

Implicit Regularization or Implicit Conditioning? Exact Risk Trajectories of SGD in High Dimensions

Supplementary material

The appendix is organized into five sections as follows:

1. Appendix A builds on the background and preliminaries discussed in the main text. We include additional related work, an extended discussion on diffusion processes and Volterra equations.
2. Appendix B expands upon the assumptions/setting around (1) and it discusses some motivating applications such as in- and out- of distribution expected risk and random features. Moreover, we discuss the equivalence of homogenized SGD and SGD, see Theorem 1.
3. Appendix C gives results for the streaming SGD setting. In particular, we establish a Volterra equation.
4. Appendix D introduces a general Volterra class of equations, called the Volterra SLD class, that encompasses homogenized SGD and its Volterra dynamics in the multi-pass setting (Section 3) and streaming (Section C). This general Volterra class allows for more types of additive noise. We prove in this section that the Volterra SLD class concentrates around its mean; thereby deriving the proof of Theorem 2.
5. We prove in Appendix E the limiting risk values for the Volterra SLD class (Theorem 12 (constant learning rate) and Theorem 13 (time dependent learning rate)). These two theorems immediately imply the limiting risk values for homogenized SGD in the multi-pass and streaming settings, see Theorems 3 and 10 respectively.
6. Appendix F discusses the exact asymptotic convergence rates for SGD and full batch momentum algorithms on high-dimensional ℓ^2 -regularized least squares problems. The results in this section (e.g., Theorems 4 and 5) were shown in a series of papers [58–60] that explored exact trajectories of loss function.
7. Appendix G contains details on the simulations.

Potential societal impacts. The results presented in this paper concern the analysis of existing methods on a simple ℓ^2 -regularized least squares problems. The results are theoretical and we do not anticipate any direct ethical and societal issues. We believe the results will be used by machine learning practitioners and we encourage them to use it to build a more just, prosperous world.

A Extended preliminaries and background.

We include some additional preliminary background information.

Related work. We highlight recent progress of advances in analyzing the excess risk of SGD. As these areas are highly active, we present below a non-exhaustive list of the current progress, and particularly focus on the least-squares setting.

In the literature, convergence guarantees and risk bounds are available for analyzing SGD and its variants [71, 56, 53, 16, 79]. A popular paradigm for analyzing risk bounds of SGD is the one-pass or streaming setting where one supposes that the gradient estimators are independent with a common distribution [24, 30]. Such a setting was considered in a series of works [30, 18] which explored ‘one-pass’ SGD on a least-squares under a design condition on the data matrix. Extending this idea, [87] provide upper and lower excess risk bounds for constant stepsize SGD on the ℓ^2 -regularized least-squares problem which extends the work of [8, 78]. These bounds are characterized by the full eigenspectrum of the population covariance matrix. Beyond the confines of the streaming setting, much less is known about the risk bounds for multi-pass SGD [45, 64, 39, 88]. Like in this work, [45, 64, 88] consider the simplified setting of analyzing the behavior of multi-pass SGD on a (high-dimensional) ℓ^2 -regularized least-squares problem. In contrast to our exact dynamics of the excess risk, previous works provide only bounds. To the best of knowledge, the only other paper which

rigorously computed exact dynamics in the high-dimensional setting was [59]. They were the first to show dynamics on the *training loss* of a least-squares problem, assuming a left orthogonal invariance condition. In this paper, we remove the left orthogonal invariance condition and extend dynamics to generalization performance. Other approaches to analyzing generalization performance of multi-pass SGD include uniform stability (see c.f. [25]). Excess risk and uniform stability are related through a (loose) triangle inequality; we preferred to focus on the excess risk in this paper and leave discussions of uniform stability for future directions.

The exact features of SGD that are responsible for the success of SGD on high-dimensional problems are the subject of extensive research. These features are often labelled under the umbrella term the *implicit regularization effects of SGD*. The implicit regularization mechanism has been primarily hypothesized to exist in highly nonconvex settings. In such a setting, empirical observations have led to the conclusion that the noise generated by small batch [32, 27] and/or large learning rate [40] SGD leads to better generalization performance. A proposed mechanism for this improvement is that the stochasticity inherent in SGD allows the optimizer to escape traps (c.f. [13] or [85]) which poorly generalize and are stationarity points for gradient flow. A related point of view on the proposed escape mechanism is that SGD demonstrates a preference for flat minima, long considered to be a preferable solution for generalization properties [26, 80]. Using a continuous time model, in a certain noise regime, [80] shows that SGD prefers flat minima which is different from the flat minima flat minimum selection of continuous time SGD with homogeneous noise. For the least-squares problem, multi-pass SGD converges to the minimal norm solution [23, 57, 84] which is widely cited as the implicit bias of SGD [33, 17]. Moreover other works have shown for the least-squares problem with ℓ^2 -regularization results similar to Lemma 1, saying that the generalization performance of SGD always performs worse than gradient descent [88].

Diffusion approximations and SGD. In this work, we analyze a new mathematical tool, homogenized SGD, introduced simultaneously in [58, 52] (a similar related SDE was also analyzed in [63]). This SDE is the high-dimensional equivalence of SGD. The technique of using SDEs to analyze SGD is not new (see, for example, stochastic modified equation (SME) [42, 48] and Langevin dynamics [15] and other SDE formulations and interpretations [31, 36, 46, 7]). Diffusion approximations to SGD have a long history. In the stochastic approximation literature, it appears as a natural counterpart to ODE methods (c.f. [36], [47]). However, these are methods that require the vanishing learning rate (such as $\gamma_k = 1/k$) and moreover, in the setup we have suggested here, the resulting SDE (which only arises in an asymptotic comparison, as is standard with stochastic approximation theory) has a vanishing diffusion term for such an aggressive learning-rate decay — the asymptotic trajectory of SGD is also approximated by the ODE gradient flow with time change $\Gamma(n) \approx \log n$.

A more natural point of comparison is the *stochastic modified equation* of [48, 42, 41] which has been rigorously compared to the behavior of SGD [43]. To make a comparison with [43, 28], we fix the learning rate γ and we rescale time to be on the order of epochs. With these changes, the SME solves,

$$d\mathcal{M}_t \stackrel{\text{def}}{=} -\gamma \nabla f(\mathcal{M}_t) dt + \gamma \sqrt{n \Sigma(\mathcal{M}_t)} dB_t, \quad \text{where} \quad \begin{cases} \Sigma(\mathbf{x}) \stackrel{\text{def}}{=} \mathbb{E}(\nabla g_I(\mathbf{x}) \otimes \nabla g_I(\mathbf{x})), \\ g_I(\mathbf{x}) \stackrel{\text{def}}{=} f_I(\mathbf{x}) - \mathbb{E} f_I(\mathbf{x}), \\ I \stackrel{\text{def}}{=} \text{Uniform}\{1, 2, \dots, n\}. \end{cases} \quad (13)$$

The diffusion matrix Σ of the SME is chosen to exactly match the covariance of the increments of SGD (4). When applied to the ℓ^2 -regularized problem (1), this matrix becomes (with \mathbf{a}_i the i -th row of \mathbf{A})

$$\Sigma(\mathbf{x}) = \frac{1}{n} \sum_{i=1}^n (\mathbf{a}_i \cdot \mathbf{x} - \mathbf{b}_i)^2 \mathbf{a}_i^T \mathbf{a}_i - \frac{1}{n^2} \mathbf{A}^T (\mathbf{A} \mathbf{x} - \mathbf{b}) (\mathbf{A} \mathbf{x} - \mathbf{b})^T \mathbf{A}.$$

HSGD and the SME can be compared by replacing second term by 0 and the first term by

$$n \Sigma(\mathbf{x}) \approx \frac{1}{n} \left(\sum_{i=1}^n (\mathbf{a}_i \cdot \mathbf{x} - \mathbf{b}_i)^2 \right) \times \left(\sum_{i=1}^n \mathbf{a}_i^T \mathbf{a}_i \right) = \frac{2}{n} \mathcal{L}(\mathbf{x}) \nabla^2 \mathcal{L}(\mathbf{x}),$$

which is the diffusion coefficient in HSGD.

The SME has been used for a variety of purposes, such as optimal learning rate scheduling [41], analysis of momentum terms [42], and prediction of test risk behavior [75]; however, analysis of

the SME is itself difficult, as the diffusion coefficient involves interactions between the functions f_i , and, to our knowledge, while the theory developed in [43] provides dimension-independent comparisons, the resulting SME has not been analyzed in any high-dimensional setting. Furthermore, the mathematical comparison which is proven in [43], on the time scale in (13), gives a comparison for time of order $\mathcal{O}(1/n)$ and with an error which is bounded by $\mathcal{O}(\gamma)$. As such, in a high-dimensional setting, the comparison that exists between SME and SGD only provides a non-vanishing error over a vanishing window of time.

Because the SME naturally matches the drift and diffusion matrix of SGD, the above remarks might lead to a conjecture that in fact no comparison is possible between SGD and SDEs; indeed, in a fixed-dimensional analysis, Yaida [83] showed that there is no small learning-rate limit of SGD that produces nontrivial stochastic behavior. In contrast, we will show that in a high-dimensional limit, this is precisely what occurs (although when univariate statistics of this high-dimensional SDE are taken, almost deterministic behavior is seen).

Background on Volterra equations. Volterra dynamics appear frequently in filtering, population dynamics, and the renewal problem, to name a few, and consequently their properties, particularly convergence and limiting behavior, are well-studied in the literature [22, 67]. Convolution-type Volterra equation, that is, equations of the form

$$\psi(t) = F(t) + \int_0^t K(t-s)\psi(s) ds, \quad t \geq 0, \quad (14)$$

where the forcing term $F : \mathbb{R}_{\geq 0} \rightarrow \widehat{\mathcal{R}}$ and the kernel is $K : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ with $\mathbb{R}_{\geq 0} = \{t \geq 0\}$, in particular, can be “solved” using Laplace transforms (must be able to invert the Laplace transform). For general Volterra equations, one replaces (14) is replaced with a general function on two variables, $K : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$.

The solution of ψ can be found by repeatedly convolving the forcing term $F(t)$ with the kernel K (provided $\sup_{t \geq 0} \sup_{s \geq 0} \mathcal{K}(t, s; \nabla^2 \mathcal{L})$ is bounded [22]), that is,

$$\psi(t) = F(t) + (K \star F)(t) + (K \star K \star F)(t) + \dots \quad \text{where} \quad \begin{cases} (K \star h)(t) \stackrel{\text{def}}{=} \int_0^t K(t, s)h(s) ds, \\ \forall h \in C([0, \infty)). \end{cases} \quad (15)$$

Moreover, numerical approximations to (15) can be found by taking a large but finite number of convolutions in the expression above. The boundedness of this solution corresponds precisely to learning rate choices for which SGD is convergent.

B Quasi-random assumptions on the data matrix, targets, and initialization

The data matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, target $\mathbf{b} \in \mathbb{R}^n$, and initialization $\mathbf{x}_0 \in \mathbb{R}^d$ may be deterministic or random; we formulate our theorems for deterministic matrix \mathbf{A} and vectors \mathbf{b} and \mathbf{x}_0 satisfying various assumptions, and in the applications of these theorems to statistical settings, we shall show that random \mathbf{A} and \mathbf{b} satisfy those assumptions. These assumptions are motivated by ERM and, in particular, when the augmented matrix $[\mathbf{A} \mid \mathbf{b}]$ has rows that are independent and sampled from some common distribution. We call these assumptions *quasi-random*.

As the problem (1) is homogeneous, we adopt the following normalization convention without loss of generality.

Assumption 3 (Data-target normalization). *There is a constant $C > 0$ independent of d and n such that the spectral norm of \mathbf{A} is bounded by C and the target vector $\mathbf{b} \in \mathbb{R}^n$ is normalized so that $\|\mathbf{b}\|^2 \leq C$.*

More importantly, we also assume that the data and targets resemble typical unstructured high-dimensional random matrices. One of the principal qualitative properties of high-dimensional random matrices is the *delocalization of their eigenvectors*, which refers to the statistical similarity of the eigenvectors to uniform random elements from the Euclidean sphere. The precise mathematical

description of this assumption is most easily given in terms of resolvent bounds. The resolvent $R(z; \mathbf{M})$ of a matrix $\mathbf{M} \in \mathbb{R}^{d \times d}$ is

$$R(z; \mathbf{M}) = (z\mathbf{I}_d - \mathbf{M})^{-1} \quad \text{for } z \in \mathbb{C}.$$

In terms of the resolvent, we suppose the following:

Assumption 4. *Suppose Ω is the contour enclosing $[0, 1 + \|\mathbf{A}\|^2]$ at distance $1/2$. Suppose there is a $\theta \in (0, \frac{1}{2})$ for which*

1. $\max_{z \in \Omega} \max_{1 \leq i \leq n} |\mathbf{e}_i^T R(z; \mathbf{A}\mathbf{A}^T) \mathbf{b}| \leq n^{\theta-1/2}$.
2. $\max_{z \in \Omega} \max_{1 \leq i \neq j \leq n} |\mathbf{e}_i^T R(z; \mathbf{A}\mathbf{A}^T) \mathbf{e}_j^T| \leq n^{\theta-1/2}$.
3. $\max_{z \in \Omega} \max_{1 \leq i \leq n} |\mathbf{e}_i^T R(z; \mathbf{A}\mathbf{A}^T) \mathbf{e}_i - \frac{1}{n} \text{tr} R(z; \mathbf{A}\mathbf{A}^T)| \leq n^{\theta-1/2}$.

Only the resolvent of $\mathbf{A}\mathbf{A}^T$ appears in these assumptions, and so in effect we are only assuming statistical properties on the left singular-vectors of \mathbf{A} . This assumption reflects the common formulation of ERM in which the rows of \mathbf{A} are independent, and so the left singular-vectors of \mathbf{A} are expected to be delocalized (under some mildness assumptions on the distributions of the rows). The first condition, which involves the interaction between $\mathbf{A}\mathbf{A}^T$ and \mathbf{b} , can be understood as requiring that \mathbf{b} is not too strongly aligned with the left singular-vectors of \mathbf{A} . The other two conditions can be viewed as corollaries of delocalization of the left singular-vectors.

As for the initialization \mathbf{x}_0 , we need to suppose that it, like \mathbf{b} , does not interact too strongly with the left singular-vectors of $\mathbf{A}^T \mathbf{A}$. In the spirit of Assumption 4, it suffices to assume the following:

Assumption 5. *Let Ω be the same contour as in Assumption 4 and let $\theta \in (0, \frac{1}{2})$. Then*

$$\max_{z \in \Omega} \max_{1 \leq i \leq d} |\mathbf{e}_i^T R(z; \mathbf{A}^T \mathbf{A}) \mathbf{x}_0| \leq n^{\theta-1/2}.$$

Note that, as a simple but common case, this assumption is surely satisfied for $\mathbf{x}_0 = \mathbf{0}$. In principle, this assumption is general enough to allow for \mathbf{x}_0 that are correlated with \mathbf{A} in a nontrivial way, but we do not have an application for such an initialization. For a large class of nonzero initializations independent from (\mathbf{A}, \mathbf{b}) , this assumption is satisfied, as a corollary of Assumption 4:

Lemma 2. *Suppose that Assumption 4 holds with some $\theta_0 \in (0, \frac{1}{2})$ and that \mathbf{x}_0 is chosen randomly, independent of (\mathbf{A}, \mathbf{b}) , and with independent coordinates in such a way that for some C independent of d or n*

$$\|\mathbb{E} \mathbf{x}_0\|_\infty \leq C/n \quad \text{and} \quad \max_i \|(\mathbf{x}_0 - \mathbb{E} \mathbf{x}_0)_i\|_{\psi_2}^2 \leq Cn^{2\theta_0-1}.$$

For any $\theta > \theta_0$, Assumption 5 holds with any $\theta > \theta_0$ on an event of probability tending to 1 as $n \rightarrow \infty$.

Note that this assumption allows for deterministic \mathbf{x}_0 having maximum norm $\mathcal{O}(1/n)$, as well as iid centered subgaussian vectors of Euclidean norm $\mathcal{O}(1)$.

To execute the mathematical comparison between SGD and HSGD, we require an additional assumption on the quadratic in the same spirit as Assumption 3.

Assumption 6 (Quadratic statistics). *Suppose $\mathcal{R} : \mathbb{R}^d \rightarrow \mathbb{R}$ is quadratic, i.e. there is a symmetric matrix $\mathbf{T} \in \mathbb{R}^{d \times d}$, a vector $\mathbf{u} \in \mathbb{R}^d$, and a constant $c \in \mathbb{R}$ so that*

$$\mathcal{R}(\mathbf{x}_t) = \frac{1}{2} \mathbf{x}_t^T \mathbf{T} \mathbf{x}_t + \mathbf{u}^T \mathbf{x}_t + c. \quad (16)$$

We also assume that \mathcal{R} satisfies Assumption 1. Moreover, we assume the following (for the same Ω and θ) as in Assumption 4:

$$\max_{z, y \in \Omega} \max_{1 \leq i \leq n} |\mathbf{e}_i^T \widehat{\mathbf{T}} \mathbf{e}_i - \frac{1}{n} \text{tr}(\widehat{\mathbf{T}})| \leq \|\mathbf{T}\|_{op} n^{-\epsilon} \quad \text{where} \quad \begin{cases} \widehat{\mathbf{T}} = R(z) \mathbf{T} R(y) + R(y) \mathbf{T} R(z), \\ R(z) = R(z; \mathbf{A}^T \mathbf{A}) \end{cases} \quad (17)$$

This assumption ensures that quadratic \mathcal{R} has a Hessian that is not too correlated with any of the left singular vectors of \mathbf{A} . Establishing Assumption 6 can be nontrivial in the cases when the quadratic has complicated dependence on \mathbf{A} . In simple cases, (especially for the case of the empirical risk and the norm) it follows automatically from Assumption 4.

Lemma 3. *Suppose that \mathcal{R} satisfies (16) with \mathbf{T} given by a polynomial p in $\mathbf{A}^T \mathbf{A}$ (especially \mathbf{I} and the monomial $\mathbf{A}^T \mathbf{A}$) having bounded coefficients, and suppose \mathbf{u} and c are norm bounded independently of n or d . Then supposing Assumptions 3 and 4 for some $\theta_0 \in (0, \frac{1}{2})$, for all n sufficiently large and for any $\theta > \theta_0$, Assumption 6 holds.*

For proofs of Lemma 2 and 3, see Section 2 in [61].

B.1 Motivating applications

Training loss and sample covariance matrices. One important (nonstatistical) quadratic statistic, which allows analysis of the optimization aspects of SGD in high dimensions, is the ℓ^2 -regularized loss function f in (1). Then provided that \mathbf{A}, \mathbf{b} satisfy Assumptions 3 and 4, \mathbf{x}_0 is iid subgaussian, Lemmas 2 and 3 and Theorems 1 and 2 show that $f(\mathbf{x}_k)$ concentrates around the solution of a Volterra integral equation. A natural setup under which Assumptions 3 and 4 are satisfied is the following:

Assumption 7. *Suppose $M > 0$ is a constant. Suppose that Σ is a positive semi-definite $d \times d$ matrix with $\text{tr } \Sigma = 1$ and $\|\Sigma\|_{op} \leq M/\sqrt{d} < \infty$. Suppose that \mathbf{A} is a random matrix $\mathbf{A} = \mathbf{Z}\sqrt{\Sigma}$ where \mathbf{Z} is an $n \times d$ matrix of independent, mean 0, variance 1 entries with subgaussian norm at most $M < \infty$, and suppose $n \leq Md$. Finally suppose that $\mathbf{b} = \mathbf{A}\beta + \xi$ for β, ξ iid centered subgaussian satisfying $\|\beta\|^2 = R$ and $\|\xi\|^2 = \frac{n}{d}\tilde{R}$.*

These assumptions naturally lead to random matrices that satisfy Assumption 7 with good probability:

Lemma 4. *If (\mathbf{A}, \mathbf{b}) satisfy Assumption 7, then (\mathbf{A}, \mathbf{b}) satisfies Assumptions 3 and 4 with probability tending to $1 - e^{-\Omega(d)}$.*

Hence, under these assumptions, we conclude:

Theorem 6. *Suppose (\mathbf{A}, \mathbf{b}) satisfy Assumption 7, $\delta > 0$ and \mathbf{x}_0 is iid centered subgaussian with $\mathbb{E} \|\mathbf{x}_0\|^2 = \hat{R}$. Then for some $\epsilon > 0$, for all $T > 0$, and for all $D > 0$ there is a $C > 0$ such that*

$$\Pr\left(\sup_{0 \leq t \leq T} \left\| \left(\frac{\mathcal{L}(\mathbf{x}_{[tn]}^{sgd})}{\frac{1}{2}\|\mathbf{x}_{[tn]}^{sgd} - \beta\|_2^2} \right) - \left(\frac{\Psi_t}{\Omega_t} \right) \right\|_2 > d^{-\epsilon} \right) \leq Cd^{-D},$$

where Ψ_t solves (10) and Ω_t solves (9) with $\mathcal{R} = \frac{1}{2}\|\cdot - \beta\|_2^2$.

We discuss generalization implications in the the next section.

Theorem 6 generalizes [59] in that it allows for varying training rates, adds a regularization parameter, and allows for non-orthogonally-invariant designs \mathbf{A} . We further note that under the assumptions of Theorem 6, we can further approximate the behavior of GF to show that

$$\begin{aligned} \Psi_t &= \mathcal{L}(\mathcal{X}_{\Gamma(t)}^{\text{gf}}) + \frac{1}{n} \int_0^t \gamma^2(s) \text{tr} \left((\mathbf{A}^T \mathbf{A})^2 e^{-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)(\Gamma(t) - \Gamma(s))} \right) \Psi_s \, ds \\ \text{where } \mathcal{L}(\mathcal{X}_{\Gamma(t)}^{\text{gf}}) &= \frac{R}{2d} \text{tr} \left[(\mathbf{A}^T \mathbf{A}) \left(\mathbf{A}^T \mathbf{A} (\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)^{-1} (\mathbf{I}_d - e^{-(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)\Gamma(t)}) - \mathbf{I}_d \right)^2 \right] \\ &\quad + \frac{\tilde{R}}{2d} \text{tr} \left[\left(\mathbf{A} (\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)^{-1} [\mathbf{I}_d - e^{-(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)\Gamma(t)}] \mathbf{A}^T - \mathbf{I}_n \right)^2 \right] \\ &\quad + \frac{\hat{R}}{2d} \text{tr} \left(\mathbf{A}^T \mathbf{A} e^{-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)\Gamma(t)} \right). \end{aligned} \tag{18}$$

For the risk $\mathcal{R}(\cdot) = 1/2\|\cdot - \beta\|_2^2$, we have following expression

$$\begin{aligned}\Omega_t &= \mathcal{R}(\mathcal{X}_{\Gamma(t)}^{\text{gf}}) + \frac{1}{n} \int_0^t \gamma^2(s) \text{tr} \left((\mathbf{A}^T \mathbf{A}) e^{-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)(\Gamma(t) - \Gamma(s))} \right) \Psi_s \, ds \\ \text{where } \mathcal{R}(\mathcal{X}_{\Gamma(t)}^{\text{gf}}) &= \frac{R}{2d} \text{tr} \left[\left(\mathbf{A}^T \mathbf{A} (\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)^{-1} (\mathbf{I}_d - e^{-(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)\Gamma(t)}) - \mathbf{I}_d \right)^2 \right] \\ &\quad + \frac{\tilde{R}}{2d} \text{tr} \left[\left((\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)^{-1} [\mathbf{I}_d - e^{-(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)\Gamma(t)}] \mathbf{A}^T \right)^2 \right] \\ &\quad + \frac{\hat{R}}{2d} \text{tr} \left(e^{-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)\Gamma(t)} \right).\end{aligned}\tag{19}$$

Under the learning rate assumptions in Theorem 3, the limiting GF terms simplify

$$\begin{aligned}\mathcal{L}(\mathcal{X}_{\infty}^{\text{gf}}) &= \frac{R}{2d} \text{tr} \left[(\mathbf{A}^T \mathbf{A}) \left(\mathbf{A}^T \mathbf{A} (\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)^{-1} - \mathbf{I}_d \right)^2 \right] \\ &\quad + \frac{\tilde{R}}{2d} \text{tr} \left[\left(\mathbf{A} (\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)^{-1} \mathbf{A}^T - \mathbf{I}_n \right)^2 \right] \\ \mathcal{R}(\mathcal{X}_{\infty}^{\text{gf}}) &= \frac{R}{2d} \text{tr} \left[\left(\mathbf{A}^T \mathbf{A} (\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)^{-1} - \mathbf{I}_d \right)^2 \right] + \frac{\tilde{R}}{2d} \text{tr} \left[\left((\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)^{-1} \mathbf{A}^T \right)^2 \right].\end{aligned}\tag{20}$$

Excess risk in linear regression. In the standard linear regression setup, we suppose that \mathbf{A} is generated by taking n independent d -dimensional samples from a centered distribution \mathcal{D}_f which we assume to be standardized (mean 0 and expected sample-norm-squared 1). We let the matrix $\Sigma_f \in \mathbb{R}^{d \times d}$ be the feature covariance of \mathcal{D}_f , that is

$$\Sigma_f \stackrel{\text{def}}{=} \mathbb{E}[\mathbf{a}\mathbf{a}^T], \quad \text{where } \mathbf{a} \sim \mathcal{D}_f.\tag{21}$$

Suppose there is a linear (“ground truth” or “signal”) function $\beta : \mathbb{R}^d \rightarrow \mathbb{R}$, which for simplicity we suppose to have $\beta(0) = 0$. In this case, we identify β with a vector using the representation $\mathbf{a} \mapsto \beta^T \mathbf{a}$. We suppose that our data is drawn from a distribution \mathcal{D} on $\mathbb{R}^d \times \mathbb{R}$, with the property that

$$\mathbb{E}[b | \mathbf{a}] = \beta^T \mathbf{a}, \quad \text{where } (\mathbf{a}, b) \sim \mathcal{D},$$

and the data $\mathbf{a} \sim \mathcal{D}_f$.

Hence we suppose that $[\mathbf{A} | \mathbf{b}]$ is a $\mathbb{R}^{n \times d} \times \mathbb{R}^{n \times 1}$ matrix on independent samples from \mathcal{D} . The vector \mathbf{x}_t represents an estimate of β , and the population risk is

$$\mathcal{R}(\mathbf{x}_t) \stackrel{\text{def}}{=} \frac{1}{2} \mathbb{E}[(b - \mathbf{x}_t^T \mathbf{a})^2 | \mathbf{x}_t] \quad \text{where } (\mathbf{a}, b) \sim \mathcal{D},$$

where (\mathbf{a}, b) is a sample independent of \mathbf{x}_t . This can be evaluated in terms of the feature covariance matrix Σ_f and the noise $\eta^2 \stackrel{\text{def}}{=} \mathbb{E}[(b - \beta^T \mathbf{a})^2]$ to give

$$\mathcal{R}(\mathbf{x}_t) = \frac{1}{2} \eta^2 + \frac{1}{2} (\beta - \mathbf{x}_t)^T \Sigma_f (\beta - \mathbf{x}_t).\tag{22}$$

It is important to note that the sequence $\{\mathbf{x}_{\lfloor tn \rfloor}\}_{t \geq 0}$ is generated from the iterates of SGD applied to the ℓ^2 -regularized least-squares problem (1).

In the case that (\mathbf{a}, b) is jointly Gaussian, it follows that we may represent

$$\mathbf{a} = \Sigma_f^{1/2} \mathbf{z}, \quad b = \beta^T \mathbf{a} + \eta w, \quad \text{where } (\mathbf{z}, w) \sim N(0, \mathbf{I}_d \oplus 1).$$

Therefore, it follows that the iterates $\mathbf{x}_{\lfloor nt \rfloor}$ are generated from the SGD algorithm applied to the problem:

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \frac{\delta}{2} \|\mathbf{x}\|_2^2 \quad \text{where } \mathbf{b} = \mathbf{A}\beta + \hat{\eta} \mathbf{w},$$

and the vector \mathbf{w} is iid $N(0, 1)$ random variables, independent of \mathbf{A} . This is also known as the generative model with noise.

Moreover, if \mathcal{D} satisfies Assumption 7 (with $\Sigma = \Sigma_f$) then the population risk $\mathcal{R}(\mathbf{x}_{\lfloor tn \rfloor})$ is well approximated by Ω :

Theorem 7. Suppose (\mathbf{A}, \mathbf{b}) satisfy Assumption 7, $\delta > 0$ and \mathbf{x}_0 is iid centered subgaussian with $\mathbb{E} \|\mathbf{x}_0\|^2 = \widehat{R}$. For some $\epsilon > 0$, for all $T > 0$, and for all $D > 0$ there is a $C > 0$ such that

$$\Pr \left(\sup_{0 \leq t \leq T} \left\| \begin{pmatrix} \mathcal{L}(\mathbf{x}_{[tn]}^{sgd}) \\ \mathcal{R}(\mathbf{x}_{[tn]}^{sgd}) \end{pmatrix} - \begin{pmatrix} \Psi_t \\ \Omega_t \end{pmatrix} \right\|_2 > d^{-\epsilon} \right) \leq Cd^{-D},$$

where Ψ_t solves (10) and Ω_t solves (9) with \mathcal{R} given by (22).

We remark that under Assumption 7 (and in-distribution) that $\eta^2 = \frac{\widehat{R}}{d}$. In the case of out-of-distribution regression (see section below), we have that $\eta^2 \neq \frac{\widehat{R}}{d}$ as the η represents the population noise.

The loss function \mathcal{L} evaluated at GF is the same as in (19) as is the limiting loss Ω_∞ . For the test risk \mathcal{R} in (22) evaluated at GF, we have the following expressions for

$$\begin{aligned} \mathcal{R}(\mathcal{X}_{\Gamma(t)}^{\text{gf}}) &= \frac{R}{2d} \text{tr} \left(\Sigma_f \left(C(t) \mathbf{A}^T \mathbf{A} - \mathbf{I}_d \right)^2 \right) + \frac{\widetilde{R}}{2d} \text{tr} \left(\Sigma_f \mathbf{A}^T \mathbf{A} C^2(t) \right) \\ &\quad + \frac{\widehat{R}}{2d} \text{tr} \left(\Sigma_f \exp \left(-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d) \Gamma(t) \right) \right) + \frac{1}{2} \eta^2 \\ \text{and } \mathcal{R}(\mathcal{X}_\infty^{\text{gf}}) &= \frac{R}{2d} \text{tr} \left(\Sigma_f (\mathbf{A}^T \mathbf{A} (\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)^{-1} - \mathbf{I}_d)^2 \right) \\ &\quad + \frac{\widetilde{R}}{2d} \text{tr} \left(\Sigma_f \mathbf{A}^T \mathbf{A} (\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)^{-2} \right) + \frac{1}{2} \eta^2 \\ \text{where } C(t) &\stackrel{\text{def}}{=} (\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)^{-1} \left(\mathbf{I}_d - \exp \left(-(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d) \Gamma(t) \right) \right). \end{aligned} \quad (23)$$

Using Theorem 3, we conclude that in the case that $\gamma(t) \rightarrow 0$ as $t \rightarrow \infty$, the excess risk of SGD tends to 0. More interestingly, in the interpolation regime, $\mathcal{L}(\mathcal{X}_\infty^{\text{gf}}) = 0$, i.e. the empirical risk tends to 0. In this case, even without taking $\gamma \rightarrow 0$, the excess risk of SGD tends to 0. If on the other hand it does not tend to 0 (i.e., $\gamma(t) \rightarrow \gamma$), we arrive at the formula for excess risk of SGD over the ridge estimator risk:

$$\begin{aligned} \Omega_\infty - \mathcal{R}(\mathcal{X}_\infty^{\text{gf}}) &= \mathcal{L}(\mathcal{X}_\infty^{\text{gf}}) \times \frac{\gamma}{2n} \frac{\text{tr}((\nabla^2 \mathcal{L}) \Sigma_f (\nabla^2 \mathcal{L} + \delta \mathbf{I}_d)^{-1})}{1 - \frac{\gamma}{2n} \text{tr}((\nabla^2 \mathcal{L})^2 (\nabla^2 \mathcal{L} + \delta \mathbf{I}_d)^{-1})} \\ &= \Psi_\infty \times \frac{\gamma}{n} \text{tr} \left(\frac{(\nabla^2 \mathcal{L})}{(\nabla^2 \mathcal{L} + \delta \mathbf{I}_d)} \Sigma_f \right). \end{aligned} \quad (24)$$

We note that the right-hand-side is proportional to Ψ_∞ (c.f. Theorem 3), and hence this excess risk due to SGD will be small if the limiting empirical risk Ψ_∞ is small. This also shows that the regularization term δ interacts with the excess risk due to SGD: if the spectrum of $\nabla^2 \mathcal{R}$ is heavy in that it has slowly decaying eigenvalues, the reduction in excess risk due to the regularization regularizer δ can be large.

(Out-of-distribution) linear regression. As before, we suppose that the data matrix \mathbf{A} is generated by taking n independent d -dimensional samples from a centered distribution \mathcal{D}_f with feature covariance Σ_f (see (21)). We also suppose, as in the previous in-distribution example, that there is a linear (“ground truth” or “signal”) function $\beta : \mathbb{R}^d \rightarrow \mathbb{R}$ which we identify with the vector $\beta \in \mathbb{R}^d$ and for which $\mathbb{E}[b|a] = \beta^T a$ where $(a, b) \sim \mathcal{D}$ and the data $a \sim \mathcal{D}_f$. We will generate our target b from the distribution $(a, b) \sim \mathcal{D}$. We then let \mathbf{x}_t be the iterates generated by SGD applied to the optimization problem

$$\min_{\mathbf{x} \in \mathbb{R}^d} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2 + \frac{\delta}{2} \|\mathbf{x}\|^2, \quad \text{where } (a_i, b_i) \sim \mathcal{D}.$$

The main distinction from the previous example is that we measure our generalization error using a different distribution than \mathcal{D} . Explicitly, there exists another centered distribution $\widehat{\mathcal{D}}_f$ (standardized)

with covariance features matrix $\widehat{\Sigma}_f \in \mathbb{R}^{d \times d}$ from which we generate a vector $\widehat{\mathbf{a}} \sim \widehat{\mathcal{D}}_f$. Moreover, we generate a test point $(\widehat{\mathbf{a}}, \widehat{b})$ from a new distribution $\widehat{\mathcal{D}}$ such that $\mathbb{E}[\widehat{b}|\widehat{\mathbf{a}}] = \beta^T \widehat{\mathbf{a}}$ with the same β as before and the distribution $\widehat{\mathcal{D}}$ has $\widehat{\mathbf{a}}$ -marginal $\widehat{\mathcal{D}}_f$. We measure the population risk, $\mathcal{R} : \mathbb{R}^d \rightarrow \mathbb{R}$ as

$$\mathcal{R}(\mathbf{x}_t) \stackrel{\text{def}}{=} \frac{1}{2} \mathbb{E}[(\widehat{b} - \mathbf{x}_t^T \widehat{\mathbf{a}})^2 | \mathbf{x}_t] = \frac{1}{2} \eta^2 + \frac{1}{2} (\mathbf{x}_t - \beta)^T \widehat{\Sigma}_f (\mathbf{x}_t - \beta) \quad (25)$$

$$\text{where } \eta^2 \stackrel{\text{def}}{=} \mathbb{E}[(\widehat{b} - \beta^T \widehat{\mathbf{a}})^2].$$

In this setting, we can again derive the limiting excess risk, which has a similar formula for $\mathcal{R}(\mathcal{X}_{\Gamma(t)}^{\text{gf}})$ as in (23) by replacing Σ_f with $\widehat{\Sigma}_f$.

Random features. A central example where the quasi-random assumptions hold is the random features setting, which was introduced in [66] for scaling kernel machines. Random features models provide a rich but tractable class of models to gain further insights into the generalization phenomena [51, 44, 3, 2, 77]. These models are particularly of interest because of their connection to neural networks where the number of random features corresponds to model complexity [29, 55, 38] and because of its use as a practical method for data analysis [66, 73].

We suppose that the data matrix \mathbf{X} is generated by taking n independent n_0 -dimensional samples from a centered distribution \mathcal{D}_f with feature covariance

$$\Sigma_f \stackrel{\text{def}}{=} \mathbb{E}[\mathbf{X}_i^T \mathbf{X}_i], \quad \text{where } \mathbf{X}_i \in \mathbb{R}^{1 \times n_0} \text{ and } \mathbf{X}_i \sim \mathcal{D}_f.$$

We suppose for simplicity that \mathbf{X} is a data matrix having dimension $n \times n_0$ whose iid rows are drawn from a multivariate Gaussian with covariance Σ_f and nice covariance structure:

Assumption 8. *The distribution \mathcal{D}_f is multivariate normal and the covariance matrix Σ_f of the random features data satisfies for some $C > 0$*

$$\frac{1}{n_0} \text{tr}(\Sigma_f) = 1 \quad \text{and} \quad \|\Sigma_f\|_{\text{op}} \leq C.$$

This allows \mathbf{X} to be represented equivalently as $\mathbf{X} = \mathbf{Z}\Sigma^{1/2}/\sqrt{n_0}$ for a iid standard Gaussian matrix \mathbf{Z} . We suppose that \mathbf{W} is an $(n_0 \times d)$ iid feature matrix having standard Gaussian entries and independent of \mathbf{Z} so that $\mathbf{Z}\Sigma^{1/2}\mathbf{W}/\sqrt{n_0}$ is a matrix whose rows are standardized.

We let σ be an activation function satisfying:

Assumption 9. *The activation function satisfies for $C_0, C_1 \geq 0$*

$$|\sigma'(x)| \leq C_0 e^{C_1|x|}, \quad \text{for all } x \in \mathbb{R}, \quad \text{and for standard normal } Z, \quad \mathbb{E}\sigma(Z) = 0.$$

We note that from the outset, the growth rate of the derivative of the activation function implies a similar bound on the growth rate of the underlying activation function σ . As before, we suppose the data $[\mathbf{X} | \mathbf{b}]$ is arranged in the matrix $\mathbb{R}^n \times (\mathbb{R}^{n_0} \times \mathbb{R})$ where each row is an independent sample from \mathcal{D} . We now transform the data $\mathbf{X} \in \mathbb{R}^{n \times n_0}$ by putting

$$\mathbf{A} = \sigma(\mathbf{X}\mathbf{W}/\sqrt{n_0}) \in \mathbb{R}^{n \times d},$$

where $\mathbf{W} \in \mathbb{R}^{n_0 \times d}$ is a matrix independent of $[\mathbf{X} | \mathbf{b}]$ of independent standard normals.³ The activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is applied element-wise.

We introduce the following notation

$$\begin{aligned} \Sigma_\sigma(\mathbf{W}) &\stackrel{\text{def}}{=} \mathbb{E}[\sigma(\mathbf{X}_i \mathbf{W}/\sqrt{n_0})^T \sigma(\mathbf{X}_i \mathbf{W}/\sqrt{n_0}) | \mathbf{W}] \\ \text{and } \widehat{\sigma}(\mathbf{W}) &\stackrel{\text{def}}{=} \mathbb{E}[\mathbf{X}_i^T \sigma(\mathbf{X}_i \mathbf{W}/\sqrt{n_0}) | \mathbf{W}]. \end{aligned} \quad (26)$$

The population risk, $\mathcal{R} : \mathbb{R}^d \rightarrow \mathbb{R}$ as a random variable in \mathbf{X} and \mathbf{W} , is

$$\begin{aligned} \mathcal{R}(\mathbf{x}_t) &\stackrel{\text{def}}{=} \mathbb{E}[(b - \mathbf{x}_t^T \sigma(\mathbf{X}_i \mathbf{W}/\sqrt{n_0}))^2 | \mathbf{x}_t, \mathbf{W}] \\ &= \eta^2 + \mathbb{E}[(\mathbf{X}_i \beta - \sigma(\mathbf{X}_i \mathbf{W}/\sqrt{n_0}) \mathbf{x}_t)^2 | \mathbf{x}_t, \mathbf{W}] \\ &= \eta^2 + \beta^T \Sigma_f \beta + \mathbf{x}_t^T \Sigma_\sigma(\mathbf{W}) \mathbf{x}_t - 2\beta^T \widehat{\sigma}(\mathbf{W}) \mathbf{x}_t, \end{aligned} \quad (27)$$

where $(\mathbf{X}_i, b) \sim \mathcal{D}$ and $\mathbb{E}[b | \mathbf{X}_i] = \mathbf{X}_i \beta$.

³In [51], the distribution of the columns are taken as independent uniform vectors on the sphere $\sqrt{d}\mathbb{S}^{d-1}$. The activation function σ is a 1-Lipschitz function from $\mathbb{R} \rightarrow \mathbb{R}$ that is applied entrywise to the underlying matrix.

The ℓ^2 -regularized least-squares problem is now

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \frac{\delta}{2} \|\mathbf{x}\|_2^2 \quad \text{where} \quad \mathbf{b} = \mathbf{X}\boldsymbol{\beta} + \eta\mathbf{w},$$

which is the random features regression. This should be compared to a two-layer neural network model, in which the hidden layer has dimension n_0 . However, the hidden layer weights are simply generated randomly in advance and are left untrained. The optimization is only performed on the final layers' weights (\mathbf{x}).

Theorem 8. *Suppose that n, d, n_0 are proportionally related. Suppose that the data matrix \mathbf{X} satisfies Assumption 8, and the random features \mathbf{W} are iid standard normal. Suppose $\mathbf{b} = \mathbf{X}\boldsymbol{\beta} + \eta\mathbf{w}$ with $\boldsymbol{\beta}, \mathbf{w}$ independent isotropic subgaussian vectors with $\mathbb{E} \|\boldsymbol{\beta}\|_2^2 = 1/n_0$ and $\mathbb{E} \|\mathbf{w}\|_2^2 = 1$ and η bounded independent of n . Suppose the activation function satisfies Assumption 9. Suppose the initialization \mathbf{x}_0 is iid centered subgaussian with $\mathbb{E} \|\mathbf{x}_0\|_2^2 = \widehat{R}$. Then for some $\epsilon > 0$, for all $T > 0$, and for all $D > 0$ there is a $C > 0$ such that*

$$\Pr \left(\sup_{0 \leq t \leq T} \left\| \begin{pmatrix} \mathcal{L}(\mathbf{x}_{[tn]}^{\text{sgd}}) \\ \mathcal{R}(\mathbf{x}_{[tn]}^{\text{sgd}}) \end{pmatrix} - \begin{pmatrix} \Psi_t \\ \Omega_t \end{pmatrix} \right\| > d^{-\epsilon} \right) \leq Cd^{-D},$$

where Ψ_t solves (10) and Ω_t solves (9) with \mathcal{R} given by (27).

Finally, as in (24) we derive the excess risk of SGD ($\gamma(t) \rightarrow \gamma$) over ridge regression:

$$\begin{aligned} \Omega_\infty - \mathcal{R}(\mathcal{X}_\infty^{\text{gf}}) &= \mathcal{L}(\mathcal{X}_\infty^{\text{gf}}) \times \frac{\gamma}{2n} \frac{\text{tr}((\nabla^2 \mathcal{L}) \boldsymbol{\Sigma}_\sigma(\mathbf{W}) (\nabla^2 \mathcal{L} + \delta \mathbf{I}_d)^{-1})}{1 - \frac{\gamma}{2n} \text{tr}((\nabla^2 \mathcal{L})^2 (\nabla^2 \mathcal{L} + \delta \mathbf{I}_d)^{-1})} \\ &= \Psi_\infty \times \frac{\gamma}{2n} \text{tr} \left(\frac{(\nabla^2 \mathcal{L})}{(\nabla^2 \mathcal{L} + \delta \mathbf{I}_d)} (\boldsymbol{\Sigma}_\sigma(\mathbf{W})) \right). \end{aligned} \quad (28)$$

B.1.1 Discussion of Theorem 1 and motivating examples

In this section, we discuss the equivalence of SGD and homogenized SGD under quadratic statistics \mathcal{R} satisfying Assumption 1 and quasi-random assumptions on the data matrix \mathbf{A} , initialization \mathbf{x}_0 , and target vector \mathbf{b} (see Appendix B). As the proof of Theorem 1 is quite mathematically involved and it does not add to the interpretation of the risk trajectories, we relegate this proof to [61, Theorem 1.3].

The proofs of Theorems 6, 7, and 8 follow immediately from Theorem 1 and Theorem 2. In each of the cases, the extra assumptions on the initialization, signal, and data matrix simplify the GF terms in (8) and (9).

C Volterra Equation for Streaming SGD

In this section we introduce constant learning rate ($\gamma(t) \equiv \gamma$) streaming SGD. Let $(\mathbf{a}_k, b_k)_{k=1}^\infty$ be iid samples from a $\mathbb{R}^d \times \mathbb{R}$ -dimensional distribution \mathcal{D} . Streaming SGD using data from \mathcal{D} , which we denote \mathcal{D} -SGD, is the algorithm

$$\mathbf{s}_{k+1} = \mathbf{s}_k - \gamma \mathbf{a}_k (\mathbf{a}_k \cdot \mathbf{s}_k - b_k) \quad \text{for} \quad k \in \{0, 1, 2, \dots\}. \quad (29)$$

This naturally describes one-pass SGD, in which data points are used only once. If \mathcal{R} is the expected risk (i.e. population risk), and if \mathcal{R} is given by $\frac{1}{2} \mathbb{E} (\mathbf{a} \cdot \mathbf{x} - b)^2$ with $(\mathbf{a}, b) \sim \mathcal{D}$, then \mathcal{D} -SGD is directly solving the population risk minimization. This is an idealized situation as one does not have access to infinite data in practice.

Deterministic behavior of streaming risks and comparison to HSGD. \mathcal{D} -SGD can encompass multi-pass SGD by letting $(\mathbf{a}_k, b_k)_{k=1}^n$ be the first n iid samples from \mathcal{D} and considering $\widehat{\mathcal{D}}$ -SGD for

$$\widehat{\mathcal{D}}_n \stackrel{\text{def}}{=} \frac{1}{n} \sum_{k=1}^n \delta_{(\mathbf{a}_k, b_k)}. \quad (30)$$

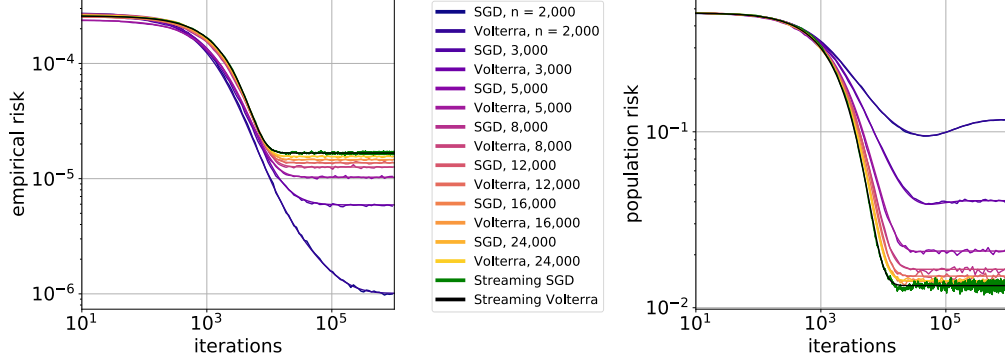


Figure 5: **Single runs of SGD vs. HSGD (Volterra) for a simple Gaussian linear regression problem**, with increasing number of samples n and $d = 2000$. Empirical risk (left) increases monotonically with n to a limit while population risk (right) decreases monotonically in n . Streaming corresponds to $n = \infty$. Covariance of Gaussian samples is \mathbf{I}_d , with target given by $b = a \cdot \beta + \eta Z$ for $\eta = 0.2$ and β a unit vector and $Z \sim N(0, 1)$. For consistency across sample sizes, time is measured in iterations.

This means that a sample from $\widehat{\mathcal{D}}_n$ (conditionally on the dataset) is distributed like (\mathbf{a}_I, b_I) , where I is a uniformly random choice of index.

To enable a comparison of streaming SGD to multi-pass SGD, we suppose for some n the matrix $[\mathbf{A} | \mathbf{b}]$ is $n \times (d + 1)$ matrix whose rows are iid samples from \mathcal{D} . We construct the empirical risk \mathcal{L} from these n samples as in (2). We then define the streaming loss

$$\mathcal{S}(\mathbf{x}) \stackrel{\text{def}}{=} \frac{1}{2} \mathbb{E}[(\mathbf{a} \cdot \mathbf{x} - b)^2] = \frac{1}{2n} \mathbb{E} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2 = \frac{1}{n} \mathbb{E} \mathcal{L}(\mathbf{x}), \quad \text{where } (\mathbf{a}, b) \sim \mathcal{D}.$$

The associated homogenized SGD representation, which we call homogenized \mathcal{D} -SGD, is

$$d\mathbf{Y}_t \stackrel{\text{def}}{=} -\gamma(\nabla \mathcal{S}(\mathbf{Y}_t)) dt + \gamma \sqrt{2\mathcal{S}(\mathbf{Y}_t)(\nabla^2 \mathcal{S})} d\mathbf{B}_t \quad \text{for } t \geq 0. \quad (31)$$

This naturally leads to Volterra dynamics in which $\nabla^2 \mathcal{L}$ is replaced by $\nabla^2 \mathcal{S}$ in Eqs. (8)-(10) whose solution we denote with Ψ_t^s and Ω_t^s (see Appendix D).

We prove a weak equivalence between HSGD and homogenized \mathcal{D} -SGD in the following theorem.

Theorem 9. *Suppose $\mathbf{X}_0 = \mathbf{Y}_0$. After rescaling time,*

$$\lim_{n \rightarrow \infty} \mathbf{X}_{t/n} = \mathbf{Y}_t \quad \text{uniformly on compact sets of time a.s.}$$

Furthermore, GF converges, $\mathcal{X}_{\gamma t/n}^{gf} \rightarrow \mathcal{X}_{\gamma t}^{s-gf}$, and the Volterra equations converge, $\Psi_{t/n} \rightarrow \Psi_t^s$ and $\Omega_{t/n} \rightarrow \Omega_t^s$, uniformly on compact sets of time.

We expect this to hold in much greater generality as suggested by Figures 1 and 5. This is an immediate consequence of the law of large numbers due to which $\frac{1}{n}\mathcal{L}$ converges almost surely to \mathcal{S} ; the applications in which we are typically interested would take time large as a function of n (the numerical results are extremely strong, see Figures 5 and 1), and we leave a deeper mathematical investigation of this point as an open question. We also note that for t very large with n , there is likely another behavior that takes hold. For example, when $t = n$, observe that $\mathbf{X}_{t/n}$ will have used approximately $(1 - 1/e)n$ samples, whereas \mathbf{Y}_n will have used n , and hence we expect a breakdown in the connection. Other works also examined the expected risk of streaming including [9, 86] but not as the limit of the dynamics of multi-pass SGD as we have done.

Features of generalization. While there are connections between streaming and multi-pass SGD, certain behavior is only accessible in the multi-pass setting. For example, early stopping can be a useful ingredient in avoiding overfitting when learning overparameterized models. However, late-time overfitting is only observable with multi-pass SGD and does not occur for streaming SGD (see the nonmonotonicity in population risk in Fig. 5).

Similarly, comparing streaming and multi-pass SGD has been suggested as a method for analyzing generalization. The bootstrap risk was introduced in [54] as an interpretable component in a decomposition of the population risk. It is defined as the excess risk of SGD for ERM \mathcal{L} with n iid samples from \mathcal{D} when compared to \mathcal{D} -SGD, i.e.

$$\underbrace{\Omega_{t/n} - \Omega_t^s}_{\text{bootstrap risk}} = \underbrace{\mathcal{R}(\mathcal{X}_{\gamma t/n}^{\text{gf}}) - \mathcal{R}(\mathcal{X}_{\gamma t}^{\text{s-gf}})}_{\text{model risk}} + \underbrace{\int_0^{t/n} K(t/n, u; \nabla^2 \mathcal{R}) \Psi_u \, du - \int_0^t K^s(t, u; \nabla^2 \mathcal{R}) \Psi_u^s \, du}_{\text{SGD bootstrap risk}}. \quad (32)$$

Thus the Volterra equations can be used to give an exact expression for the bootstrap risk. In particular, we can predict the iteration at which the bootstrap risk becomes large as the streaming and multi-pass risks bifurcate. This allows for quantitative prediction of the stopping time in Claim 1 of [54], a central conjecture of their paper.

On taking time to infinity, we can further evaluate the SGD bootstrap risk.

Theorem 10. *If $\gamma \leq \min\{2(\frac{1}{n} \text{tr}(\mathbf{A}^T \mathbf{A}))^{-1}, 2(\text{tr} \nabla^2 \mathcal{S})^{-1}\}$, then the limiting bootstrap risk is given by*

$$\lim_{t \rightarrow \infty} \Omega_{t/n} - \Omega_t^s = \mathcal{R}(\mathcal{X}_\infty^{\text{gf}}) - \mathcal{R}(\mathcal{X}_\infty^{\text{s-gf}}) + \frac{\gamma \Psi_\infty}{2n} \times \text{tr} \left\{ \frac{(\nabla^2 \mathcal{R})(\nabla^2 \mathcal{L})}{\nabla^2 \mathcal{L} + \delta \mathbf{I}_d} \right\}$$

Here Ψ_∞ is the limiting training losses given by Theorem 3, i.e.

$$\Psi_\infty = \mathcal{L}(\mathcal{X}_\infty^{\text{gf}}) \times \left(1 - \frac{\gamma}{2n} \text{tr} \left\{ \frac{(\nabla^2 \mathcal{L})^2}{\nabla^2 \mathcal{L} + \delta \mathbf{I}_d} \right\} \right)^{-1}.$$

D The Volterra SLD class and Concentration of HSGD

First, we state the Volterra equation for streaming SGD that we referenced in the main text.

Volterra Dynamics (streaming). The following deterministic dynamical system is the high-dimensional equivalent for $\mathcal{L}(\mathbf{Y}_t)$ and $\mathcal{R}(\mathbf{Y}_t)$, respectively,

$$\Psi_t^s = \mathcal{S}(\mathcal{X}_{\gamma t}^{\text{s-gf}}) + \int_0^t K^s(t, u; \nabla^2 \mathcal{S}) \Psi_u^s \, du, \quad \text{for } t \geq 0 \quad (\text{Empirical risk}) \quad (33)$$

$$\Omega_t^s = \mathcal{R}(\mathcal{X}_{\gamma t}^{\text{s-gf}}) + \int_0^t K^s(t, u; \nabla^2 \mathcal{R}) \Psi_u^s \, du \quad \text{for } t \geq 0 \quad (\text{Population risk}) \quad (34)$$

where the kernel K , for any $d \times d$ matrix \mathbf{P} , is

$$K^s(t, u; \mathbf{P}) = \gamma^2 \text{tr} \left((\nabla^2 \mathcal{S}) \mathbf{P} \exp(-2\gamma(\nabla^2 \mathcal{S})(t - u)) \right) \quad (35)$$

and GF for streaming, $\mathcal{X}_t^{\text{s-gf}}$, is the solution of

$$d\mathcal{X}_t^{\text{s-gf}} \stackrel{\text{def}}{=} -\nabla \mathcal{S}(\mathcal{X}_t^{\text{s-gf}}) dt \quad \text{and} \quad \mathcal{X}_0^{\text{s-gf}} = \mathbf{Y}_0.$$

We will enlarge the class of SLDs that we consider to what we will call as the *Volterra SLD class* defined as

$$d\mathbf{X}_t \stackrel{\text{def}}{=} -\gamma(t) \nabla f(\mathbf{X}_t) dt + \gamma(t) \sqrt{\mathcal{L}(\mathbf{X}_t) \mathcal{M}_t + \mathcal{A}_t} d\mathbf{B}_t, \quad (36)$$

where \mathcal{M}_t and \mathcal{A}_t are two deterministic positive definite functions which we assume to be normalized to satisfy:

Assumption 10. *The covariance processes \mathcal{M} and \mathcal{A} satisfy for some absolute constants $c > 0$ and $\epsilon > 0$*

$$\sup_{t \geq 0} (\text{tr} \mathcal{M}_t + \text{tr} \mathcal{A}_t) \leq c < \infty \quad \text{and} \quad \sup_{t \geq 0} (\|\mathcal{M}_t\|_{op} + \|\mathcal{A}_t\|_{op}) < d^{-\epsilon}.$$

The \mathcal{M}_t represents noise in the data either because the data is randomly sampled or the data is transformed by multiplicative transformation. In contrast \mathcal{A}_t represents an additive noise at each step. It is any noise which does not multiple the state \mathbf{x} , for example, label noise.

Volterra SLD class. The Volterra SLD class is so-named because the expected loss satisfies a Volterra type integral equation. Define for $t \geq s \geq 0$ and positive semidefinite \mathbf{P} ,

$$\Gamma(t) = \int_0^t \gamma(s) ds, \text{ and } \left\{ \begin{array}{l} K(t, s; \mathbf{P}) = \gamma^2(s) \text{tr} \left(\mathcal{M}_s \mathbf{P} \exp(-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)(\Gamma(t) - \Gamma(s))) \right) \\ A(t, s; \mathbf{P}) = \gamma^2(s) \text{tr} \left(\mathcal{A}_s \mathbf{P} \exp(-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)(\Gamma(t) - \Gamma(s))) \right) \end{array} \right\}. \quad (37)$$

We shall suppose throughout that $\mathcal{X}_t^{\text{gf}}$ is the canonical GF

$$d\mathcal{X}_t^{\text{gf}} = -\nabla f(\mathcal{X}_t^{\text{gf}}) \quad \text{and} \quad \mathcal{X}_0^{\text{gf}} = \mathbf{X}_0.$$

The loss $\mathcal{L}(\mathbf{X}_t)$ concentrates around the solution Ψ_t of the Volterra integral equation (see Theorem 11 for a precise formulation):

$$\Psi_t = \mathcal{L}(\mathcal{X}_{\Gamma(t)}^{\text{gf}}) + \int_0^t A(t, s; \mathbf{A}^T \mathbf{A}) ds + \int_0^t K(t, s; \mathbf{A}^T \mathbf{A}) \Psi_s ds. \quad (38)$$

We give a formal proof of the concentration result in Section D.1. For other quadratics \mathcal{R} , the loss $\mathcal{R}(\mathbf{X}_t)$ concentrates around

$$\Omega_t = \mathcal{R}(\mathcal{X}_{\Gamma(t)}^{\text{gf}}) + \int_0^t A(t, s; \nabla^2 \mathcal{R}) ds + \int_0^t K(t, s; \nabla^2 \mathcal{R}) \Psi_s ds. \quad (39)$$

D.1 Volterra Concentration

In this section, we prove the concentration result. In this section, we prove that homogenized SGD concentrates around its mean provided that the expected loss is in the Volterra SLD class. The result, in this section, Theorem 11, is more general than Theorem 2 which follows by setting $\mathcal{M}_t \equiv \frac{1}{n} \nabla^2 \mathcal{L}$ and $\mathcal{A}_t \equiv 0$.

Theorem 11. *Under Assumption 10, the loss $\mathcal{L}(\mathbf{X}_t)$ concentrates around Ψ_t the solution of the Volterra equation*

$$\Psi_t = \mathcal{L}(\mathcal{X}_{\Gamma(t)}^{\text{gf}}) + \int_0^t A(t, s; \mathbf{A}^T \mathbf{A}) ds + \int_0^t K(t, s; \mathbf{A}^T \mathbf{A}) \Psi_s ds,$$

in that for any $T, D > 0$ there is a $C(T, D, \|\mathbf{A}\|_{\text{op}}, \|\mathbf{b}\|, \epsilon) > 0$ sufficiently large that

$$\Pr \left[\sup_{0 \leq t \leq T} |\mathcal{L}(\mathbf{X}_t) - \Psi_t| > Cd^{-\epsilon/2} \right] \leq Cd^{-D}.$$

Furthermore, for another quadratic $\mathcal{R}(\mathbf{x}) = 1/2 \mathbf{x}^T \mathbf{T} \mathbf{x} + \mathbf{u}^T \mathbf{x} + c$ with \mathbf{T} symmetric matrix having $\|\nabla^2 \mathcal{R}\|_{\text{op}} \leq C$, $\|\nabla \mathcal{R}\|_2 \leq C$, $\|\nabla \mathcal{R}(0)\| \leq 1$ are independent of the Brownian motion,

$$\Pr \left[\sup_{0 \leq t \leq T} \left| -\mathcal{R}(\mathbf{X}_t) + \mathcal{R}(\mathcal{X}_{\Gamma(t)}^{\text{gf}}) + \int_0^t A(t, s; \nabla^2 \mathcal{R}) ds + \int_0^t K(t, s; \nabla^2 \mathcal{R}) \Psi_s ds \right| > Cd^{-\epsilon/2} \right] \leq Cd^{-D}.$$

Proof. Step 1. Volterra equation for the expected loss.

Define $\mathbf{Q}_t = \exp((\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)\Gamma(t))$ and apply Itô's rule to $\mathbf{Q}_t \mathbf{X}_t$, derive

$$d(\mathbf{Q}_t \mathbf{X}_t) = \gamma(t) \mathbf{Q}_t \mathbf{A}^T \mathbf{b} dt + \gamma(t) \mathbf{Q}_t \sqrt{\mathcal{L}(\mathbf{X}_t) \mathcal{M}_t + \mathcal{A}_t} d\mathbf{B}_t.$$

Hence

$$\mathbf{Q}_t \mathbf{X}_t = \mathbf{Q}_0 \mathbf{X}_0 + \int_0^t \gamma(s) \mathbf{Q}_s \mathbf{A}^T \mathbf{b} ds + \int_0^t \gamma(s) \mathbf{Q}_s \sqrt{\mathcal{L}(\mathbf{X}_s) \mathcal{M}_s + \mathcal{A}_s} d\mathbf{B}_s.$$

Note that on setting $\mathcal{M}_s = \mathcal{A}_s = 0$, this gives GF $\mathcal{X}_{\Gamma(t)}^{\text{gf}}$, and hence we have representation

$$\mathbf{X}_t = \mathcal{X}_{\Gamma(t)}^{\text{gf}} + \mathbf{Q}_t^{-1} \int_0^t \gamma(s) \mathbf{Q}_s \sqrt{\mathcal{L}(\mathbf{X}_s) \mathcal{M}_s + \mathcal{A}_s} d\mathbf{B}_s.$$

Expanding the quadratic,

$$\begin{aligned} \mathcal{L}(\mathbf{X}_t) &= \mathcal{L}(\mathbf{x}_{\Gamma(t)}^{\text{gf}}) + \nabla \mathcal{L}(\mathbf{x}_{\Gamma(t)}^{\text{gf}})^T \mathbf{Q}_t^{-1} \int_0^t \gamma(s) \mathbf{Q}_s \sqrt{\mathcal{L}(\mathbf{X}_s) \mathcal{M}_s + \mathcal{A}_s} d\mathbf{B}_s \\ &\quad + \frac{1}{2} \left\| \mathbf{A} \mathbf{Q}_t^{-1} \int_0^t \gamma(s) \mathbf{Q}_s \sqrt{\mathcal{L}(\mathbf{X}_s) \mathcal{M}_s + \mathcal{A}_s} d\mathbf{B}_s \right\|^2. \end{aligned} \quad (40)$$

It follows that with \mathcal{F}_t the sigma-algebra generated by $(\mathbf{X}_0, \mathcal{L}, (\mathbf{B}_s : 0 \leq s \leq t))$ if we compute the \mathcal{F}_0 -conditional expectation, the Brownian integral vanishes, and we are left with two contributions from the second norm-squared process

$$\begin{aligned} \mathbb{E}[\mathcal{L}(\mathbf{X}_t) \mid \mathcal{F}_0] &= \mathcal{L}(\mathbf{x}_{\Gamma(t)}^{\text{gf}}) + \frac{1}{2} \int_0^t \gamma^2(s) \text{tr} \left(\mathbf{A}^T \mathbf{A} \mathbf{Q}_t^{-2} \mathbf{Q}_s^2 \mathcal{A}_s \right) ds \\ &\quad + \frac{1}{2} \int_0^t \gamma^2(s) \text{tr} \left(\mathbf{A}^T \mathbf{A} \mathbf{Q}_t^{-2} \mathbf{Q}_s^2 \mathcal{M}_s \right) \mathbb{E}[\mathcal{L}(\mathbf{X}_s) \mid \mathcal{F}_0] ds. \end{aligned} \quad (41)$$

This is the claimed Volterra equation (see e.g., $\Psi_t \stackrel{\text{def}}{=} \mathbb{E}[\mathcal{L}(\mathbf{X}_t) \mid \mathcal{F}_0]$ in (8); here $\mathcal{A}_t = 0$.)

Step 2. High probability boundedness of \mathcal{L} . We observe before beginning that many of the quantities that appear in the expressions above are bounded. The GF $\mathbf{x}_{\Gamma(t)}^{\text{gf}}$ satisfies a uniform bound, solely in terms of its initial conditions and, in particular, the boundedness of $\mathcal{L}(\mathbf{x}_{\Gamma(t)}^{\text{gf}})$ satisfies $\mathcal{L}(\mathbf{x}_{\Gamma(t)}^{\text{gf}}) \leq \mathcal{L}(\mathbf{X}_0)$. The matrix $\mathbf{Q}_t^{-1} \mathbf{Q}_s$ is uniformly bounded in norm by 1 for all $t \geq s$. We have also assumed that $\|\mathbf{A}\|_{\text{op}}$ and $\|\mathbf{b}\|_2$ are bounded. By applying Itô's formula to the norm $u_t = \frac{1}{2} \|\mathbf{X}_t\|_2^2$, we have from (36) that

$$du_t = -\gamma(t) \mathbf{X}_t^T (\mathbf{A}^T (\mathbf{A} \mathbf{X}_t - \mathbf{b})) dt + \gamma(t) \mathbf{X}_t^T \sqrt{\mathcal{L}(\mathbf{X}_t) \mathcal{M}_t + \mathcal{A}_t} d\mathbf{B}_t + \frac{\gamma^2(t)}{2} \text{tr}(\mathcal{L}(\mathbf{X}_t) \mathcal{M}_t + \mathcal{A}_t) dt$$

From the norm boundedness of \mathbf{A} and \mathbf{b} , we can bound $\mathcal{L}(\mathbf{X}_t) \leq 2\|\mathbf{A}\|_{\text{op}}^2 u_t + 2\|\mathbf{b}\|_2^2 \leq C(u_t + 1)$. Likewise, increasing C as need be, using the boundedness of $\gamma(t)$, $\text{tr} \mathcal{M}_t$ and $\text{tr} \mathcal{A}_t$, we conclude

$$d\langle u_t \rangle = \gamma^2(t) \text{tr}(\mathbf{X}_t \mathbf{X}_t^T (\mathcal{L}(\mathbf{X}_t) \mathcal{M}_t + \mathcal{A}_t)) \leq C(u_t + 1)^2.$$

It follows that $z_t \stackrel{\text{def}}{=} \log(1 + u_t) - Ct$ is supermartingale with $\langle z_t \rangle \leq C$ for some sufficiently large C and all $t \leq T$. Hence with probability at least $1 - e^{-2C(T)(\log d)^{3/2}}$,

$$z_t \leq (\log d)^{3/4}$$

for all $t \leq T$. On this same event it follows for a sufficiently large constant $C > 0$

$$f(t) = \mathcal{L}(\mathbf{X}_t) + \delta u_t \quad \text{and} \quad \mathcal{L}(\mathbf{X}_t) \leq C(u_t + 1) \leq C^2 e^{Ct + (\log d)^{3/4}}$$

for all $t \leq T$.

Step 3. Concentration of the loss. We may now control the difference of the loss from its expectation.

Specifically, in comparing (40) and (41), we may express the difference $\Delta_t \stackrel{\text{def}}{=} \mathcal{L}(\mathbf{X}_t) - \mathbb{E}[\mathcal{L}(\mathbf{X}_t) \mid \mathcal{F}_0]$ as

$$\begin{aligned} \Delta_t &= \mathbf{M}_t^{(1)} + \mathbf{M}_t^{(2)} + \frac{1}{2} \int_0^t \gamma^2(s) \text{tr} \left(\mathbf{A}^T \mathbf{A} \mathbf{Q}_t^{-2} \mathbf{Q}_s^2 \mathcal{M}_s \right) \Delta_s ds, \quad \text{where} \\ \mathbf{M}_t^{(1)} &= \nabla \mathcal{L}(\mathbf{x}_{\Gamma(t)}^{\text{gf}})^T \mathbf{Q}_t^{-1} \int_0^t \gamma(s) \mathbf{Q}_s \sqrt{\mathcal{L}(\mathbf{X}_s) \mathcal{M}_s + \mathcal{A}_s} d\mathbf{B}_s, \quad \text{and} \\ \mathbf{M}_t^{(2)} &= \frac{1}{2} \left\| \mathbf{A} \mathbf{Q}_t^{-1} \int_0^t \gamma(s) \mathbf{Q}_s \sqrt{\mathcal{L}(\mathbf{X}_s) \mathcal{M}_s + \mathcal{A}_s} d\mathbf{B}_s \right\|^2 \\ &\quad - \frac{1}{2} \int_0^t \gamma^2(s) \text{tr} \left(\mathbf{A}^T \mathbf{A} \mathbf{Q}_t^{-2} \mathbf{Q}_s^2 (\mathcal{L}(\mathbf{X}_s) \mathcal{M}_s + \mathcal{A}_s) \right) ds. \end{aligned} \quad (42)$$

We claim that both processes $\mathbf{M}_t^{(1)}$ and $\mathbf{M}_t^{(2)}$ are small, whose proof we defer. Specifically, with probability $1 - C(T)e^{-(\log d)^{3/2}}$ we have

$$\max_{0 \leq t \leq T} \{ |\mathbf{M}_t^{(1)}| + |\mathbf{M}_t^{(2)}| \} \leq d^{-3\epsilon/4}.$$

From the uniform boundedness in norm of $\mathbf{A}^T \mathbf{A}$, we then conclude from (42) for all $t \leq T$

$$|\Delta_t| \leq d^{-3\epsilon/4} + \int_0^t \|\mathbf{A}^T \mathbf{A}\|_{\text{op}} |\Delta_s| \, ds.$$

Using Gronwall's inequality,

$$|\Delta_t| \leq \|\mathbf{A}^T \mathbf{A}\|_{\text{op}}^{-1} (e^{\|\mathbf{A}^T \mathbf{A}\|_{\text{op}} t} - 1) d^{-3\epsilon/4}.$$

Thus we conclude by increasing the constants in the claimed bound that the desired inequality holds.

Step 4 (Deferred). Concentration of the martingales. We introduce two martingales, for each fixed $t \in [0, T]$,

$$\begin{aligned} M_u^{(1,t)} &= \nabla \mathcal{L}(\mathbf{X}_{\Gamma(t)}^{\text{gf}})^T \mathbf{Q}_t^{-1} \int_0^u \gamma(s) \mathbf{Q}_s \sqrt{\mathcal{L}(\mathbf{X}_s) \mathcal{M}_s + \mathcal{A}_s} \, d\mathbf{B}_s. \\ M_u^{(2,t)} &= \frac{1}{2} \left\| \mathbf{A} \mathbf{Q}_t^{-1} \int_0^u \gamma(s) \mathbf{Q}_s \sqrt{\mathcal{L}(\mathbf{X}_s) \mathcal{M}_s + \mathcal{A}_s} \, d\mathbf{B}_s \right\|^2 \\ &\quad - \frac{1}{2} \int_0^u \gamma^2(s) \text{tr} \left(\mathbf{A}^T \mathbf{A} \mathbf{Q}_t^{-2} \mathbf{Q}_s^2 (\mathcal{L}(\mathbf{X}_s) \mathcal{M}_s + \mathcal{A}_s) \right) \, ds. \end{aligned}$$

We first show that if we fix any $t \leq T$, then for all d sufficiently large with respect to T and with probability at least $1 - 2e^{-(\log d)^{3/2}}$,

$$\max_{0 \leq u \leq t} \{|M_u^{(1,t)}| + |M_u^{(2,t)}|\} \leq d^{-7\epsilon/8}.$$

We will then need to use a meshing argument to complete the argument. We show the details for the first. Those for the second are similar.

We simply need to bound the quadratic variation of each. Note

$$\langle M_u^{(1,t)} \rangle = \int_0^u \gamma^2(s) \text{tr}(\mathbf{Q}_t^{-1} \mathbf{Q}_s \nabla \mathcal{L}(\mathbf{X}_{\Gamma(t)}^{\text{gf}}) \nabla \mathcal{L}(\mathbf{X}_{\Gamma(t)}^{\text{gf}})^T \mathbf{Q}_t^{-1} \mathbf{Q}_s (\mathcal{L}(\mathbf{X}_s) \mathcal{M}_s + \mathcal{A}_s)) \, du.$$

Here we use the norm boundedness of $\mathcal{M}_t + \mathcal{A}_t$, by $2d^{-\epsilon}$. We further bound the other terms in norm to produce

$$\langle M_u^{(1,t)} \rangle \leq 2d^{-\epsilon} \|\mathbf{A}^T \mathbf{A}\|_{\text{op}}^2 (C^2 e^{Cu + (\log d)^{3/4}}) \mathcal{L}(\mathbf{X}_{\Gamma(t)}^{\text{gf}}).$$

We note that $\mathcal{L}(\mathbf{X}_{\Gamma(t)}^{\text{gf}}) \leq \mathcal{L}(\mathbf{X}_0^{\text{gf}})$. Hence with probability at least $1 - e^{-(\log d)^{3/2}}$ (for all d sufficiently large with respect to $T, \|\mathbf{A}\|_{\text{op}}, \|\mathbf{b}\|_2, \epsilon$),

$$\max_{0 \leq u \leq t} |M_u^{(1,t)}| \leq d^{-7\epsilon/8}/2.$$

Step 5 (Deferred). Mesh argument. Finally, we use a union bound to gain the control from Step 4 over a mesh of $[0, T]$ of spacing d^{-100} . From the union bound, we therefore have for all these mesh points $\{t_k\}$

$$\max_k \max_{0 \leq u \leq t_k} \{|M_u^{(1,t_k)}| + |M_u^{(2,t_k)}|\} \leq d^{-7\epsilon/8},$$

and this holds with probability $1 - 2Td^{100} e^{-(\log d)^{3/2}}$. For $t \in [t_k, t_{k+1}]$, we just use that

$$\|\nabla \mathcal{L}(\mathbf{X}_{\Gamma(t)}^{\text{gf}})^T \mathbf{Q}_t^{-1} - \nabla \mathcal{L}(\mathbf{X}_{\Gamma(t_{k+1})}^{\text{gf}})^T \mathbf{Q}_{t_{k+1}}^{-1}\| \leq C(T, \mathbf{A}, \mathbf{b}) d^{-100},$$

and thus on the event that $\mathcal{L}(\mathbf{X}_s)$ is bounded, we have for $t \in [t_k, t_{k+1}]$

$$|M_t^{(1)} - M_t^{(1,t_{k+1})}| \leq C(T, \mathbf{A}, \mathbf{b}) d^{-50}$$

for all d sufficiently large with respect to $T, \|\mathbf{A}\|_{\text{op}}$, and $\|\mathbf{b}\|_2$.

Step 6. Other quadratics. Hence, if we take Ψ_t as a solution to the Volterra equation

$$\Psi_t = \mathcal{L}(\mathbf{X}_{\Gamma(t)}^{\text{gf}}) + \frac{1}{2} \int_0^t \gamma^2(s) \text{tr} \left(\mathbf{A}^T \mathbf{A} \mathbf{Q}_t^{-2} \mathbf{Q}_s^2 \mathcal{A}_s \right) \, ds + \frac{1}{2} \int_0^t \gamma^2(s) \text{tr} \left(\mathbf{A}^T \mathbf{A} \mathbf{Q}_t^{-2} \mathbf{Q}_s^2 \mathcal{M}_s \right) \Psi_s \, ds,$$

then we have a high-quality approximation for the loss $\mathcal{L}(\mathbf{X}_t)$, and moreover, applying Itô's equation, we may always represent another quadratic $\mathcal{R} : \mathbb{R}^d \rightarrow \mathbb{R}$ of the SLD by (analogously to (40))

$$\begin{aligned} \mathcal{R}(\mathbf{X}_t) &= \mathcal{R}(\mathbf{X}_{\Gamma(t)}^{\text{gf}}) + \nabla \mathcal{R}(\mathbf{X}_{\Gamma(t)}^{\text{gf}})^T \mathbf{Q}_t^{-1} \int_0^t \gamma(s) \mathbf{Q}_s \sqrt{\mathcal{L}(\mathbf{X}_s) \mathcal{M}_s + \mathcal{A}_s} d\mathbf{B}_s \\ &+ \frac{1}{2} \left(\mathbf{Q}_t^{-1} \int_0^t \gamma(s) \mathbf{Q}_s \sqrt{\mathcal{L}(\mathbf{X}_s) \mathcal{M}_s + \mathcal{A}_s} d\mathbf{B}_s \right)^T (\nabla^2 \mathcal{R}) \mathbf{Q}_t^{-1} \int_0^t \gamma(s) \mathbf{Q}_s \sqrt{\mathcal{L}(\mathbf{X}_s) \mathcal{M}_s + \mathcal{A}_s} d\mathbf{B}_s \end{aligned}$$

By comparing this to the same expression, where we replace the losses $\mathcal{L}(\mathbf{X}_s)$ by Ψ_s and compute expectations over the Brownian terms, we arrive at (compare (41))

$$\begin{aligned} \mathcal{R}(\mathbf{X}_t) &= M_t^{(3)} + \mathcal{R}(\mathbf{X}_{\Gamma(t)}^{\text{gf}}) + \frac{1}{2} \int_0^t \gamma^2(s) \text{tr} \left((\nabla^2 \mathcal{R}) \mathbf{Q}_t^{-2} \mathbf{Q}_s^2 \mathcal{A}_s \right) ds \\ &+ \frac{1}{2} \int_0^t \gamma^2(s) \text{tr} \left((\nabla^2 \mathcal{R}) \mathbf{Q}_t^{-2} \mathbf{Q}_s^2 \mathcal{M}_s \right) \Psi_s ds. \end{aligned}$$

Provided the Hessian $(\nabla^2 \mathcal{R})$ and gradient $\nabla \mathcal{R}(0)$ are bounded independently of d uniformly on T , the concentration of $M_t^{(3)}$ now follows exactly as in Steps 4 and 5. \square

E Limiting values of the excess risk

In this section, we prove the limiting excess risk values, Theorem 3. We will, in fact, prove a more general version of Theorem 3 which holds for a wider class, so called the Volterra SLD class, as discussed in (37). Theorems 12 (constant learning rate) and 13 (time dependent learning rate) immediately imply Theorem 3 by setting $\mathcal{M}_t \equiv \frac{1}{n} \nabla^2 \mathcal{L}$ and $\mathcal{A}_t \equiv 0$. By using the Volterra SLD class, we also recover the result for streaming, Theorem 10.

Excess risk in the constant case. Under the stronger assumptions of constant learning rate, and constant variance profile, this can be further simplified. That is, suppose

Assumption 11. *Suppose that the covariance processes \mathcal{M} and \mathcal{A} are constant and satisfy for some absolute constants $c > 0$ and $\epsilon > 0$*

$$(\gamma(t), \mathcal{M}_t, \mathcal{A}_t) \equiv (\gamma, \mathcal{M}, \mathcal{A}), \quad \text{where} \quad (\text{tr } \mathcal{M} + \text{tr } \mathcal{A}) \leq c < \infty \quad \text{and} \quad (\|\mathcal{M}\|_{op} + \|\mathcal{A}\|_{op}) < d^{-\epsilon}.$$

Under this assumption the kernels in the Volterra equation simplify to be:

$$\begin{aligned} K(t, s; \mathbf{P}) &= K(t - s; \mathbf{P}) = \gamma^2 \text{tr} \left(\mathcal{M} \mathbf{P} \exp(-2\gamma(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)(t - s)) \right), \\ A(t, s; \mathbf{P}) &= A(t - s; \mathbf{P}) = \gamma^2 \text{tr} \left(\mathcal{A} \mathbf{P} \exp(-2\gamma(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)(t - s)) \right) \end{aligned} \quad (43)$$

The theory of convolution-type Volterra equations is substantially simpler than those of non-convolution type. In particular, we can completely rates of convergence and the limiting loss, as well as convergence guarantees (note that if the training loss of the underlying GF does not tend to 0 and or $\mathcal{A} \neq 0$, then the loss does not tend to 0, and so this is neighborhood convergence).

Theorem 12 (Limit risk values, constant learning rate). *Suppose the learning rate is constant, $\gamma(t) \equiv \gamma$. Under Assumption 11, the Volterra SLD is (neighborhood) convergent if and only if*

$$\mathcal{I}(\gamma) \stackrel{\text{def}}{=} \int_0^\infty K(t; \mathbf{A}^T \mathbf{A}) dt = \frac{\gamma}{2} \text{tr} \left(\mathcal{M} (\mathbf{A}^T \mathbf{A}) (\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)^{-1} \right) < 1. \quad (44)$$

In the case that $\mathcal{I}(\gamma) < 1$, Ψ_t converges as $t \rightarrow \infty$ to

$$\Psi_\infty \stackrel{\text{def}}{=} (1 - \mathcal{I})^{-1} \left(\mathcal{L}(\mathbf{X}_\infty^{\text{gf}}) + \frac{\gamma}{2} \text{tr}(\mathcal{A} (\mathbf{A}^T \mathbf{A}) (\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)^{-1}) \right). \quad (45)$$

Likewise, the population risk Ω_t converges as $t \rightarrow \infty$ to

$$\Omega_\infty \stackrel{\text{def}}{=} \mathcal{R}(\mathbf{X}_\infty^{\text{gf}}) + \frac{\gamma}{2} \text{tr}((\mathcal{A} + \mathcal{M} \Psi_\infty) (\nabla^2 \mathcal{R}) (\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)^{-1}). \quad (46)$$

Proof of Theorem 12. follows immediately from limiting values of renewal equations (see [22] and [5]). \square

By setting $\mathcal{M} = \frac{1}{n} \nabla^2 \mathcal{L}$ and $\mathcal{A} \equiv 0$, we recover the multi-pass SGD setting discussed in the main portion of this paper. When the learning rate satisfies $\gamma(t) \equiv \gamma \in (0, 2(\frac{1}{n} \text{tr} \{ \frac{(\mathbf{A}^T \mathbf{A})^2}{\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d} \})^{-1})$, it follows that $\mathcal{I}(\gamma) < 1$. Consequently, Theorem 12 proves Theorem 3 when the learning rate is constant.

Excess risk when learning rate is time dependent. In the case that the learning rate is time dependent, we prove the following result for the limiting dynamics under the expanded Volterra SLD class. Here we will still assume that the covariance processes \mathcal{M} and \mathcal{A} are constant.

Assumption 12. *Suppose that the covariance processes \mathcal{M} and \mathcal{A} are constant and satisfy for some absolute constants $c > 0$ and $\epsilon > 0$*

$$(\mathcal{M}_t, \mathcal{A}_t) \equiv (\mathcal{M}, \mathcal{A}), \quad \text{where} \quad (\text{tr } \mathcal{M} + \text{tr } \mathcal{A}) \leq c < \infty \quad \text{and} \quad (\|\mathcal{M}\|_{op} + \|\mathcal{A}\|_{op}) < d^{-\epsilon}.$$

The time-dependent learning rate excess risk is given below.

Theorem 13 (Time infinity risk values for SLD class). *Suppose Assumption 12 holds for the Volterra SLD class and the integrated learning rate satisfies $\Gamma(t) \rightarrow \infty$ and $\gamma(t) \rightarrow \gamma \in [0, \infty)$. Let the limiting learning rate value γ be chosen such that the kernel norm is less than 1, that is,*

$$\mathcal{I}(\gamma) \stackrel{\text{def}}{=} \int_0^\infty K(t, s; \mathbf{A}^T \mathbf{A}) dt = \frac{\gamma}{2} \text{tr} \left(\mathcal{M}(\mathbf{A}^T \mathbf{A})(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)^{-1} \right) < 1. \quad (47)$$

Then with Ψ_∞ given by the limiting empirical risk,

$$\Psi_\infty = \left(1 - \frac{\gamma}{2} \text{tr} \left\{ \frac{\mathcal{M} \mathbf{A}^T \mathbf{A}}{\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d} \right\} \right)^{-1} \times \left(\mathcal{L}(\mathcal{X}_\infty^{\text{gf}}) + \frac{\gamma}{2} \text{tr} \left\{ \frac{\mathcal{A} \mathbf{A}^T \mathbf{A}}{\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d} \right\} \right), \quad (48)$$

the excess risk converges to

$$\Omega_t - \mathcal{R}(\mathcal{X}_{\Gamma(t)}^{\text{gf}}) \rightarrow \frac{\gamma}{2} \times \text{tr} \left\{ \frac{(\mathcal{A} + \mathcal{M} \Psi_\infty)(\nabla^2 \mathcal{R})}{\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d} \right\}.$$

Proof. First suppose that the limiting loss value of Ψ_t , defined in (38), is bounded and it exists at infinity. We show under this condition on Ψ_t that the limiting risk value holds for Ω_∞ , defined in (39). A simple computation with a change of variables gives

$$\begin{aligned} & \lim_{t \rightarrow \infty} \Omega_t - \mathcal{R}(\mathcal{X}_{\Gamma(t)}^{\text{gf}}) \\ &= \lim_{t \rightarrow \infty} \int_0^t \gamma^2(s) \text{tr} \left(\nabla^2 \mathcal{R} \mathcal{A} \exp(-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)(\Gamma(t) - \Gamma(s))) \right) ds \\ & \quad + \lim_{t \rightarrow \infty} \int_0^t \gamma^2(s) \text{tr} \left((\nabla^2 \mathcal{R}) \mathcal{M} \exp(-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)(\Gamma(t) - \Gamma(s))) \right) \Psi_s ds \\ &= \lim_{t \rightarrow \infty} \int_0^{\Gamma(t)} \gamma(s) \text{tr} \left(\nabla^2 \mathcal{R} \mathcal{A} \exp(-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)(\Gamma(t) - s)) \right) ds \\ & \quad + \lim_{t \rightarrow \infty} \int_0^{\Gamma(t)} \gamma(s) \text{tr} \left((\nabla^2 \mathcal{R}) \mathcal{M} \exp(-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)(\Gamma(t) - s)) \right) \Psi_{\Gamma^{-1}(s)} ds \\ &= \lim_{t \rightarrow \infty} \int_0^{\Gamma(t)} \gamma(\Gamma(t) - v) \text{tr} \left(\nabla^2 \mathcal{R} \mathcal{A} \exp(-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)v) \right) ds \\ & \quad + \lim_{t \rightarrow \infty} \int_0^{\Gamma(t)} \gamma(\Gamma(t) - v) \text{tr} \left((\nabla^2 \mathcal{R}) \mathcal{M} \exp(-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)v) \right) \Psi_{\Gamma^{-1}(\Gamma(t)-v)} dv. \end{aligned} \quad (49)$$

Dominated convergence theorem allows us to interchange the integral and limit as Ψ_t and $\gamma(t)$ are bounded. We pull out the limiting values of $\lim_{t \rightarrow \infty} \gamma(t) = \gamma$ and Ψ_∞ . By integrating, we deduce

$$\begin{aligned}
& \lim_{t \rightarrow \infty} \Omega_t - \mathcal{R}(\mathcal{X}_{\Gamma(t)}^{\text{gf}}) \\
&= \lim_{t \rightarrow \infty} \int_0^{\Gamma(t)} \gamma(\Gamma(t) - v) \operatorname{tr} \left((\nabla^2 \mathcal{R}) \mathcal{A} \exp(-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)v) \right) ds \\
&\quad + \lim_{t \rightarrow \infty} \int_0^{\Gamma(t)} \gamma(\Gamma(t) - v) \operatorname{tr} \left((\nabla^2 \mathcal{R}) \mathcal{M} \exp(-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)v) \right) \Psi_{\Gamma^{-1}(\Gamma(t)-v)} dv \\
&= \gamma \int_0^\infty \operatorname{tr} \left((\nabla^2 \mathcal{R}) \mathcal{A} \exp(-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)v) \right) dv \\
&\quad + \gamma \Psi_\infty \int_0^\infty \operatorname{tr} \left((\nabla^2 \mathcal{R}) \mathcal{M} \exp(-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)v) \right) dv \\
&= \gamma \operatorname{tr} \left((\nabla^2 \mathcal{R}) (\mathcal{M} \Psi_\infty + \mathcal{A}) (2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d))^{-1} \right).
\end{aligned} \tag{50}$$

The result for the limiting risk value $\lim_{t \rightarrow \infty} \Omega_t - \mathcal{R}(\mathcal{X}_{\Gamma(t)}^{\text{gf}})$ follows.

It remains to show that Ψ_t is bounded and exists at infinity with its limiting value given by (48). Recall the loss kernel for Ψ_t given by

$$K(t, s) \stackrel{\text{def}}{=} K(t, s; \mathbf{A}^T \mathbf{A}) = \gamma^2(s) \operatorname{tr} \left(\mathcal{M} \mathbf{A}^T \mathbf{A} \exp(-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)(\Gamma(t) - \Gamma(s))) \right), \tag{51}$$

so that Ψ_t is the solution to the Volterra equation

$$\Psi_t = \mathcal{L}(\mathcal{X}_{\Gamma(t)}^{\text{gf}}) + \int_0^t K(t, s) \Psi_s ds. \tag{52}$$

Under the kernel norm bounded by 1, (47), we show that the kernel $K(s, t)$ is of L^∞ -type on $[0, \infty)$. A kernel is L^∞ -type if $\|K\|_{L^\infty(J)} < \infty$ for a set $J \subset \mathbb{R}$ where $\|K\|_{L^\infty(J)} = \sup_{t \in J} \int_J |K(s, t)| ds$ [22, Chapter 9.2]. For this, we see that for each t and s

$$K(t, s) \leq \hat{\gamma} \cdot \gamma(s) \operatorname{tr} \left(\mathcal{M} \mathbf{A}^T \mathbf{A} \exp(-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)(\Gamma(t) - \Gamma(s))) \right). \tag{53}$$

This implies by change of variables that

$$\begin{aligned}
\int_0^t K(t, s) ds &\leq \int_0^{\Gamma(t)} \hat{\gamma} \operatorname{tr} \left(\mathcal{M} \mathbf{A}^T \mathbf{A} \exp(-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)(\Gamma(t) - s)) \right) ds \\
&\leq \frac{\hat{\gamma}}{2} \operatorname{tr} (\mathcal{M} \mathbf{A}^T \mathbf{A} (\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)^{-1}) < \infty.
\end{aligned} \tag{54}$$

Hence, it follows that the kernel K is L^∞ -type on $[0, \infty)$. To prove the boundedness assumption of Ψ_t , we will need something slightly stronger. We show that there exists a finite number of intervals J_i such that $\cup_i J_i = [0, \infty)$ and $\|K\|_{L^\infty(J_i)} \leq 1$. From this and Theorem 9.3.13 in [22], it will follow that the resolvent is also of type L^∞ on $[0, \infty)$. Since $\gamma(t) \rightarrow \gamma$, there exists a t_0 such that for all $t \geq t_0$, $\gamma(t) \leq \gamma + \varepsilon$. This $\varepsilon > 0$ can be chosen sufficiently small such that $\gamma + \varepsilon < 2(\operatorname{tr}(\mathcal{M} \mathbf{A}^T \mathbf{A} (\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)^{-1}))^{-1}$ (see (47) which gives an upper bound on γ). First, we observe that

$$\sup_{t \geq 0} \sup_{0 \leq s \leq t_0} K(t, s) \leq \hat{\gamma}^2 \operatorname{tr} (\mathcal{M} (\mathbf{A}^T \mathbf{A}) e^{2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)\Gamma(t_0)}) < \infty.$$

We break up the interval $[0, t_0]$ into finitely many intervals of length each of which has a length strictly less than $(\hat{\gamma}^2 \operatorname{tr} (\mathcal{M} \mathbf{A}^T \mathbf{A} e^{2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)\Gamma(t_0)})^{-1})$. If we denote these intervals by J_i , then it immediately follows by bounding the integral using the sup of K multiplied by the length of the interval J_i that

$$\|K\|_{L^\infty(J_i)} = \sup_{t \in J_i} \int_{J_i} K(t, s) ds < 1.$$

It only remains to show on the tail, that is, $J_\infty \stackrel{\text{def}}{=} (t_0, \infty)$, for which $\|K\|_{L^\infty(J_\infty)} < 1$. Using the same change of variables as in (53) and our choice of t_0 , we have that for all $t \geq t_0$

$$\begin{aligned} \int_{t_0}^t K(t, s) \, ds &\leq \int_{t_0}^t (\gamma + \varepsilon) \gamma(s) \operatorname{tr} \left(\mathcal{M} \mathbf{A}^T \mathbf{A} \exp \left(-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)(\Gamma(t) - \Gamma(s)) \right) \right) \, ds \\ &= \int_{\Gamma(t_0)}^{\Gamma(t)} (\gamma + \varepsilon) \operatorname{tr} \left(\mathcal{M} \mathbf{A}^T \mathbf{A} \exp \left(-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)(\Gamma(t) - s) \right) \right) \, ds \\ &\leq \frac{\gamma + \varepsilon}{2} \operatorname{tr} \left(\mathcal{M} \mathbf{A}^T \mathbf{A} (\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)^{-1} \right) < 1. \end{aligned}$$

The last inequality following by our assumption on $\gamma + \varepsilon$ being sufficiently small. By Theorem 9.3.13 in [22], we have that the resolvent is also of type L^∞ on $[0, \infty)$. We also have that $K(t, s)$ is of bounded type, that is the kernel is bounded (see [22, Definition 9.5.2] for precise definition). Since the forcing term $\mathcal{L}(\mathcal{X}_{\Gamma(t)}^{\text{gf}})$ and $\int_0^t A(t, s; \mathbf{A}^T \mathbf{A}) \, ds$ are bounded, then it follows by [22, Theorem 9.5.4] that the solution to the Volterra equation (52), Ψ_t , is bounded.

We now show that Ψ_t exists at infinity. Fix a $\varepsilon > 0$. By the assumptions on the learning rate, there exists a $t_0 > 0$ such that for all sufficiently large $t \geq s \geq t_0$

$$\gamma - \varepsilon \leq \gamma(t) \leq \gamma + \varepsilon \quad \text{and} \quad (\gamma - \varepsilon)(t - s) \leq \Gamma(t) - \Gamma(s) \leq (\gamma + \varepsilon)(t - s). \quad (55)$$

Using these inequalities for $\gamma(t)$, we get an upper bound and lower bound on the kernel $K(t, s)$ which we denote by $\overline{K}(t, s)$ and $\underline{K}(t, s)$, respectively. Specifically for all $t, s \geq t_0$,

$$\begin{aligned} K(t, s) &\leq \overline{K}(t, s) \stackrel{\text{def}}{=} (\gamma + \varepsilon)^2 \operatorname{tr} \left(\mathcal{M} \mathbf{A}^T \mathbf{A} \exp \left(-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)(\gamma - \varepsilon)(t - s) \right) \right) \\ K(t, s) &\geq \underline{K}(t, s) \stackrel{\text{def}}{=} (\gamma - \varepsilon)^2 \operatorname{tr} \left(\mathcal{M} \mathbf{A}^T \mathbf{A} \exp \left(-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)(\gamma + \varepsilon)(t - s) \right) \right). \end{aligned} \quad (56)$$

The kernels $\overline{K}(t, s)$ and $\underline{K}(t, s)$ are substantially nicer than the original $K(t, s)$ because they are proper convolution kernels. Here one can define $\overline{K} : [0, \infty) \rightarrow \mathbb{R}$ by

$$\overline{K}(t) \stackrel{\text{def}}{=} (\gamma + \varepsilon)^2 \operatorname{tr} \left(\mathcal{M} \mathbf{A}^T \mathbf{A} \exp \left(-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)(\gamma - \varepsilon)t \right) \right).$$

Then it follows that $\overline{K}(t, s) = \overline{K}(t - s)$. A similar result holds for $\underline{K}(t, s)$.

For ease of notation, define the forcing function: for $t \geq t_0$

$$F(t) \stackrel{\text{def}}{=} \mathcal{L}(\mathcal{X}_{\Gamma(t)}^{\text{gf}}) + \int_0^{t_0} K(t, s) \Psi_s \, ds + \int_0^t A(t, s; \mathbf{A}^T \mathbf{A}) \, ds, \quad (57)$$

where Ψ_t is a solution to (52). Similar to the definitions of $\overline{K}(t)$ and $\underline{K}(t)$, we define $\overline{F}(t)$ and $\underline{F}(t)$ respectively as

$$\underline{F}(t) \leq F(t) \leq \overline{F}(t), \quad (58)$$

$$\begin{aligned} \text{where } \overline{F}(t) &\stackrel{\text{def}}{=} \mathcal{L}(\mathcal{X}_{\Gamma(t)}^{\text{gf}}) + \int_0^{t_0} K(t, s) \Psi_s \, ds + \int_0^{t_0} A(t, s; \mathbf{A}^T \mathbf{A}) \, ds \\ &\quad + \int_0^t (\gamma + \varepsilon)^2 \operatorname{tr} \left(\mathcal{M} \mathbf{A}^T \mathbf{A} \exp \left(-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)(\gamma - \varepsilon)(t - s) \right) \right) \, ds \\ \text{and } \underline{F}(t) &\stackrel{\text{def}}{=} \mathcal{L}(\mathcal{X}_{\Gamma(t)}^{\text{gf}}) + \int_0^{t_0} K(t, s) \Psi_s \, ds + \int_0^{t_0} A(t, s; \mathbf{A}^T \mathbf{A}) \, ds \\ &\quad + \int_0^t (\gamma - \varepsilon)^2 \operatorname{tr} \left(\mathcal{M} \mathbf{A}^T \mathbf{A} \exp \left(-2(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)(\gamma + \varepsilon)(t - s) \right) \right) \, ds \end{aligned} \quad (59)$$

Because Ψ_s is bounded, it follows that $\lim_{t \rightarrow \infty} \int_0^{t_0} K(t, s) \Psi_s \, ds = \lim_{t \rightarrow \infty} A(t, s; \mathbf{A}^T \mathbf{A}) = 0$. Also it is clear that the $F(t)$, $\overline{F}(t)$, and $\underline{F}(t)$ are bounded.

Using the upper/lower bound on the kernel (56), we can squeeze the value of Ψ_t between two expressions: for $t, s \geq t_0$,

$$\underline{F}(t) + \int_{t_0}^t \underline{K}(t, s) \Psi_s \, ds \leq \Psi_t \leq \overline{F}(t) + \int_{t_0}^t \overline{K}(t, s) \Psi_s \, ds. \quad (60)$$

Using a similar argument for $K(t, s)$ and choosing ε sufficiently small, $\overline{K}(t, s)$ and $\underline{K}(t, s)$ are L^∞ -type on $[0, \infty)$. Moreover using a similar argument as we did for K itself, the norms $\|\overline{K}\|_{L^\infty([0, \infty))} < 1$ and $\|\underline{K}\|_{L^\infty([0, \infty))} < 1$. Here we used the upper bound on γ in (47) and a sufficiently small ε . Note we do not need to break up into finite intervals. As before, the resolvent then is of L^∞ -type on $[0, \infty)$ [22, Corollary 9.3.10]. Further because of non-negativity, Proposition 9.8.1 in [22] yields that the resolvents are also non-negative.

Consider the upper bound (a similar argument will hold for the lower bound). We can apply Gronwall's inequality (60) [22, Theorem 9.8.2]. It follows that Ψ_t is upper bounded (lower bounded) by the solutions $\overline{\Psi}_t$ ($\underline{\Psi}_t$) to the following convolution Volterra equations

$$\overline{\Psi}_t = \overline{F}(t) + \int_{t_0}^t \overline{K}(t, s) \overline{\Psi}_s \, ds \quad \text{and} \quad \underline{\Psi}_t = \underline{F}(t) + \int_{t_0}^t \underline{K}(t, s) \underline{\Psi}_s \, ds.$$

Specifically, we have $\underline{\Psi}_t \leq \Psi_t \leq \overline{\Psi}_t$ for all $t \geq t_0$. Since $\overline{\Psi}_t$ and $\underline{\Psi}_t$ are solutions to a proper convolution-type Volterra equation and both functions $\overline{F}(t)$, and $\underline{F}(t)$ have limits at infinity ($\overline{F}(\infty) \stackrel{\text{def}}{=} \lim_{t \rightarrow \infty} \overline{F}(t)$ and $\underline{F}(\infty) \stackrel{\text{def}}{=} \lim_{t \rightarrow \infty} \underline{F}(t)$), by [5], for $t \geq t_0$

$$\limsup_{t \rightarrow \infty} \Psi_t \leq \limsup_{t \rightarrow \infty} \overline{\Psi}_t = \overline{F}(\infty) (1 - \|\overline{K}\|_{L^\infty([0, \infty))})^{-1} \leq \overline{F}(\infty) (1 - \|\overline{K}\|_{L^\infty([0, \infty))})^{-1}, \quad (61)$$

and similarly, the lower bound gives

$$\liminf_{t \rightarrow \infty} \Psi_t \geq \liminf_{t \rightarrow \infty} \underline{\Psi}_t \leq \underline{F}(\infty) (1 - \|\underline{K}\|_{L^\infty([0, \infty))})^{-1}. \quad (62)$$

A simple computation yields that

$$\begin{aligned} \|\overline{K}\|_{L^\infty([0, \infty))} &= \frac{(\gamma + \varepsilon)^2}{\gamma - \varepsilon} G(\mathcal{M}) \quad \text{and} \quad \|\underline{K}\|_{L^\infty([0, \infty))} = \frac{(\gamma - \varepsilon)^2}{\gamma + \varepsilon} G(\mathcal{M}) \\ \text{and} \quad \overline{F}(\infty) &= \mathcal{L}(\mathbf{X}_\infty^{\text{gf}}) + \frac{(\gamma + \varepsilon)^2}{\gamma - \varepsilon} G(\mathcal{A}) \quad \text{and} \quad \underline{F}(\infty) = \mathcal{L}(\mathbf{X}_\infty^{\text{gf}}) + \frac{(\gamma - \varepsilon)^2}{\gamma + \varepsilon} G(\mathcal{A}) \end{aligned} \quad (63)$$

where $G(\mathcal{H}) \stackrel{\text{def}}{=} \frac{1}{2} \text{tr}(\mathcal{H} \mathbf{A}^T \mathbf{A} (\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)^{-1})$.

So for any sufficiently small $\varepsilon > 0$, we have that

$$\begin{aligned} \left(1 - \frac{(\gamma - \varepsilon)^2}{\gamma + \varepsilon} G(\mathcal{M})\right)^{-1} \times \left\{ \mathcal{L}(\mathbf{X}_\infty^{\text{gf}}) + \frac{(\gamma - \varepsilon)^2}{\gamma + \varepsilon} G(\mathcal{A}) \right\} &\leq \liminf_{t \rightarrow \infty} \Psi_t \\ &\leq \limsup_{t \rightarrow \infty} \Psi_t \leq \left(1 - \frac{(\gamma + \varepsilon)^2}{\gamma - \varepsilon} G(\mathcal{M})\right)^{-1} \times \left\{ \mathcal{L}(\mathbf{X}_\infty^{\text{gf}}) + \frac{(\gamma + \varepsilon)^2}{\gamma - \varepsilon} G(\mathcal{A}) \right\}. \end{aligned} \quad (64)$$

As this holds for any sufficiently small ε , the result follows by sending $\varepsilon \rightarrow 0$. \square

F Algorithmic regularization

In this section, we discuss the exact asymptotic convergence rates for SGD and full batch momentum algorithms on high-dimensional ℓ^2 -regularized least squares problems. The results in this section (e.g., Theorems 4 and 5) were shown in a series of papers [58–60] that explored exact trajectories of loss function.

F.1 Convergence rates of SGD

To characterize the rates, we define λ_{\min} as the smallest non-zero eigenvalue of $\mathbf{A}^T \mathbf{A}$. Then for generic initial conditions, (in particular almost surely if \mathbf{X}_0 is isotropic norm 1), then

$$\lim_{t \rightarrow \infty} \left(\mathcal{L}(\mathbf{X}_t^{\text{gf}}) - \mathcal{L}(\mathbf{X}_\infty^{\text{gf}}) \right)^{1/t} = \begin{cases} e^{-\gamma(\lambda_{\min} + \delta)}, & \text{if } \delta > 0, \\ e^{-2\gamma\lambda_{\min}}, & \text{otherwise.} \end{cases}$$

The rate of convergence of Ψ_t to Ψ_∞ is given by, (for small γ), the rate above. For larger γ , another rate can frustrate the convergence. Recall the *Malthusian exponent* of the convolution Volterra equation in (11) is given by

$$\lambda_* = \inf \left\{ x : 1 = \int_0^\infty e^{xt} K(t; \mathbf{A}^T \mathbf{A}) dt = \gamma^2 \int_0^\infty e^{xt} \text{tr} \left(\mathcal{M} \mathbf{A}^T \mathbf{A} \exp(-2\gamma(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I})t) \right) dt \right\}. \quad (65)$$

The set may be empty, in which case the infimum is ∞ . We recall below Theorem 4.

Theorem 14. For $\gamma > 0$ satisfying $\mathcal{I}(\gamma) < 1$ (see (44)), define

$$\Xi(\gamma) \stackrel{\text{def}}{=} \begin{cases} \min\{\gamma(\lambda_{\min} + \delta), \lambda_*(\gamma)\} & \text{if } \delta > 0, \\ \min\{2\gamma\lambda_{\min}, \lambda_*(\gamma)\} & \text{if } \delta = 0. \end{cases} \quad (66)$$

Then the rates of convergence of both the training and test loss are

$$\lim_{t \rightarrow \infty} (\Psi_t - \Psi_\infty)^{1/t} = e^{-\Xi(\gamma)} = \lim_{t \rightarrow \infty} (\Omega_t - \Omega_\infty)^{1/t}$$

Furthermore, when $\gamma = n / \text{tr}(\mathbf{A}^T \mathbf{A})$, we have the rate guarantee $\Xi(\gamma) \geq \frac{\lambda_{\min} n}{2 \text{tr}(\mathbf{A}^T \mathbf{A})}$.

Proof. See [59, Theorem 1.2] for proof. \square

E.2 Momentum GD (MGD) rates

In this section, we consider a popular *deterministic* or *full-batch* algorithm for solving the ridge regression problem in (1), that is, gradient descent with momentum (a.k.a Polyak momentum). Throughout this section, we use the notation, $\mathbf{x}_t^{\text{m-gd}} = \mathbf{x}_t$. Gradient descent with momentum (MGD), initialized at $\mathbf{x}_0 \in \mathbb{R}^d$ and $\mathbf{x}_1 = \mathbf{x}_0 - \frac{\gamma}{1+m} \nabla f(\mathbf{x}_0)$, iterates for $k \geq 1$

$$\mathbf{x}_{k+1} = \mathbf{x}_k + m(\mathbf{x}_k - \mathbf{x}_{k-1}) - \gamma \nabla f(\mathbf{x}_k), \quad (67)$$

where $\gamma, m > 0$ are the stepsize and momentum parameters respectively. From Proposition 3.1 in [60], there exists k -degree polynomials P_k and Q_k such that the iterates of GD+M satisfy the following

$$\mathbf{x}_k = P_k(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}) \mathbf{x}_0 + Q_k(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}) \mathbf{A}^T \mathbf{b}, \quad \text{with } P_k(\lambda) = 1 - (\lambda) Q_k(\lambda) \quad (68)$$

and the coefficients of P_k and Q_k only depend on the largest and smallest eigenvalue of $\mathbf{A}^T \mathbf{A}$. For Polyak, similar to the work in [60, Section 3.1], we can give an explicit representation for these polynomials P_k and Q_k .

Proposition 1 (Polynomial representation of MGD). *Suppose $\mathbf{x}_0 \in \mathbb{R}^d$ and fix a stepsize $\gamma > 0$ and momentum parameter $m > 0$. For the iterates of GD+M on (1) with ridge parameter $\delta > 0$, we have the following representation for the polynomials*

$$\mathbf{x}_k = P_k(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}) \mathbf{x}_0 + Q_k(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}) \mathbf{A}^T \mathbf{b}, \quad (69)$$

where P_k and Q_k are k -degree polynomials satisfying

$$P_k(\lambda) = m^{k/2} \left(\frac{2m}{1+m} T_k(\sigma(\lambda)) + \left(1 - \frac{2m}{1+m} \right) U_k(\sigma(\lambda)) \right) \quad \text{and} \quad Q_k(\lambda) = \frac{1 - P_k(\lambda)}{\lambda}$$

$$\text{where } \sigma(\lambda) = \frac{1 + m - \gamma\lambda}{2\sqrt{m}}$$

and T_k, U_k are Chebyshev polynomials of the 1st and 2nd kind respectively.

(70)

Proof. The proof can be found in [60, Appendix A.2] or [19, Chapter 11]. We include a sketch of the proof. From the recurrence in (67) and the gradient of the ridge regression, the polynomials P_k that generate GD+M satisfy the following three-term recurrence,

$$\begin{aligned} P_{k+1}(\lambda) &= (1 + m - \gamma\lambda) P_k(\lambda) - m P_{k-1}(\lambda) \\ P_k(\lambda) &= \frac{P_{k+1}(\lambda) + m P_k(\lambda)}{1 + m - \gamma\lambda}. \end{aligned} \quad (71)$$

We define the polynomial generating function for P_k as $\mathcal{G}(\lambda, t) = \sum_{k=0}^{\infty} t^k P_k(\lambda)$. Using the recurrence in (71), we get that

$$\begin{aligned}\mathcal{G}(\lambda, t) &= 1 + \frac{1}{t(1-m+\gamma\lambda)} \sum_{k=2}^{\infty} t^k P_k(\lambda) - \frac{mt}{1-m+\gamma\lambda} \sum_{k=0}^{\infty} t^k P_k(\lambda) \\ &= 1 + \frac{1}{t(1-m+\gamma\lambda)} [\mathcal{G}(\lambda, t) - 1 - t(1 - \frac{\gamma}{1+m}\lambda)] - \frac{mt}{1-m+\gamma\lambda} \mathcal{G}(\lambda, t).\end{aligned}$$

By solving this expression for the generating polynomial, we have

$$\mathcal{G}(\lambda, t) = \frac{1 + t(m - (\gamma + \frac{\gamma}{1+m}))}{1 - t(1 - m + \gamma\lambda) - mt^2}.$$

This generating function for MGD closely resembles the generating function for Chebyshev polynomials of the 1st and 2nd kind. Under simple transformations (e.g., $t \mapsto \frac{t}{\sqrt{m}}$), this is exactly the case. These transformations yield the expression in (70). \square

The role of $\sigma(\lambda)$ is to transform the eigenvalues of $\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d$ within a specific range controlled by the learning rate and momentum. It is known that the Chebyshev polynomials are well-behaved on the interval of $[-1, 1]$ and grow exponentially off of this region.

Moreover for a generic quadratic applied to \mathbf{x}_k , the rate of convergence will be controlled by $P_k(\lambda)$. Using standard asymptotic behavior of Chebyshev polynomials, we can derive asymptotic rates based on $\lambda_{\min} \stackrel{\text{def}}{=} \lambda_{\min}(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)$ and $\lambda_{\max} \stackrel{\text{def}}{=} \lambda_{\max}(\mathbf{A}^T \mathbf{A} + \delta \mathbf{I}_d)$, the smallest (non-zero) and largest eigenvalues of $\mathbf{A}^T \mathbf{A}$ respectively. We record this result below

Proposition 2 (Asymptotic rates of MGD). *The asymptotic rate of MGD is*

$$\limsup_{k \rightarrow \infty} \sqrt[k]{P_k} = \begin{cases} \sqrt{m} & \text{if } \gamma \in \left[\frac{(1-\sqrt{m})^2}{\lambda_{\min}}, \frac{(1+\sqrt{m})^2}{\lambda_{\max}} \right] \\ \sqrt{m} (|\sigma(\lambda_{\min})| + \sqrt{\sigma(\lambda_{\min})^2 - 1}) & \text{if } \gamma \in \left[0, \min \left\{ \frac{2(1+m)}{\lambda_{\min} + \lambda_{\max}}, \frac{(1-\sqrt{m})^2}{\lambda_{\min}} \right\} \right] \\ \sqrt{m} (|\sigma(\lambda_{\max})| + \sqrt{\sigma(\lambda_{\max})^2 - 1}) & \text{if } \gamma \in \left[\max \left\{ \frac{2(1+m)}{\lambda_{\min} + \lambda_{\max}}, \frac{(1+\sqrt{m})^2}{\lambda_{\max}} \right\}, \frac{2(1+m)}{\lambda_{\max}} \right] \\ \geq 1 & \text{otherwise.} \end{cases} \quad (72)$$

Proof. See [62] for a complete proof. The result follows from knowing that the iterates are given by Chebyshev polynomials and then applying well-known asymptotics of Chebyshev polynomials to get the convergence rate. \square

We can minimize over the rate to find the optimal parameters. In this case, they become the parameters used in the Heavy-Ball algorithm [65] where

$$m = \left(\frac{\sqrt{\lambda_{\max}} - \sqrt{\lambda_{\min}}}{\sqrt{\lambda_{\max}} + \sqrt{\lambda_{\min}}} \right)^2 \quad \text{and} \quad \gamma = \left(\frac{2}{\sqrt{\lambda_{\max}} + \sqrt{\lambda_{\min}}} \right)^2. \quad (73)$$

A simple computation yields that the asymptotic rate for Heavy-Ball is $\frac{\sqrt{\lambda_{\max}} - \sqrt{\lambda_{\min}}}{\sqrt{\lambda_{\max}} + \sqrt{\lambda_{\min}}}$.

G Numerical simulations

To illustrate our theoretical results and conjectures we report simulations and experiments using SGD with constant learning rate on the ℓ^2 -regularized least squares problem. In all simulations for the random ℓ^2 -regularized least-square problem, the vectors $\boldsymbol{\eta}$, and $\boldsymbol{\beta}$ are sampled from a standard Gaussian and the initialization vector $\mathbf{x}_0 = \mathbf{0}$ (for Figures 1 and 5) and $N(0, 4\mathbf{I}_d)$ (Figure 6). For the random features model (see Section B.1 and Figure 1, a standardized ReLu activation function was applied, that is

$$\sigma(\cdot) = \frac{\max\{\cdot, 0\} - 0.5(\pi)^{-1}}{0.5 - 0.5\pi^{-1}}. \quad (74)$$

The entries of the hidden weight matrix $\mathbf{W} \in \mathbb{R}^{n_0 \times d}$ in the random feature model are standard normal.




Figure 6: **Concentration of mean squared error (MSE) and expected test loss, $\frac{1}{2}\|x - \beta\|_2^2$, for SGD on a Gaussian random ℓ^2 -regularized least-squares problem (Section B) where $\beta \sim N(\mathbf{0}, \mathbf{I}_d)$ is the ground truth signal and a generative model $\mathbf{b} = \mathbf{A}\beta + \eta$ where entries of η iid standard normal with $\|\eta\|_2^2 = 2.25$, $n = 0.9d$ with ℓ^2 -regularization parameter $\delta = 0.1$. SGD with constant learning rate $\gamma = 0.8$ was initialized at $\mathbf{x}_0 \sim N(\mathbf{0}, 4\mathbf{I}_d)$ (independent of \mathbf{A} , β); an 80% confidence interval (shaded region) over 10 runs for each n . Any quadratic statistic, such as the MSE, becomes non-random in the large limit and all runs of SGD converge to a deterministic function Ω_t (red) solving a Volterra equation (9). This is an illustration of Theorem 1 and Theorem 2.**

Volterra: computing theoretical dynamics. When the entries of \mathbf{A} are generated by standard Gaussians, a celebrated work [50] gives an explicit limiting density for the eigenvalues when d and n are proportional. In this case, the Volterra equation (8) for the loss function \mathcal{L} is computable without needing to input the empirical eigenvalues of the data matrix $\mathbf{A}\mathbf{A}^T$. Since the covariance of standard Gaussians is explicitly $\Sigma_f = \mathbf{I}_d$ (see Appendix B.1), one can also directly solve for the expected risk (9) for applications such as in-distribution expected risk. As such, the Volterra equation is completely determined. To solve it, a Chebyshev quadrature was used to derive a numerical approximation for the kernel, K , (8). The size of the grid points used to compute the numerical integration does effect the Volterra equations convergence to the theoretical limit. We suggest that the number of epochs be equal to the number of grid points used in the numerical quadrature rule. Next, to generate the solution \mathcal{L} of the Volterra equation, we implement a Picard iteration which finds a fix point to the Volterra equation by repeatedly convolving the kernel and adding the forcing term. Despite the numerical approximations to integrals, the resulting solutions to the Volterra equation (Ψ and Ω) model the true behavior of SGD remarkably well. Similarly, by evaluating contour integrals, random features with Gaussian \mathbf{X} and \mathbf{W} known explicit formulas for the limiting densities of eigenvalues and eigenvectors (see e.g., [2]). This approach was used to compute the theoretical dynamics in Figure 6.

When the limiting eigenvalues and eigenvectors are unavailable, as in the case of real data sets, an empirical Volterra equation solver was used. We computed the svd of the data matrix \mathbf{A} and calculated an empirical covariance for Σ_f (see Appendix B.1). The singular values and vectors of \mathbf{A} and Σ_f were then used to compute the forcing term (i.e., the GF terms $\mathcal{L}(\mathcal{X}^{\text{sf}})$ and $\mathcal{R}(\mathcal{X}^{\text{sf}})$) and kernel K (10). As before, a Chebyshev quadrature was used to derive the integral for the kernel K and a Picard iteration to find the fix point of the Volterra was applied. This method was used to compute the theoretical dynamics Ω_t and Ψ_t in Figures 1 and 5.

Real data. The CIFAR-5m [54] example (Figures 1 is shown to demonstrate that large-dimensional random matrix predictions often work for large dimensional real data. Random features models were used to predict the car/plane class vector which has approximately 1 million samples. The data sets were all standardized and pre-processed to have mean 0 and variance 1 before applying the random features model with standardized ReLu.

We give specific simulation/experimental details below:

- *CIFAR-5m streaming, Figure 1:* Plots of single runs of SGD on CIFAR-5m [54] using the car/plane class vector (samples $n = 1$ million, features $n_0 = 32 \times 32 \times 3$) on a random features model with standardized ReLu (see (74)). CIFAR-5m car/plane data set was standardized so that entries were mean 0 and variance 1. Standard Gaussian weight matrix $\mathbf{W} \in \mathbb{R}^{n_0 \times d}$ with fixed $d = 6,000$ used in the random features set-up (see Appendix B.1). Multi-pass SGD with constant learning rate $\gamma = 0.8$ applied to various sample size $n = 1000 \cdot [4, 6, 10, 20, 40]$

Table 1: **Summary of the eigenvalues in ICR** with normalized trace equal to 1.0, i.e., $\frac{1}{n} \text{tr}(\mathbf{A}\mathbf{A}^T) = 1.0$. All data sets were standardized before applying any transformations (e.g., random features). For random features (RF), standard Gaussian $\mathbf{W} \in \mathbb{R}^{n_0 \times d}$ applied to the data set followed by entry-wise application of standardized ReLu (see (74) and Appendix B.1 and Appendix G for exact set-up).

Data set	Samples (n)	Features (d)	Eigenvalues of $\mathbf{A}\mathbf{A}^T$	
			Largest	Smallest
CIFAR-10 ¹ (all)	50,000	3,072	11,118.80	$4.7 \cdot 10^{-4}$
CIFAR-10 ¹ RF large d	50,000	5,551	8,162.84	$2.8 \cdot 10^{-1}$
CIFAR-10 ¹ RF small d	50,000	452	8,403.31	13.18
CIFAR-5m ² (all)	5 million	3,072	1,195,595.52	$1.03 \cdot 10^{-1}$
CIFAR-5m ² (car/plane)	1 million	3,072	258,599.09	$1.7 \cdot 10^{-2}$
Gaussian under parameterized	2,000	100	29.35	12.4
Gaussian equal	2,000	1,930	4.06	$3.4 \cdot 10^{-4}$
Gaussian over parameterized	2,000	100,000	1.30	$7.4 \cdot 10^{-1}$
Gaussian-RF under parameterized	2,000	100	66.38	3.95
Gaussian-RF equal	2,000	1,467	27.15	$6.2 \cdot 10^{-3}$
Gaussian-RF over parameterized	2,000	316,227	21.77	$6.6 \cdot 10^{-2}$
MNIST ³ (all)	60,000	784	5,562.79	$1.1 \cdot 10^{-2}$
MNIST ³ RF large d	60,000	5,551	4,249.29	$2.8 \cdot 10^{-1}$
MNIST ³ RF small d	60,000	452	4,564.77	15.09

¹ [34] ² [54] ³ [37]

on (1) with $\delta = 0.01$. Empirical volterra solver was applied to match the multi-pass setting using the same variables. An empirical covariance $\Sigma_\sigma(\mathbf{W})$ computed using all 1 million samples. Streaming SGD using constant learning rate $\gamma = 0.8$ applied to the expected risk using the empirical covariance $\Sigma_\sigma(\mathbf{W})$. As the ℓ^2 regularization parameter δ is hit by a factor of n , in the streaming setting, the regularization is set to 0.0. Empirical Volterra using the eigenvalues of $\Sigma_\sigma(\mathbf{W})$ with $\gamma = 0.8$ and $\delta = 0.0$ matched the SGD steaming setting.

- *Random features theory.*
- *ICR, Figure 2:* Graph of the ICR under the assumption that the normalized trace of $\nabla^2 \mathcal{L}$ is 1.0, that is, $\frac{1}{n} \text{tr}(\nabla^2 \mathcal{L}) = 1.0$. All data sets, MNIST, CIFAR-10, and CIFAR-5m are standardized (i.e., entries normalized so that mean 0.0 and variance 1.0). Largest and smallest (non-zero) eigenvalues of the feature covariance reported. For the random features set-up (RF), standard Gaussian matrix $\mathbf{W} \in \mathbb{R}^{n_0 \times d}$ where n_0 is the underlying number of features from the data set and d ranged from $10^{2.5}$ to $10^{3.9}$ was applied to the data set followed by an entry-wise activation standardized ReLu. Reported (dashed lines) are the largest and smallest eigenvalues after applying the standardized ReLu and making the normalized trace equal to 1.0. In the Gaussian set-up, the number of samples n was fixed at 2000 and d ranged from 10^2 to 10^5 ; entries of \mathbf{A} standard Gaussians. In the random features Gaussian (Gaussian-RF), we fixed the samples $n = 2000$ and $n_0 = 100$ and varied

the $d = 10^2$ to $10^{5.5}$. Largest and smallest eigenvalues of $\sigma(\mathbf{X}\mathbf{W})^T \sigma(\mathbf{X}\mathbf{W})$ reported after making the normalized trace 1.0.

- *Gaussian linear regression streaming, Figure 5:* Simple linear regression with targets from a generative model, $\mathbf{b} = \mathbf{A}\boldsymbol{\beta} + \boldsymbol{\xi}$; signal $\boldsymbol{\beta} \sim N(\mathbf{0}, \frac{1}{d}\mathbf{I}_d)$ and noise $\boldsymbol{\xi} \sim N(\mathbf{0}, \frac{0.04}{d}\mathbf{I})$. A $(n \times 2000)$ data matrix \mathbf{A} with $A_{ij} \sim N(0, 1/2000)$ with various n values (see figure). SGD with constant learning rate $\gamma = 0.8$ initialized at $\mathbf{x}_0 = \mathbf{0}$ was applied to the linear regression problem with a regularization parameter of 0.01, see training loss and excess risk in linear regression in Appendix B.1. In this setting, the covariance of the expected risk is explicitly given by \mathbf{I}_d/d . A new data point $\mathbf{a} \sim N(0, \frac{1}{d}\mathbf{I}_d)$ and $b = \mathbf{a}\boldsymbol{\beta} + 0.2Z$ with $Z \sim N(0, 1)$ generated and the expected risk computed as $(\mathbf{a}\mathbf{x}_t - b)^2$ where \mathbf{x}_t are the iterates of SGD. Empirical volterra solver used with grid points \approx number of iterations of SGD.
- *Gaussian linear regression concentration, Figure 6:* Simple linear regression with targets from generative model, $\mathbf{b} = \mathbf{A}\boldsymbol{\beta} + \boldsymbol{\xi}$; signal $\boldsymbol{\beta} \sim N(\mathbf{0}, \frac{1}{d}\mathbf{I}_d)$, noise $\boldsymbol{\xi} \sim N(\mathbf{0}, \frac{1.5^2}{n}\mathbf{I}_n)$. Matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ is row normalized and $\frac{d}{n} = 0.9$ for $n = \{100, 400, 1600, 6400\}$. 10 runs of SGD with constant learning $\gamma = 0.8$ started at $\mathbf{x}_0 \sim N(0, \frac{4}{n}\mathbf{I}_d)$ applied to the ℓ^2 -regularized least squares problem with $\delta = 0.1$, see training loss and excess risk in linear regression in Appendix B.1. 80% confidence interval (shaded) depicted in Figure 6. Volterra equation solver used with grid points approximately the same as epochs. Expected risk computed as in Figure 5. Concentration around the Volterra equation occurs as n (or d) $\rightarrow \infty$ across different risk functions.