &= 2 \sum_{t=0}^S \text{Tr}[H^u A_t^2 \mathbb{E}[\xi_t \otimes \xi_t]] + 2 \sum_{t=S+1}^{T-1} \text{Tr}[H^u \tilde{A}_t^2 \mathbb{E}[\xi_t \otimes \xi_t]] \\ &\leq \underbrace{2\sigma^2 \sum_{t=0}^S \text{Tr}[H^{u+1} A_t^2]}_{\mathcal{J}_1} + \underbrace{2\sigma^2 \sum_{t=S+1}^{T-1} \text{Tr}[H^{u+1} \tilde{A}_t^2]}_{\mathcal{J}_2}. \end{aligned} \quad (53)$$

We proceed bounding the individual terms by applying (52). This gives

$$A_t = \frac{H^{-1}}{T-S} Q^{S-t} (1 - Q^{T-S}) \preceq \frac{H^{-1}}{T-S} (1 - Q^{T-S}).$$

Furthermore,

$$\begin{aligned} \text{Tr}[H^{u-1} (1 - Q^{T-S})^2] &= \sum_{j \in \mathbb{N}} \sigma_j^{u-1} (1 - (1 - \gamma \sigma_j)^{T-S})^2 \\ &\leq \sum_{j \in \mathbb{N}} \sigma_j^{u-1} ((T-S) \gamma \sigma_j)^{1-u+\alpha} \\ &= \gamma^{1-u+\alpha} \text{Tr}[H^\alpha] (T-S)^{1-u+\alpha}, \end{aligned}$$

where in the inequality we use that for any  $x \in [0, 1]$ ,  $u \in [0, 1 + \alpha]$  one has

$$(1 - (1-x)^t)^2 \leq 1 - (1-x)^t \leq (tx)^{1-u+\alpha}.$$

As a result,

$$\begin{aligned} \mathcal{J}_1 &\leq 2\sigma^2 \frac{S+1}{(T-S)^2} \text{Tr}[H^{u-1} (1 - Q^{T-S})^2] \\ &\leq 2\sigma^2 \gamma^{1-u+\alpha} \text{Tr}[H^\alpha] (S+1) (T-S)^{\alpha-1-u}. \end{aligned} \quad (54)$$

Similarly,

$$\tilde{A}_t = \frac{H^{-1}}{T-S} (1 - Q^{T-t})$$

and

$$\begin{aligned} \text{Tr}[H^{u-1}(1 - Q^{T-t})^2] &= \sum_{j \in \mathbb{N}} \sigma_j^{u-1} (1 - (1 - \gamma \sigma_j)^{T-t})^2 \\ &\leq \gamma^{1-u+\alpha} \text{Tr}[H^\alpha] (T-t)^{1-u+\alpha}. \end{aligned}$$

Hence, since  $1 - u + \alpha > 0$  we find

$$\begin{aligned} \mathcal{J}_2 &\leq \gamma^{1-u+\alpha} \frac{2\sigma^2}{(T-S)^2} \text{Tr}[H^\alpha] \sum_{t=S+1}^{T-1} (T-t)^{1-u+\alpha} \\ &= \gamma^{1-u+\alpha} \frac{2\sigma^2}{(T-S)^2} \text{Tr}[H^\alpha] \sum_{t=1}^{T-S-1} t^{1-u+\alpha} \\ &\leq \gamma^{1-u+\alpha} \frac{2\sigma^2}{(T-S)^2} \text{Tr}[H^\alpha] (T-S-1)^{2-u+\alpha} \\ &\leq 2\sigma^2 \gamma^{1-u+\alpha} \text{Tr}[H^\alpha] (T-S)^{\alpha-u}. \end{aligned} \tag{55}$$

The result follows by combining (55), (54) and (53).  $\square$

## D.2 Proof of Proposition 5

**Perturbation Argument.** Relating the semi-stochastic recursion (51) to the fully stochastic recursion in (48) is based on the perturbation idea from [1], which has been also applied in [10] and in [26] in a similar context. For sake of completeness we give a brief summary.

For  $r \geq 0$  we introduce the sequence  $(\mu_t^r)_t$

$$\mu_{t+1}^r = (I - H)\mu_t^r + \gamma \Xi_{t+1}^r,$$

where  $\Xi_t^0 = \xi_t$  and for  $r \geq 0$

$$\Xi_{t+1}^{r+1} = (H - \hat{H}_t)\mu_t^r.$$

We further let  $\eta_t^r = \mu_t - \sum_{j=0}^r \mu_t^j$  which follows the recursion

$$\eta_{t+1}^r = (I - \hat{H}_t)\eta_t^r + \gamma \Xi_{t+1}^{r+1}.$$

From Lemma 2 in [26]<sup>3</sup> we have for any  $r \geq 0$

$$\mathbb{E}[\mu_t^r \otimes \mu_t^r] \preceq \gamma^{r+1} \kappa^{2r} \sigma^2 I. \tag{56}$$

and

$$\mathbb{E}[\Xi_t^r \otimes \Xi_t^r] \preceq \gamma^r \kappa^{2r} \sigma^2 H. \tag{57}$$

Bounding  $(\eta_t^r)_t$  is then done by applying the next Lemma, being an easy extension of Lemma 3 in [26] to tail-averaging.

---

<sup>3</sup>Lemma 2 in [26] is shown in the special case where  $\hat{H}_t = z_t \otimes z_t$  for i.i.d. observations  $z_t \in \mathcal{H}$ , but the proof of (56) and (57) is literally the same.

**Lemma 5** (Rough Bound SGD Recursion). *Consider the SGD recursion given in (48), satisfying (50). Assume further  $\gamma\kappa^2 < 1$ . For any  $1 \leq T$ ,  $0 \leq S \leq T - 1$  we have*

$$\mathbb{E} \left[ \|H^{\frac{u}{2}} \bar{\mu}_{S,T}\|^2 \right] \leq \sigma^2 \gamma^2 \kappa^u \frac{\text{Tr}[H]}{2(T-S)} \Delta(S, T).$$

where  $\Delta(S, T) = T(T+1) - S(S+1)$ . In particular, given  $1 \leq K$  and if  $0 \leq S \leq \frac{K-1}{K+1} T$  one has

$$\Delta(S, T) = T(T+1) - S(S+1) \leq K(T-S)^2.$$

*Proof of Lemma 5.* Following the arguments given in the proof of Lemma 3 in [26] we get

$$\mathbb{E} \left[ \|H^{\frac{u}{2}} \mu_t\|^2 \right] \leq \sigma^2 \gamma^2 \kappa^u \text{Tr}[H] t.$$

By convexity, this leads to

$$\begin{aligned} \mathbb{E} \left[ \|H^{\frac{u}{2}} \bar{\mu}_{S,T}\|^2 \right] &\leq \frac{1}{T-S} \sum_{t=S+1}^T \mathbb{E} \left[ \|H^{\frac{u}{2}} \mu_t\|^2 \right] \\ &\leq \sigma^2 \gamma^2 \kappa^u \frac{\text{Tr}[H]}{T-S} \sum_{t=S+1}^T t \\ &= \sigma^2 \gamma^2 \kappa^u \frac{\text{Tr}[H]}{2(T-S)} (T(T+1) - S(S+1)). \end{aligned}$$

□

**Proof of Proposition 5.** With these preparations we prove Proposition 5, applying the above described perturbation method. More precisely, we decompose

$$\bar{\mu}_{S,T} = \sum_{j=0}^r \bar{\mu}_{S,T}^j + \bar{\eta}_{S,T}^r$$

and have

$$\mathbb{E} \left[ \|H^{\frac{u}{2}} \bar{\mu}_{S,T}\|^2 \right]^{1/2} \leq \sum_{j=0}^r \mathbb{E} \left[ \|H^{\frac{u}{2}} \bar{\mu}_{S,T}^j\|^2 \right]^{1/2} + \mathbb{E} \left[ \|H^{\frac{u}{2}} \bar{\eta}_{S,T}^r\|^2 \right]^{1/2}. \quad (58)$$

The first term in (58) we apply Lemma 4 and (57). Denoting

$$\Lambda(S, T) = 4\sigma^2 \text{Tr}[H^\alpha] \gamma^{1-u+\alpha} (T-S)^{\alpha-u} \left( 1 + \frac{S+1}{T-S} \right)$$

we get with  $\gamma\kappa^2 \leq 1/4$

$$\begin{aligned} \sum_{j=0}^r \mathbb{E} \left[ \|H^{\frac{u}{2}} \bar{\mu}_{S,T}^j\|^2 \right]^{1/2} &\leq \sum_{j=0}^r (\gamma^j \kappa^{2j} \Lambda(S, T))^{1/2} \\ &= \sqrt{\Lambda(S, T)} \sum_{j=0}^r (\gamma\kappa^2)^{j/2} \\ &\leq \frac{\sqrt{\Lambda(S, T)}}{1 - \sqrt{\gamma\kappa^2}} \\ &\leq 2 \sqrt{\Lambda(S, T)}. \end{aligned} \quad (59)$$

For bounding the second term in (58) we apply the rough SGD recursion bound from Lemma 5 and (57). Since  $\gamma\kappa^2 < 1$ , we find as  $r \rightarrow \infty$

$$\mathbb{E} \left[ \|H^{\frac{u}{2}} \bar{\eta}_{S,T}^r\|^2 \right]^{1/2} \leq \left( \gamma^{2+r} \kappa^{u+2r} \sigma^2 \frac{\text{Tr}[H]}{2(T-S)} \Delta(S,T) \right)^{1/2} \rightarrow 0. \quad (60)$$

The final result follows by combining (60) and (59) with (58).

## E SGD Variance Term

Given  $b \in [n]$  the mini-batch SGD recursion is given by

$$w_{t+1} = w_t + \gamma \frac{1}{b} \sum_{i=b(t-1)+1}^{bt} (\langle w_t, x_{j_i} \rangle_{\mathcal{H}} - y_{j_i}) x_{j_i}, \quad t = 1, \dots, T,$$

with  $w_0 = 0$ ,  $\gamma > 0$  a constant step-size<sup>4</sup> and where  $j_1, \dots, j_{bT}$  are i.i.d. random variables, distributed according to the uniform distribution on  $[n]$ .

We analyze tail-averaged mini-batch SGD. More precisely, for  $0 \leq S \leq T - 1$  the algorithm under consideration is

$$\bar{w}_{S,T} := \frac{1}{T-S} \sum_{t=S+1}^T w_t.$$

For ease of notation we suppress dependence on  $b$ .

Recall the GD recursion

$$v_{t+1} = v_t - \gamma \frac{1}{n} \sum_{j=1}^n (\langle v_t, x_j \rangle_{\mathcal{H}} - y_j) x_j.$$

Denoting

$$\hat{\Sigma}_t = \frac{1}{b} \sum_{i=b(t-1)+1}^{bt} x_{j_i} \otimes x_{j_i}, \quad \hat{h}_t = \frac{1}{b} \sum_{i=b(t-1)+1}^{bt} y_{j_i} x_{j_i}$$

for any  $t \geq 1$ , we have

$$w_{t+1} - v_{t+1} = \left( I - \gamma \hat{\Sigma}_{t+1} \right) (w_t - v_t) + \gamma \xi_{t+1},$$

where we define  $\xi_{t+1} = \xi_{t+1}^{(1)} + \xi_{t+1}^{(2)}$  and

$$\xi_{t+1}^{(1)} = (\hat{\Sigma} - \hat{\Sigma}_{t+1}) v_t, \quad \xi_{t+1}^{(2)} = \hat{h}_{t+1} - \hat{h}. \quad (61)$$

Denoting by  $\mathcal{G}_n$  the  $\sigma$ -field generated by the data, we have for any  $t \geq 1$

$$\mathbb{E} \left[ \xi_{t+1}^{(1)} \mid \mathcal{F}_t, \mathcal{G}_n \right] = \mathbb{E} \left[ \xi_{t+1}^{(2)} \mid \mathcal{F}_t, \mathcal{G}_n \right] = 0$$

almost surely. Thus, the difference  $(\mu_t)_t = (w_t - v_t)_t$  follows a recursion as in (48), with  $\hat{Q}_t = I - \gamma \hat{\Sigma}_{t+1}$ .

<sup>4</sup>constant means independent of the iteration  $t$ , but possibly depending on  $n$

**Proposition 6.** Let  $\alpha \in (0, 1]$ ,  $\gamma\kappa^2 \leq 1/4$  and  $n$  be sufficiently large. Set  $L = T - S$ .

$$\mathbb{E} \left[ \|\Sigma^{\frac{1}{2}}(\bar{w}_{S,T} - \bar{v}_{S,T})\|^2 \right] \leq 32C_* \frac{\gamma^\alpha T r[\Sigma^\alpha]}{bL^{1-\alpha}} \Upsilon(S, T) + 32\gamma^2 \kappa^4 M^2 \frac{\tilde{\Delta}(S, T)^2}{L} \delta_n ,$$

with  $C_* = \kappa^4(2\|w_*\| + 1)^2 + M^2$ ,

$$\Upsilon(S, T) = 1 + \frac{S+1}{L} ,$$

$$\tilde{\Delta}(S, T) = \frac{1}{6}(T(T+1)(2T+1) - S(S+1)(2S+1))$$

and

$$\delta_n = 2 \exp \left( -a \sqrt{\frac{n}{\gamma T \mathcal{N}(1/\gamma T)}} \right) ,$$

for some  $a > 0$ . If additionally  $1 \leq K$  and  $1 \leq T$ ,  $0 \leq S \leq T-1$  satisfy  $S \leq \frac{K-1}{K+1} T$ , we have

$$\mathbb{E} \left[ \|\Sigma^{\frac{1}{2}}(\bar{w}_{S,T} - \bar{v}_{S,T})\|^2 \right] \leq 64C_* K \frac{\gamma^\alpha T r[\Sigma^\alpha]}{bL^{1-\alpha}} + 128\gamma^2 \kappa^4 M^2 K^4 L^5 \delta_n .$$

## E.1 Proof of Proposition 6

For proving Proposition 6 we aim at applying Proposition 5 and show that all assumptions are satisfied by stating a series of Lemmata. The first one provides an upper bound for the covariance of the noise process.

**Lemma 6.** Assume  $|Y| \leq M$  a.s. . For any  $t = S, \dots, T$  we have almost surely

$$\mathbb{E} \left[ \xi_{t+1}^{(1)} \otimes \xi_{t+1}^{(1)} \mid \mathcal{G}_n \right] \preceq \frac{\kappa^4}{b} \|v_t\|^2 \hat{\Sigma}$$

and

$$\mathbb{E} \left[ \xi_{t+1}^{(2)} \otimes \xi_{t+1}^{(2)} \mid \mathcal{G}_n \right] \preceq \frac{M^2}{b} \hat{\Sigma} .$$

Here, expectation is taken with respect to the  $b$ -fold uniform distribution on  $[n]$  in step  $t+1$ .

*Proof of Lemma 6.* Recall that

$$\xi_{t+1}^{(1)} = (\hat{\Sigma} - \hat{\Sigma}_{t+1})v_t = \frac{1}{b} \sum_{i=bt+1}^{b(t+1)} \tilde{\xi}_i ,$$

with

$$\tilde{\xi}_i := \hat{\Sigma}v_t - \langle v_t, x_{j_i} \rangle x_{j_i} .$$

By independence, we have

$$\begin{aligned} \mathbb{E} \left[ \xi_{t+1}^{(1)} \otimes \xi_{t+1}^{(1)} \mid \mathcal{G}_n \right] &= \frac{1}{b^2} \sum_{i, i'} \mathbb{E} \left[ \tilde{\xi}_i \otimes \tilde{\xi}_{i'} \mid \mathcal{G}_n \right] \\ &= \frac{1}{b^2} \sum_i \mathbb{E} \left[ \tilde{\xi}_i \otimes \tilde{\xi}_i \mid \mathcal{G}_n \right] . \end{aligned}$$

The first part follows then by

$$\mathbb{E} \left[ \tilde{\xi}_i \otimes \tilde{\xi}_i \mid \mathcal{G}_n \right] \preceq \mathbb{E} \left[ \langle v_t, x_{j_i} \rangle_{\mathcal{G}}^2 x_{j_i} \otimes x_{j_i} \mid \mathcal{G}_n \right] \preceq \kappa^4 \|v_t\|^2 \hat{\Sigma} .$$

The second part of the Lemma follows by writing

$$\xi_{t+1}^{(2)} = \hat{h}_{t+1} - \hat{h} = \frac{1}{b} \sum_{i=b(t-1)+1}^{bt} \xi'_{j_i},$$

with  $\xi'_{j_i} = y_{j_i} x_{j_i} - \hat{h}$  and observing that

$$\mathbb{E}[\xi'_{j_i} \otimes \xi'_{j_i} | \mathcal{G}_n] \preceq \mathbb{E}[|y_{j_i}|^2 x_{j_i} \otimes x_{j_i} | \mathcal{G}_n] \preceq M^2 \hat{\Sigma}. \quad (62)$$

□

The next Lemma provides a uniform for the GD updates, leading to a uniform bound for the noise process.

**Lemma 7** (Uniform Bound Gradient Descent updates). *Assume  $|Y| \leq M$  a.s. and let  $\tilde{M} = \max(M, \kappa \|w_*\|)$  and  $\bar{\sigma} := 2\tilde{M}$ . For any  $\delta \in (0, 1]$  and for any  $S + 1 \leq t \leq T$ , with probability at least  $1 - \delta$  one has*

$$\|v_t\| \leq 2 \|w_*\| + 1,$$

provided

$$n \geq 64 \max\{\bar{\sigma}^2, \kappa \tilde{M}\} \log^2(2\delta^{-1}) \gamma T \max\{1, \mathcal{N}(1/\gamma T)\}.$$

Moreover, with probability at least  $1 - (T - S)\delta$  one has

$$\sup_{S+1 \leq t \leq T} \|v_t\| \leq 2 \|w_*\| + 1. \quad (63)$$

*Proof of Lemma 7.* We decompose

$$\|v_t\| \leq \|v_t - w_*\| + \|w_*\|.$$

For bounding the first term we apply the results in [5], decomposition (5.9) with eq. (5.17) and (5.22) for  $\lambda = \frac{1}{\gamma t}$ <sup>5</sup>. For that we need to ensure a moment condition

$$\mathbb{E}\left[|Y - \langle w_*, X \rangle|^l | X\right] \leq \frac{1}{2} l! \bar{\sigma}^2 \tilde{M}^{l-2} \quad \text{a.s.}, \quad (64)$$

for some  $\bar{\sigma}^2 > 0$ ,  $\tilde{M} < \infty$  and for any  $l \geq 2$ . Indeed, since  $|Y| \leq M$  a.s. and  $|\langle w_*, X \rangle| \leq \kappa \|w_*\|$ , we easily derive

$$\begin{aligned} \mathbb{E}\left[|Y - \langle w_*, X \rangle|^l | X\right] &\leq 2^{l-1} \left( \mathbb{E}\left[|Y|^l | X\right] + |\langle w_*, X \rangle|^l \right) \\ &\leq 2^{l-1} \left( M^l + (\kappa \|w_*\|)^l \right) \\ &\leq \frac{1}{2} l! \bar{\sigma}^2 \tilde{M}^{l-2} \quad \text{a.s.}, \end{aligned}$$

with  $\tilde{M} = \max(M, \kappa \|w_*\|)$  and  $\bar{\sigma} := 2\tilde{M}$ . Thus, with probability at least  $1 - \delta$

$$\begin{aligned} \|v_t - w_*\| &\leq \|w_*\| + 2 \log(2\delta^{-1}) \left( \frac{\kappa \tilde{M} \gamma t}{n} + \bar{\sigma} \sqrt{\frac{\gamma t \mathcal{N}(1/\gamma t)}{n}} \right) \\ &\leq \|w_*\| + 2 \log(2\delta^{-1}) \left( \frac{\kappa \tilde{M} \gamma T}{n} + \bar{\sigma} \sqrt{\frac{\gamma T \mathcal{N}(1/\gamma T)}{n}} \right). \end{aligned}$$

<sup>5</sup>The constant in eq. (5.17) equals one in case of GD.

Assuming

$$n \geq 64 \max\{\bar{\sigma}^2, \kappa \tilde{M}\} \log^2(2\delta^{-1}) \gamma T \max\{1, \mathcal{N}(1/\gamma T)\} \quad (65)$$

we find

$$2 \log(2\delta^{-1}) \bar{\sigma} \sqrt{\frac{\gamma T \mathcal{N}(1/\gamma T)}{n}} \leq \frac{1}{4}.$$

Moreover, the same condition also implies

$$n \geq 64 \kappa \tilde{M} \log(2\delta^{-1}) \gamma T$$

owing to the fact that  $2 \log(2\delta^{-1}) > 1$  and thus

$$2 \log(2\delta^{-1}) \frac{\kappa \tilde{M} \gamma T}{n} \leq \frac{1}{32}.$$

Hence,

$$\|v_t\| \leq 2\|w_*\| + \frac{1}{32} + \frac{1}{4} \leq 2\|w_*\| + 1,$$

with probability at least  $1 - \delta$ . The uniform bound in (63) follows from taking a union bound, i.e.

$$\left\{ \sup_{S+1 \leq t \leq T} \|v_t\| \geq 2\|w_*\| + 1 \right\} \subseteq \bigcup_{t=S+1}^T \{\|v_t\| \geq 2\|w_*\| + 1\}.$$

□

**Lemma 8** ([6], eq. (47)). *For any  $\delta \in (0, 1]$  and  $\lambda > 0$  satisfying*

$$n\lambda \geq 64\kappa^2 \log^2(2\delta^{-1}) \max\{1, \mathcal{N}(\lambda)\} \quad (66)$$

one has

$$\left\| \left( \hat{\Sigma} + \lambda \right)^{-1} (\Sigma + \lambda) \right\| \leq 2$$

with probability at least  $1 - \delta$ .

The following Lemma provides a rough bound for the tail-averaged updates, generalized from [26] to tail-averaging.

**Lemma 9** (Rough bound for averaged SGD variance). *Assume  $|Y| \leq M$  a.s. and  $\gamma\kappa^2 < 1$ . One has almost surely*

$$\|\bar{w}_{S,T} - \bar{v}_{S,T}\| \leq 4\gamma\kappa M \frac{\tilde{\Delta}(S,T)}{T-S},$$

where

$$\tilde{\Delta}(S,T) = \sum_{t=S+1}^T t^2 = \frac{1}{6} (T(T+1)(2T+1) - S(S+1)(2S+1)).$$

Moreover, if  $1 \leq K$ ,  $1 \leq T$ ,  $0 \leq S \leq T-1$  satisfy  $S \leq \frac{K-1}{K+1} T$ , one has

$$\tilde{\Delta}(S,T) \leq 2K^2(T-S)^3$$

and

$$\|\bar{w}_{S,T} - \bar{v}_{S,T}\| \leq 8\gamma\kappa M K^2 (T-S)^2,$$

almost surely.



*Proof of Lemma 9.* Recall that the gradient updates are given by  $v_0 = 0$  and

$$v_{t+1} = v_t - \gamma(\hat{\Sigma}v_t - \hat{h}) = \hat{Q}v_t + \gamma\hat{h},$$

with  $\hat{Q} = (1 - \gamma\hat{\Sigma})$ ,  $\|\hat{Q}\| < 1$  and  $\|\hat{h}\| \leq \kappa M$ . Thus,

$$\|v_{t+1}\| \leq \|v_t\| + \gamma\kappa M$$

and inductively one obtains

$$\|v_t\| \leq \gamma\kappa M t. \quad (67)$$

Let  $\mu_t = w_t - v_t$ . Starting with  $\mu_0 = 0$ , then  $(\mu_t)_t$  follows the recursion

$$\mu_{t+1} = \hat{Q}_{t+1}\mu_t + \gamma\xi_{t+1}, \quad \hat{Q}_{t+1} = (I - \gamma\hat{\Sigma}_{t+1}),$$

where  $\xi_{t+1} = \xi_{t+1}^{(1)} + \xi_{t+1}^{(2)}$  is defined in (22). By (67) and since  $\gamma\kappa^2 < 1$  we have

$$\|\xi_{t+1}^{(1)}\| \leq \|(\hat{\Sigma} - \hat{\Sigma}_{t+1})\| \|v_t\| \leq 2\gamma\kappa^3 M t < 2\kappa M t.$$

Furthermore,

$$\|\xi_{t+1}^{(2)}\| = \|\hat{h}_{t+1} - \hat{h}\| \leq 2\kappa M.$$

Using  $\|\hat{Q}_{t+1}\| < 1$ , one easily calculates

$$\|\mu_t\| \leq \gamma \sum_{j=1}^t \|\xi_j\| \leq 4\gamma\kappa M t^2.$$

Thus,

$$\|\bar{\mu}_{S,T}\| \leq \frac{4\gamma\kappa M}{T-S} \sum_{t=S+1}^T t^2 = 4\gamma\kappa M \frac{\tilde{\Delta}(S,T)}{T-S},$$

with

$$\tilde{\Delta}(S,T) = \sum_{t=S+1}^T t^2 = \frac{1}{6}(T(T+1)(2T+1) - S(S+1)(2S+1)).$$

Finally,

$$\tilde{\Delta}(S,T) \leq 2K^2(T-S)^3,$$

implied by  $S \leq T-1$  and  $S \leq \frac{K-1}{K+1} T$ .  $\square$

*Proof of Proposition 6.* We define the events

$$\mathcal{E}_1 = \left\{ \mathbf{x} \in \mathcal{X}^n : \left\| \left( \hat{\Sigma} + \lambda \right)^{-1/2} (\Sigma + \lambda)^{1/2} \right\|^2 \leq 2 \right\}, \quad (68)$$

where we set  $\lambda = \frac{1}{\gamma L}$  and

$$\mathcal{E}_2 = \left\{ (\mathbf{x}, \mathbf{y}) \in \mathcal{X}^n \times \mathcal{Y}^n : \forall t = S+1, \dots, T : \mathbb{E}[\xi_{t+1} \otimes \xi_{t+1} \mid \mathcal{G}_n] \preceq \frac{C_*}{b} \hat{\Sigma} \right\}, \quad (69)$$

with  $C_* = \kappa^4(2\|w_*\| + 1)^2 + M^2$ .

Denoting  $\bar{\mathcal{N}}(\lambda) = \max\{1, \mathcal{N}(\lambda)\}$ , Lemma 8 gives  $\mathbb{P}[\mathcal{E}_1^c] \leq \delta_1$ , provided

$$n \geq 64\kappa^2 \log^2(2\delta_1^{-1}) \gamma L \bar{\mathcal{N}}(1/\gamma L)$$

or, equivalently,

$$\delta_1 \geq 2 \exp\left(-a_1 \sqrt{\frac{n}{\gamma L \bar{N}(1/\gamma L)}}\right), \quad (70)$$

with  $a_1 = \frac{1}{8\kappa}$ . Similarly, applying Lemma 7 and Lemma 6 gives  $\mathbb{P}[\mathcal{E}_2^c] \leq L\delta_2$  if

$$n \geq C_{\kappa, \bar{M}, \bar{\sigma}} \log^2(2\delta_2^{-1}) \gamma T \bar{N}(1/\gamma T), \quad C_{\kappa, \bar{M}, \bar{\sigma}} = 64 \max\{\bar{\sigma}^2, \kappa \bar{M}\}$$

or equivalently

$$\delta_2 \geq 2 \exp\left(-a_2 \sqrt{\frac{n}{\gamma T \bar{N}(1/\gamma T)}}\right). \quad (71)$$

with  $a_2 = \frac{1}{\sqrt{C_{\kappa, \bar{M}, \bar{\sigma}}}}$ .

Setting  $\bar{\mu}_{S,T} = \bar{w}_{S,T} - \bar{v}_{S,T}$ , we decompose

$$\mathbb{E}\left[\|\Sigma^{\frac{1}{2}} \bar{\mu}_{S,T}\|^2\right] \leq \mathbb{E}\left[\|\Sigma^{\frac{1}{2}} \bar{\mu}_{S,T}\|^2 1_{\mathcal{E}_1 \cap \mathcal{E}_2}\right] + \mathbb{E}\left[\|\Sigma^{\frac{1}{2}} \bar{\mu}_{S,T}\|^2 1_{\mathcal{E}_1^c}\right] + \mathbb{E}\left[\|\Sigma^{\frac{1}{2}} \bar{\mu}_{S,T}\|^2 1_{\mathcal{E}_2^c}\right]. \quad (72)$$

For bounding the first term note that

$$\Sigma^{\frac{1}{2}} = \Sigma^{\frac{1}{2}} (\Sigma + \lambda)^{-\frac{1}{2}} (\Sigma + \lambda)^{\frac{1}{2}} (\hat{\Sigma} + \lambda)^{-\frac{1}{2}} (\hat{\Sigma} + \lambda)^{\frac{1}{2}},$$

where

$$\|\Sigma^{\frac{1}{2}} (\Sigma + \lambda)^{-\frac{1}{2}}\| \leq 1.$$

Thus, by definition of  $\mathcal{E}_1$  and  $\mathcal{E}_2$ , using  $\|(\hat{\Sigma} + \lambda)^{\frac{1}{2}} u\|^2 = \|\hat{\Sigma}^{\frac{1}{2}} u\|^2 + \lambda \|u\|^2$ , we find with  $\lambda = \frac{1}{\gamma(T-S)}$  and Proposition 5 with  $\sigma^2 = C_*/b$

$$\begin{aligned} \mathbb{E}\left[\|\Sigma^{\frac{1}{2}} \bar{\mu}_{S,T}\|^2 1_{\mathcal{E}_1 \cap \mathcal{E}_2}\right] &\leq 2\mathbb{E}\left[\|\hat{\Sigma}^{\frac{1}{2}} \bar{\mu}_{S,T}\|^2\right] + \frac{2}{\gamma(T-S)} \mathbb{E}\left[\|\bar{\mu}_{S,T}\|^2\right] \\ &\leq 32C_* \frac{\gamma^\alpha \Upsilon(S, T)}{bL^{1-\alpha}} \mathbb{E}\left[\text{Tr}\left[\hat{\Sigma}^\alpha\right]\right] \\ &\leq 32C_* \frac{\gamma^\alpha \Upsilon(S, T)}{bL^{1-\alpha}} \text{Tr}\left[\Sigma^\alpha\right]. \end{aligned} \quad (73)$$

In the last step we apply Jensen's inequality, giving  $\mathbb{E}\left[\text{Tr}\left[\hat{\Sigma}^\alpha\right]\right] \leq \text{Tr}\left[\Sigma^\alpha\right]$ .

For bounding the second and third term recall that  $\|\Sigma^{\frac{1}{2}}\|^2 \leq \kappa^2$ . We have by Lemma 9

$$\|\Sigma^{\frac{1}{2}} \bar{\mu}_{S,T}\|^2 \leq 16\gamma^2 \kappa^4 M^2 \frac{\tilde{\Delta}(S, T)^2}{L^2}.$$

Hence,

$$\mathbb{E}\left[\|\Sigma^{\frac{1}{2}} \bar{\mu}_{S,T}\|^2 1_{\mathcal{E}_1^c}\right] \leq 16\gamma^2 \kappa^4 M^2 \frac{\tilde{\Delta}(S, T)^2}{L^2} \delta_1 \quad (74)$$

and

$$\mathbb{E}\left[\|\Sigma^{\frac{1}{2}} \bar{\mu}_{S,T}\|^2 1_{\mathcal{E}_2^c}\right] \leq 16\gamma^2 \kappa^4 M^2 \frac{\tilde{\Delta}(S, T)^2}{L^2} L \delta_2. \quad (75)$$

The result follows from collecting (75), (74), (73) and (72) and by choosing

$$\delta_n := \max\{\delta_1, \delta_2\} = 2 \exp\left(-a \sqrt{\frac{n}{\gamma T \bar{N}(1/\gamma T)}}\right), \quad (76)$$

with  $a = \min\{a_1, a_2\}$ . Note that we also use the fact that  $\gamma t \bar{N}(1/\gamma t)$  is increasing in  $t$  and  $L \leq T$ .

If additionally  $1 \leq K$  and  $1 \leq T$ ,  $0 \leq S \leq T - 1$  satisfy  $S \leq \frac{K-1}{K+1} T$ , we have

$$\Upsilon(S, T) \leq 1 + K \leq 2K, \quad \tilde{\Delta}(S, T) \leq 2K^2 L^3.$$

this gives

$$\mathbb{E} \left[ \|\Sigma^{\frac{1}{2}} \bar{\mu}_{S, T}\|^2 \right] \leq 64C_* K \frac{\gamma^\alpha T r[\Sigma^\alpha]}{bL^{1-\alpha}} + 128\gamma^2 \kappa^4 M^2 K^4 L^5 \delta_n.$$

□

## F Main Results Tail-Averaging SGD

From Theorem 2 and Proposition 6 combined with decomposition (28) we obtain

**Theorem 3.** *Let  $\alpha \in (0, 1]$ ,  $1 \leq T$ ,  $0 \leq S \leq T - 1$  and Assumptions 1, 2 hold. Assume  $\gamma\kappa^2 < 1/4$ . Set  $L = T - S$  and*

$$\delta_n = 2 \exp\left(-a \sqrt{\frac{n}{\gamma T \mathcal{N}(1/\gamma T)}}\right),$$

with  $a > 0$  given in (76). Then

$$\begin{aligned} \mathbb{E} \left[ \|\Sigma^{\frac{1}{2}}(\bar{w}_{S, T} - w_*)\|^2 \right] &\lesssim \frac{\gamma^\alpha}{bL^{1-\alpha}} Tr[\Sigma^\alpha] + (\gamma L)^{-2(r+1/2)} + \frac{\mathcal{N}(\frac{1}{\gamma L})}{n} \\ &+ \frac{\gamma L}{n^2} + \frac{(\gamma L)^{-2r}}{n} + \frac{1}{n} + \gamma^2 L^5 \delta_n, \end{aligned}$$

under each of the following assumptions:

1.  $0 \leq r \leq 1/2$  and  $1 \leq K$ ,  $0 \leq S \leq \frac{K-1}{K+1} T$ ,
2.  $1/2 < r$ ,  $1 < K$ ,  $0 < S \leq \frac{K-1}{K+1} T$  and  $T \leq (K+1)S$ .

The constant hidden in  $\lesssim$  in the above bound depends on the model parameters  $\kappa, M, r, R, K$  given in the assumptions.

*Proof of Corollary 1.* Plugging in Assumptions 2 and 3 gives in either case

$$\begin{aligned} \mathbb{E} \left[ \|\Sigma^{1/2}(\bar{w}_{L_n} - w_*)\|^2 \right] &\lesssim \frac{\gamma_n^\alpha}{b_n L_n^{1-\alpha}} + (\gamma_n L_n)^{-2(r+1/2)} + \frac{(\gamma_n L_n)^\nu}{n} \\ &+ \frac{\gamma_n L_n}{n^2} + \frac{(\gamma_n L_n)^{-2r}}{n} + \frac{1}{n} + \gamma_n^2 L_n^5 \delta_n, \end{aligned}$$

As in the proof of Corollary 2 we have as  $n \rightarrow \infty$

$$n^{-1} = o\left((\gamma_n L_n)^{-2(r+1/2)}\right), \quad \frac{\gamma_n L_n}{n^2} = o\left((\gamma_n L_n)^{-2(r+1/2)}\right)$$

and

$$\frac{(\gamma_n L_n)^{-2r}}{n} = o\left((\gamma_n L_n)^{-2(r+1/2)}\right),$$

so we can disregard the terms  $n^{-1}$ ,  $\frac{\gamma_n L_n}{n^2}$ ,  $\frac{(\gamma_n L_n)^{-2r}}{n}$  for  $n$  large enough. Furthermore,  $\delta_n$  satisfies

$$\delta_n \lesssim \exp\left(-a \sqrt{\frac{n}{(\gamma_n T_n)^{\nu+1}}}\right) = \exp\left(-a n^{\frac{1}{2}(1-\frac{\nu+1}{2r+1+\nu})}\right),$$

showing

$$\gamma_n^2 L_n^5 \delta_n = o\left((\gamma_n L_n)^{-2(r+1/2)}\right)$$

as  $n \rightarrow \infty$  since  $1 - \frac{\nu+1}{2r+1+\nu} > 0$  and  $\delta_n$  decreases exponentially fast (note that we require  $S_n$  to be of the same order as  $T_n$ ). Furthermore, the choice

$$\gamma_n L_n \simeq n^{\frac{1}{2r+1+\nu}}$$

precisely balances the two terms  $(\gamma_n L_n)^{-2(r+1/2)}$  and  $\frac{(\gamma_n L_n)^\nu}{n}$ , so the remaining leading order terms are

$$\mathbb{E}\left[\|\Sigma^{1/2}(\bar{w}_{L_n} - w_*)\|^2\right] \lesssim \frac{\gamma_n^\alpha}{b_n L_n^{1-\alpha}} + (\gamma_n L_n)^{-2(r+1/2)}.$$

Finally, choosing  $\alpha = \nu$ , a calculation shows that all choices of  $b_n, (\gamma_n L_n)$  are balancing the two remaining terms.  $\square$

## G Auxiliary Technical Lemmata

### G.1 Probabilistic Ones

**Proposition 7** ([14], Proposition 1). *Define*

$$\mathcal{B}_n(\lambda) := \left[1 + 4\kappa^2 \left(\frac{\kappa}{n\lambda} + \sqrt{\frac{\mathcal{N}(\lambda)}{n\lambda}}\right)^2\right]. \quad (77)$$

For any  $\lambda > 0$ ,  $\delta \in (0, 1]$ , with probability at least  $1 - \delta$  one has

$$\left\|\left(\hat{\Sigma} + \lambda\right)^{-1}(\Sigma + \lambda)\right\| \leq 8 \log^2(2\delta^{-1}) \mathcal{B}_n(\lambda). \quad (78)$$

**Corollary 3.** *Let  $\delta \in (0, 1]$  and assume that*

$$n\lambda \geq 16\kappa^2 \max\{1, \mathcal{N}(\lambda)\}. \quad (79)$$

Then

$$\mathcal{B}_n(\lambda) \leq 2.$$

In particular,

$$\left\|\left(\hat{\Sigma} + \lambda\right)^{-1}(\Sigma + \lambda)\right\| \leq 16 \log^2(2\delta^{-1})$$

holds with probability at least  $1 - \delta$ .

*Proof of Corollary 3.* Assumption (79) immediately gives

$$\sqrt{\frac{\mathcal{N}(\lambda)}{n\lambda}} \leq \frac{1}{4\kappa}$$

as well as

$$\frac{\kappa}{\lambda n} \leq \frac{1}{4\kappa}.$$

The result then follows by plugging these bounds into (77).  $\square$

**Lemma 10.** *Let  $\lambda > 0$  and assume that*

$$n\lambda \geq 16\kappa^2 \max\{1, \mathcal{N}(\lambda)\}. \quad (80)$$

*For any  $w \in \mathcal{H}$  and  $\delta \in (0, 1]$ , one has with probability at least  $1 - \delta$*

$$\|\Sigma^{\frac{1}{2}}w\| \leq 4\log(2\delta^{-1}) \|(\hat{\Sigma} + \lambda)^{\frac{1}{2}}w\|.$$

*Proof of Lemma 10.* Applying Corollary 3, we find

$$\begin{aligned} \|\Sigma^{\frac{1}{2}}w\| &\leq \|\Sigma^{\frac{1}{2}}(\Sigma + \lambda)^{-\frac{1}{2}}\| \|(\hat{\Sigma} + \lambda)^{-\frac{1}{2}}(\Sigma + \lambda)^{\frac{1}{2}}\| \|(\hat{\Sigma} + \lambda)^{\frac{1}{2}}w\| \\ &\leq 4\log(2\delta^{-1}) \|(\hat{\Sigma} + \lambda)^{\frac{1}{2}}w\|. \end{aligned}$$

□

**Lemma 11.** *Let  $X$  be a nonnegative random variable with  $\mathbb{P}[X > C \log^u(k\delta^{-1})] < \delta$  for any  $\delta \in (0, 1]$ . Then  $\mathbb{E}[X] \leq \frac{C}{k}u\Gamma(u)$ , where  $\Gamma$  denotes the Gamma-function.*

*Proof.* Apply  $\mathbb{E}[X] = \int_0^\infty \mathbb{P}[X > t]dt$ .

□

## G.2 Miscellaneous

**Lemma 12.** *For any  $0 \leq S \leq T$  and for any  $a \in [0, 1]$  one has*

$$T^{a+1} - S^{a+1} \leq (T + S)(T - S)^a. \quad (81)$$

*Proof of Lemma 12.* Rewriting (81) to

$$1 - \left(\frac{S}{T}\right)^{a+1} \leq \left(1 + \frac{S}{T}\right) \left(1 - \frac{S}{T}\right)^a$$

shows that it is sufficient to show that

$$h_a(u) := (1 + u)(1 - u)^a + u^{a+1} - 1 \geq 0$$

for any  $u \in [0, 1]$ . This follows by observing that  $h_0(u) \equiv 0$ ,  $h_1(u) = 2u$ . Moreover,  $h_a$  is concave if  $a \in (0, 1)$ , satisfying  $h_a(0) = h_a(1) = 0$ . □

**Lemma 13.** *Let  $(a_k)_k$  and  $(\xi_k)_k$  be two sequences, then*

$$\sum_{t=S+1}^T \sum_{k=0}^{t-1} a_{t-1-k} \xi_k = \sum_{t=0}^S \left( \sum_{k=S-t}^{T-(t+1)} a_k \right) \xi_t + \sum_{t=S+1}^{T-1} \left( \sum_{k=0}^{T-(t+1)} a_k \right) \xi_t.$$

**Lemma 14.** 1. *Let  $\varphi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$  monotonically non-decreasing. Then*

$$\sum_{t=S}^T \varphi(t) \leq \int_S^{T+1} \varphi(t) dt \leq \sum_{t=S}^T \varphi(t+1).$$

2. *Let  $\varphi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$  monotonically non-increasing. Then*

$$\sum_{t=S}^T \varphi(t+1) \leq \int_S^{T+1} \varphi(t) dt \leq \sum_{t=S}^T \varphi(t).$$