

Halting Time is Predictable for Large Models: A Universality Property and Average-case Analysis

Courtney Paquette*[†] Bart van Merriënboer* Elliot Paquette[†]
Fabian Pedregosa*

October 27, 2020

Abstract

Average-case analysis computes the complexity of an algorithm averaged over all possible inputs. Compared to worst-case analysis, it is more representative of the typical behavior of an algorithm, but remains largely unexplored in optimization. One difficulty is that the analysis can depend on the probability distribution of the inputs to the model. However, we show that this is not the case for a class of large-scale problems trained with first-order methods including random least squares and one-hidden layer neural networks with random weights. In fact, the halting time exhibits a *universality property*: it is independent of the probability distribution. With this barrier for average-case analysis removed, we provide the first explicit average-case convergence rates showing a tighter complexity not captured by traditional worst-case analysis. Finally, numerical simulations suggest this universality property holds for a more general class of algorithms and problems.

Key words. universality, random matrix theory, optimization

1 Introduction

Traditional worst-case analysis of optimization algorithms provides complexity bounds for any input, no matter how unlikely [Nemirovski, 1995, Nesterov, 2004]. It gives convergence guarantees, but the bounds are not always representative of the typical runtime of an algorithm. In contrast, average-case analysis gives sharper runtime estimates when some or all of its inputs are random. This is often paired with concentration bounds that quantify the spread of those estimates. In this way, it is representative of the typical behavior.

Yet, average-case analysis is rarely used in optimization because the complexity of algorithms is assumed to depend on the specific probability distribution of the inputs. Surprisingly, simulations reveal this is not the case for large-scale problems (see Figure 1).




Figure 1: As the model grows (x -axis), the standard deviation (shaded region) in the halting time of gradient descent on random least squares vanishes and the halting time becomes **predictable**. Note also a **universality** phenomenon, that is, the halting time limit is the same for problems generated from different distributions.

*Google Research, Brain Team

[†]Department of Mathematics and Statistics, McGill University, Montreal, QC, Canada, H3A 0B9; CP is a CIFAR AI chair; <https://cypaquette.github.io/>. Research by EP was supported by a Discovery Grant from the Natural Science and Engineering Research Council (NSERC) of Canada.; <https://elliottaquette.github.io/>.

We show that almost all instances of high-dimensional data are indistinguishable to first-order algorithms. Particularly, the *halting time*, i.e. the number of iterations to reach a given accuracy, for any first-order method converges to a deterministic value which is independent of the input distribution (see Figure 1). Since the halting time is deterministic, the empirical complexity coincides almost surely with the average-case rates.

Table 1: **Comparison of convergence guarantees for non-strongly convex objectives** in terms of asymptotic behavior of $\|\nabla f(\mathbf{x}_k)\|^2$ as problem size and iteration are large in the isotropic features model. The average-case guarantees are strictly faster than the traditional worst-case and adversarial rates. Furthermore, the traditional worst-case complexity bounds depend on the distance to the optimum which under our constant signal-to-noise model, grows as the problem size, or dimension, d , increases. The rates are stated in terms of an absolute constant C , the amount of signal R and noise \tilde{R} , and the maximum λ^+ and minimum λ^- eigenvalues. Denote $\|J_1^2(x)\|_\infty$ the maximum value of the squared Bessel function of the first kind ($J_1(x)$). See Section 1.2 and 5 for derivations.

Method		Non-strongly cvx w/o noise	Non-strongly cvx w/ noise ⁵
Gradient descent (GD)	Worst ¹	$\frac{1}{(k+1)^2} \cdot R^2(\lambda^+)^2$	$\frac{d}{(k+1)^2} \cdot \tilde{R}^2(\lambda^+)^2 C$
	Adversarial ⁴	$\frac{1}{(k+1)^2} \cdot \frac{R^2(\lambda^+)^2}{e^2}$	$\frac{1}{k} \cdot \frac{\tilde{R}^2 \lambda^+}{2}$
	Average	$\frac{1}{k^{5/2}} \cdot \frac{R^2(\lambda^+)^2 \Gamma(5/2)}{2^{3/2} \pi}$	$\frac{1}{k^{3/2}} \cdot \frac{\tilde{R}^2 \lambda^+ \Gamma(3/2)}{2^{1/2} \pi}$
Nesterov accelerated method ²	Worst ³	$\frac{1}{k(k+2)^2} \cdot 8R^2(\lambda^+)^2$	$\frac{d}{k(k+2)^2} \cdot 8\tilde{R}^2(\lambda^+)^2 C$
	Adversarial	$\frac{1}{k^{7/2}} \cdot \frac{8e^{-1/2}}{\sqrt{2\pi}} R^2(\lambda^+)^2$	$\frac{1}{k^2} \cdot \ J_1^2(x)\ _\infty \tilde{R}^2 \lambda^+$
	Average	$\frac{1}{k^4} \cdot \frac{8R^2(\lambda^+)^2}{\pi^2}$	$\frac{\log(k)}{k^3} \cdot \frac{4\tilde{R}^2 \lambda^+}{\pi^2}$

¹ Taylor et al. [2017] ² Beck and Teboulle [2009], Nesterov [2004] ³ Nesterov [2012] ⁴ Adversarial model maximizes the norm of the gradient subject to a fixed condition number (see Section 1.2). ⁵ When noise is added, the convergence rates are dominated by this noise term.

1.1 Main results

In this paper, we analyze the halting time and develop the first explicit average-case analysis for first-order methods on quadratic objectives. Quadratic objective functions are rich enough to reproduce the dynamics that arise in more complex models, yet simple enough to be understood in closed form. Quadratic models are receiving renewed interest in the machine learning community as recent advances have shown that over-parameterized models, including neural networks, have training dynamics similar to those of quadratic problems [Arora et al., 2019, Chizat et al., 2019, Jacot et al., 2018, Novak et al., 2019].

The precise form of the quadratic problem we consider is

$$\arg \min_{\mathbf{x} \in \mathbb{R}^d} \left\{ f(\mathbf{x}) \stackrel{\text{def}}{=} \frac{1}{2n} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2 \right\}, \quad \text{with } \mathbf{b} \stackrel{\text{def}}{=} \mathbf{A}\tilde{\mathbf{x}} + \boldsymbol{\eta}, \quad (1)$$

Table 2: **Comparison of convergence guarantees for strongly convex objectives** in terms of asymptotic behavior of $\|\nabla f(\mathbf{x}_k)\|^2$ as problem size and iteration are large in the isotropic features model. Average-case matches the worst-case asymptotic guarantees multiplied by an additional *polynomial correction term* (green). This polynomial term has little effect on the complexity compared to the linear rate. However as the matrix $\mathbf{H} = \frac{1}{n}\mathbf{A}^T\mathbf{A}$ becomes ill-conditioned ($r \rightarrow 1$), the polynomial correction starts to dominate the average-case complexity. Indeed this shows that the support of the spectrum does not fully determine the rate. *Many* eigenvalues contribute meaningfully to the average-rate. See Section 5 for derivations and Table 1 for definition of terms in the rates.

Method	Strongly cvx w/ noise	
Gradient descent (GD)	Worst ¹	$(1 - \frac{\lambda^-}{\lambda^+})^{2k} (\lambda^+)^2$
	Average	$(1 - \frac{\lambda^-}{\lambda^+})^{2k} \frac{1}{k^{3/2}} [R^2\lambda^- + \tilde{R}^2r] \cdot C$
Polyak [Polyak, 1964]	Worst	$(1 - \frac{2\sqrt{\lambda^-}}{\sqrt{\lambda^+} + \sqrt{\lambda^-}})^{2k} \cdot C$
	Average	$(1 - \frac{2\sqrt{\lambda^-}}{\sqrt{\lambda^+} + \sqrt{\lambda^-}})^{2k} [\frac{(\lambda^+ - \lambda^-)}{2} R^2 + \tilde{R}^2r] \cdot C$
Nesterov accelerated method [Nesterov, 2004]	Worst	$(1 - \frac{2\sqrt{\lambda^-}}{\sqrt{\lambda^+} + \sqrt{\lambda^-}})^k (1 - \frac{\lambda^-}{\lambda^+})^k \cdot C$
	Average	$(1 - \frac{2\sqrt{\lambda^-}}{\sqrt{\lambda^+} + \sqrt{\lambda^-}})^k (1 - \frac{\lambda^-}{\lambda^+})^k [\frac{1}{k^{1/2}} \cdot R^2\lambda^- + \frac{1}{k^{1/2}} \cdot \tilde{R}^2r] \cdot C$

¹ Taylor et al. [2017]

where $\mathbf{A} \in \mathbb{R}^{n \times d}$ is the data matrix, $\tilde{\mathbf{x}} \in \mathbb{R}^d$ is the signal vector, and $\boldsymbol{\eta} \in \mathbb{R}^n$ is a source of noise. All of these inputs will possibly be random and the target $\mathbf{b} = \mathbf{A}\tilde{\mathbf{x}} + \boldsymbol{\eta}$ is produced by a generative model corrupted by noise.

We work in the following setting. Both the number of features (d) and data dimension (n) grow to infinity while d/n tends to a fixed $r \in (0, \infty)$. We use $\tilde{R}^2 = \frac{1}{d}\mathbb{E}[\|\boldsymbol{\eta}\|^2]$ to denote the magnitude of the noise and implicitly define $R^2 \approx \frac{1}{d}\|\mathbf{b}\|^2 - \tilde{R}^2$ to measure the strength of the signal. We make this precise in Assumption 1 of Section 2. Throughout, the signal-to-noise ratio in $[0, \infty]$ is held constant as the problem size grows. Moreover, we assume that the data matrix \mathbf{A} is *independent* of both the signal, $\tilde{\mathbf{x}}$, and noise $\boldsymbol{\eta}$. Note this, together with the generative model, allows for some amount of dependence between \mathbf{A} and the target \mathbf{b} . We will also assume that $\mathbf{H} \stackrel{\text{def}}{=} \frac{1}{n}\mathbf{A}^T\mathbf{A}$ has a well-defined *limiting spectral density*, denoted by $d\mu$, as $n, d \rightarrow \infty$ (made precise in Assumption 2 of Section 2).

Our first contribution is a framework to analyze the average-case complexity of gradient-based methods in the described setting. Our framework highlights how the algorithm, signal and noise levels interact with each other to produce different average-case convergence guarantees. The culmination of this framework is the average-case convergence rates for first-order methods (see Tables 1 and 2).

Our framework is broad enough to facilitate multiple perspectives on average-case analysis. Our motivating and central application is the *fully-average-case*, in which we assume that all inputs are random. The quintessential random data model is *isotropic features*. This supposes the entries of \mathbf{A} are i.i.d. random variables with zero mean, equal variance, and bounded fourth moments, that is, $\mathbb{E}[A_{ij}] = 0, \mathbb{E}[A_{ij}^2] = \sigma^2, \mathbb{E}[A_{ij}^4] < \infty$ for all i, j . In a celebrated theorem of Marčenko and Pastur [1967], the spectrum of $\mathbf{H} = \frac{1}{n}\mathbf{A}^T\mathbf{A}$ converges to a compactly supported measure as the problem size grows *without any further assumptions on the distribution of the entries of \mathbf{A}* . This limiting spectral distribution is known as the Marčenko-Pastur law:

$$d\mu_{\text{MP}}(\lambda) \stackrel{\text{def}}{=} \delta_0(\lambda) \max\{1 - \frac{1}{r}, 0\} + \frac{\sqrt{(\lambda - \lambda^-)(\lambda^+ - \lambda)}}{2\pi\lambda\sigma^2 r} \mathbb{1}_{[\lambda^-, \lambda^+]}, \quad (2)$$

where $\lambda^- \stackrel{\text{def}}{=} \sigma^2(1 - \sqrt{r})^2$ and $\lambda^+ \stackrel{\text{def}}{=} \sigma^2(1 + \sqrt{r})^2$.

However, our framework is built to be vastly more general. To start, the framework covers a fully-average-case analysis with other data models, such as the one-hidden layer network with random weights and the correlated features model (see Section 2.2). More to the point, this framework also allows for a type of semi-average-case analysis, in which only \mathbf{b} is taken to be random. When we do this and then choose \mathbf{A} in such a way as to maximize the halting time, we call this the *adversarial average-case*. See Section 1.2 for further details and motivations.

We now discuss the contents of this framework in detail, which is to say we survey how Assumptions 1 and 2 combine to show the halting time is concentrated and deterministic. The first step is to express the conditional expectation of the gradient at the k -th iterate as a sum of expected traces of polynomials in the matrix $\mathbf{H} = \frac{\mathbf{A}^T\mathbf{A}}{n}$ (c.f. Proposition 4.3):

$$\mathbb{E}[\|\nabla f(\mathbf{x}_k)\|^2 | \mathbf{H}] = \frac{R^2}{d} \text{tr}(\mathbf{H}^2 P_k^2(\mathbf{H})) + \frac{\tilde{R}^2}{n} \text{tr}(\mathbf{H} P_k^2(\mathbf{H})). \quad (3)$$

The polynomial P_k , known as the *residual polynomial*, is a k -th degree polynomial associated with the gradient-based algorithm. This tool of associating each algorithm with polynomials is a classic technique in numerical iterative methods for solving linear systems [Fischer, 1996, Flanders and Shortley, 1950, Golub and Varga, 1961, Rutishauser, 1959]. Such polynomials are used to prove convergence of some of the most celebrated algorithms like the conjugate gradient method [Hestenes and Stiefel, 1952]. Explicit expressions of the residual polynomials for Nesterov’s accelerated methods [Beck and Teboulle, 2009, Nesterov, 2004], both convex and strongly convex, as well as, gradient descent and Polyak’s momentum (a.k.a Heavy-ball) [Polyak, 1964] are derived in Section 3. These polynomials may be of independent interest.

The result in (3) gives an *exact expression* for the expected gradient depending only on traces of powers of \mathbf{H} , which in turn can be expressed in terms of its *eigenvalues*. Our second main assumption (Assumption 2) then ensures that these traces converge to integrals against the spectral density $d\mu$. In summary, the squared gradient norm concentrates to a deterministic quantity, $\overset{12}{\mathcal{E}}_{d \rightarrow \infty} [\|\nabla f(\mathbf{x}_k)\|^2]$:

¹In many situations this deterministic quantity $\mathcal{E}_{d \rightarrow \infty} [\|\nabla f(\mathbf{x}_k)\|^2]$ is in fact the limiting expectation of the squared-norm of the gradient. However, under the assumptions that we are using, this does not immediately follow. It is however always the limit of the median of the squared-norm of the gradient.

²Technically, there is no need to assume the measure μ has a density – the theorem holds just as well for any limiting spectral measure μ . In fact, a version of this theorem can be formulated at finite n just as well, thus dispensing entirely with Assumption 2 – c.f. Proposition 4.3.




Figure 2: The spectrum of matrices $\frac{1}{n}\mathbf{A}^T\mathbf{A}$ under the isotropic features model converges as $n, d \rightarrow \infty$ to the Marčenko-Pastur distribution, shown here for different values of $r = d/n$.

Theorem 1.1 (Concentration of the gradient). *Under Assumptions 1 and 2 the norm of the gradient concentrates around a deterministic value:*

$$\|\nabla f(\mathbf{x}_k)\|^2 \xrightarrow[d \rightarrow \infty]{\Pr} \underbrace{\overbrace{R^2}^{\text{signal}} \int \underbrace{\lambda^2 P_k^2(\lambda)}_{\text{algorithm}} \underbrace{d\mu}_{\text{model}}}_{\text{algorithm}} + \underbrace{\overbrace{\tilde{R}^2}^{\text{noise}} r}_{\text{algorithm}} \int \underbrace{\lambda P_k^2(\lambda)}_{\text{algorithm}} \underbrace{d\mu}_{\text{model}} \stackrel{\text{def}}{=} \mathcal{E}_{d \rightarrow \infty} [\|\nabla f(\mathbf{x}_k)\|^2]. \quad (4)$$

Notably, the deterministic value for which the gradient concentrates around depends only on \mathbf{H} through its eigenvalues.

The concentration of the norm of the gradient above yields a candidate for the limiting value of the halting time, or the first time the gradient $\|\nabla f(\mathbf{x}_k)\|^2$ falls below some predefined ε . We define this candidate for the halting time τ_ε from $\mathcal{E}_{d \rightarrow \infty} [\|\nabla f(\mathbf{x}_k)\|^2]$ and we denote the halting time T_ε , as

$$\tau_\varepsilon \stackrel{\text{def}}{=} \inf \{k > 0 : \mathcal{E}_{d \rightarrow \infty} [\|\nabla f(\mathbf{x}_k)\|^2] \leq \varepsilon\} \quad \text{and} \quad T_\varepsilon \stackrel{\text{def}}{=} \inf \{k > 0 : \|\nabla f(\mathbf{x}_k)\|^2 \leq \varepsilon\}. \quad (5)$$

We note that the deterministic value τ_ε is, by definition, the average complexity of the first-order algorithm. This leads to our second main result that states the almost sure convergence of the halting time to a constant value.

Theorem 1.2 (Halting time universality). *Let ε be such that $\mathcal{E}_{d \rightarrow \infty} [\|\nabla f(\mathbf{x}_k)\|^2] > \varepsilon$, $\|\nabla f(\mathbf{x}_0)\|^2 > \varepsilon$ (at least one step is performed), and $\mathcal{E}_{d \rightarrow \infty} [\|\nabla f(\mathbf{x}_k)\|^2] \neq \varepsilon$ for all k . Under Assumptions 1 and 2,*

$$\lim_{d \rightarrow \infty} \Pr(T_\varepsilon = \tau_\varepsilon) = 1. \quad (6)$$

1.2 Comparison between average and worst-case

The average-case analysis we develop in this paper is effective in the large problem size limit, whereas worst-case analysis is performed for a fixed matrix size. This implies that there are potentially *dimension-dependent* quantities which must be addressed when making a comparison.

For example, all the first-order methods considered here converge linearly for the finite-dimensional least squares problem: the rate is determined by the gap between the smallest nonzero eigenvalue of the matrix \mathbf{H} and 0. However this could very well be meaningless in the context of a high-dimensional problem, as this gap becomes vanishingly small as the problem size grows.

In the context of the isotropic features model, when the ratio of features to samples r is 1, this is precisely what occurs: the smallest eigenvalues tend to 0 as the matrix size grows. In contrast, when r is bounded away from 1, the least squares problem in (1) has an almost sure *dimension-independent* lower bound on the Hessian, (c.f. Figure 2). However, for the comparison we do here, there is another dimension-dependent quantity which will have a greater impact on the worst-case bounds.

Worst-case rates and the distance to optimality. Typical worst-case upper bounds for first-order algorithms depend on the distance to optimality, $\|\mathbf{x}_0 - \mathbf{x}^*\|^2$. For example, let us consider gradient descent (GD). Tight worst-case bounds for GD in the strongly convex and convex setting [Taylor et al., 2017], respectively, are

$$\begin{aligned} \|\nabla f(\mathbf{x}_k)\|^2 &\leq (\lambda_{\mathbf{H}}^+)^2 \|\mathbf{x}_0 - \mathbf{x}^*\|^2 \left(1 - \frac{\lambda_{\mathbf{H}}^-}{\lambda_{\mathbf{H}}^+}\right)^{2k} \stackrel{\text{def}}{=} \text{UB}_{\text{sc}}(\|\nabla f(\mathbf{x}_k)\|^2) \\ \text{and} \quad \|\nabla f(\mathbf{x}_k)\|^2 &\leq \frac{(\lambda_{\mathbf{H}}^+)^2 \|\mathbf{x}_0 - \mathbf{x}^*\|^2}{(k+1)^2} \stackrel{\text{def}}{=} \text{UB}_{\text{cvx}}(\|\nabla f(\mathbf{x}_k)\|^2), \end{aligned}$$

where \mathbf{x}^* is the solution to (1) found by the algorithm.

To formulate a comparison between the fully-average-case rates, where \mathbf{A} follows isotropic features, and the worst-case rates, we must make an estimate of this distance to optimality $\|\mathbf{x}_0 - \mathbf{x}^*\|^2$. In the noiseless setting ($\tilde{R} = 0$), the expectation of $\|\mathbf{x}_0 - \mathbf{x}^*\|^2$ is a constant multiple of R^2 . In particular, it is independent of the dimension. Similarly when we have dimension-independent-strong-convexity, ($r \neq 1$), even with noisy targets \mathbf{b} ($\tilde{R} > 0$), the distance to the optimum is well-behaved and $\mathbb{E}[\|\mathbf{x}_0 - \mathbf{x}^*\|^2]$ is a constant involving \tilde{R}^2 and R^2 . Hence, a direct comparison between worst and average-case is relatively simple.

For the ill-conditioned case when $r = 1$, the situation is more complicated with noisy targets. To maintain a fixed and finite signal-to-noise ratio, the distance to optimality will behave like $\|\mathbf{x}^* - \mathbf{x}_0\|^2 \approx d\tilde{R}^2$; that is, it is dimension-dependent.³ So the worst-case rates have a dimension-dependent constant whereas the average-case rates are dimension-independent. This dimension-dependent term can be seen in the last column of Table 1. Conversely, if one desires to make $\mathbb{E}[\|\mathbf{x}_0 - \mathbf{x}^*\|^2]$ constant across dimensions using a generative model with noise, one is forced to scale $\boldsymbol{\eta}$ to go to zero as $d \rightarrow \infty$, thus reducing the full generative model to the noiseless regime.

Adversarial model. As mentioned above, the comparison with existing worst-case bounds is problematic due to dimension-dependent factors. To overcome this, we consider the following *adversarial model*. First, we assume a noisy generative model for \mathbf{b} (Assumption 1 holds). Next, our adversary chooses the matrix \mathbf{A} without knowledge of \mathbf{b} to *maximize the norm of the gradient* subject to the constraint that the convex hull of the eigenvalues of $\mathbf{H} = \frac{1}{n}\mathbf{A}^T\mathbf{A}$ equals $[\lambda^-, \lambda^+]$. For comparison to the average-case analysis with isotropic features, we would choose λ^\pm to be the endpoints of the Marcenko-Pastur law. In light of Theorem 1.1, the adversarial model seeks to solve the constrained optimization problem

$$\lim_{d \rightarrow \infty} \max_{\mathbf{H}} \mathbb{E}[\|\nabla f(\mathbf{x}_k)\|^2] = \max_{\lambda \in [\lambda^-, \lambda^+]} \{R^2\lambda^2 P_k^2(\lambda) + \tilde{R}^2 r \lambda P_k^2(\lambda)\}. \quad (7)$$

We call this expression the *adversarial average-case guarantee*.

The main distinction between worst-case and adversarial average-case is that traditional worst-case maximizes the gradient over *all* inputs – both targets \mathbf{b} and data matrix \mathbf{A} . This leads to dimension-dependent complexity, as there are usually exceptional target vectors that are heavily dependent on the data matrix \mathbf{A} (such as those built from extremal singular vectors of \mathbf{A}) and cause the algorithm to perform exceptionally slowly.

In contrast, the adversarial average-case keeps the randomness of the target \mathbf{b} while maximizing over the data matrix \mathbf{A} . This is a more meaningful worst-case comparison: for example, in the setting of linear regression, the response and measurements of the independent variables are typically generated through different means and have different and independent sources of noise (see for example [Walpole and Myers, 1989, Example 10.1]). Hence the independence of the noise intervenes to limit how truly bad the data matrix \mathbf{A} can be. Furthermore, the complexity of the adversarial average-case is dimension-independent. Table 1 shows these adversarial complexities for non-strongly convex objectives (1). Similar results can also be derived for strongly convex objectives but are omitted for brevity.

Comparison with adversarial and worst-case complexities. By construction, the average-case convergence rates are at least as good as the worst-case and adversarial guarantees. The average-case complexity in the convex, noiseless setting ($r = 1, \tilde{R} = 0$) for Nesterov’s accelerated method (convex) [Beck and Teboulle, 2009, Nesterov, 2004] and gradient descent (GD) are an order of magnitude faster in k than the worst case rates (see Table 1, first column). It may appear at first glance (Table 1) that there is a discrepancy between the average-case and exact worst-case

³Precisely, we show that $\frac{d\tilde{R}^2}{\|\mathbf{x}^* - \mathbf{x}_0\|^2}$ is tight (see Section 5, Lemma 5.1).




Figure 3: **Average-case vs worst-case in least squares** with isotropic features ($r = 1, d = 4096$). **Left:** 8000 runs of GD, standard deviation (shaded region, undetectable), and theoretical rates (dashed lines). Empirical runs precisely match the theoretical average-case rates. **Right:** Ratio of the upper bound in worst-case to empirical gradient after $k = 4096$ iterations, $\text{UB}_{\text{cvx}}(\|\nabla f(\mathbf{x}_k)\|^2)/\|\nabla f(\mathbf{x}_k)\|^2$. From the concentration of gradients (left), this implies that *the norms of the gradient for worst-cases are always larger than average*. The distribution of worst-case gradient rates with noise has large variance contrasting the little variance in (left) and makes the expected worst-case unpredictable in contrast with the noiseless.

rate in the noisy, $r = 1$. As noted in the previous section, the worst-case rates have dimension-dependent constants. Provided the dimension is bigger than the iteration counter ($d \geq k^{1/2}$ for GD and $d \geq \log(k)$ for Nesterov), the average complexity indeed yields a faster rate of convergence. Average-case is always strictly better than adversarial rates (see Table 1). This improvement in the average rate indeed highlights that *the support of the spectrum does not fully determine the rate*. Many eigenvalues contribute meaningfully to the average rate. Hence, our results are not and cannot be purely explained by the support of the spectrum.

The average-case complexity in the strongly convex case matches the worst-case guarantees multiplied by an additional *polynomial correction term* (green in Table 2). This polynomial term has little effect on the complexity compared to the linear rate. However as the matrix \mathbf{H} becomes ill-conditioned ($r \rightarrow 1$), the polynomial correction starts to dominate the average-case complexity. The sublinear rates in Table 1 show this effect and it accounts for the improved average-case rates.

Our average-case rates accurately predict the empirical convergence observed in simulations, in contrast to the worst-case rates (see Figure 3). Although our rates only hold on average, surprisingly, even a single instance of GD exactly matches the theoretical predictions. Moreover, the noisy non-strongly convex worst-case is highly unpredictable due to the instability in \mathbf{x}^* across runs. As such, the worst-case analysis is not representative of typical behavior (see Figure 3).

These theoretical results are supported by simulations and empirically extended to other models, such as logistic regression, as well as other algorithms, such as stochastic gradient descent (SGD) (see Section 6). This suggests that this universality property holds for a wider class of problems.

Related work. The average-case analysis has a long history in computer science and numerical analysis. Often it is used to justify the superior performance of algorithms such as Quicksort (sorting) [Hoare, 1962] and the simplex method in linear programming, see for example [Borgwardt, 1986, Smale, 1983, Spielman and Teng, 2004] and references therein. Despite this rich history, it is challenging to transfer these ideas into continuous optimization due to the ill-defined notion of a typical continuous optimization problem. Recently Lacotte and Pilanci [2020], Pedregosa and Scieur [2020] derived a framework for average-case analysis of gradient-based methods and developed optimal algorithms with respect to the average-case. The class of problems they consider is a special case of

(1) with vanishing noise. We use a similar framework – extending the results to all first-order methods and noisy quadratics while also providing concentration and explicit average-case convergence guarantees.

A natural criticism of a simple average-case analysis is that the complexity is data model dependent and thus it only has predictive power for a small subset of real world phenomena. Because of this, it becomes important to show that any modeling choices made in defining the data ensemble have limited effect. Paquette and Trogdon [2020] showed that the halting time for conjugate gradient becomes deterministic as the dimension grows and it exhibits a universality property, that is, for the class of Wishart matrices, the halting times are identical (see also Deift and Trogdon [2020]). It is conjectured that this property holds in greater generality – for more distributions and more algorithms [Deift and Trogdon, 2018, Deift et al., 2014]). In Sagun et al. [2017], empirical evidence confirms this for neural networks and spin glass models. Our paper is in the same spirit as these – definitively answering the question that all first-order methods share this universality property for the halting time on quadratic problems.

Our work draws heavily upon classical polynomial based iterative methods. Originally designed for the Chebyshev iterative method [Flanders and Shortley, 1950, Golub and Varga, 1961], the polynomial approach for analyzing algorithms was instrumental in proving worst-case complexity for the celebrated conjugate gradient method [Hestenes and Stiefel, 1952]. For us, the polynomial approach gives an explicit equation relating the eigenvalues of the data matrix to the iterates which, in turn, allows the application of random matrix theory.

Notation. We denote vectors in lowercase boldface (\mathbf{x}) and matrices in uppercase boldface (\mathbf{H}). The norm $\|\mathbf{x}\|_2^2 = \mathbf{x}^T \mathbf{x}$ gives the usual Euclidean 2-norm and $\|\mathbf{H}\|_{\text{op}} = \max$ singular value of \mathbf{H} is the usual operator-2 norm. Given a matrix $\mathbf{H} \in \mathbb{R}^{d \times d}$, the largest eigenvalue of \mathbf{H} is $\lambda_{\mathbf{H}}^+$ and its smallest eigenvalue is $\lambda_{\mathbf{H}}^-$. A sequence of random variables $\{y_d\}_{d=0}^{\infty}$ converges in probability to y , indicated by $y_d \xrightarrow[\text{Pr}]{d \rightarrow \infty} y$, if for any $\varepsilon > 0$, $\lim_{d \rightarrow \infty} \Pr(|y_d - y| > \varepsilon) = 0$. In other words, the probability that y_d is far from y goes to 0 as d increases. Probability measures are denoted by μ and their densities by $d\mu$. We say a sequence of random measures μ_d converges to μ weakly in probability if for any bounded continuous function f , we have $\int f d\mu_d \rightarrow \int f d\mu$ in probability.

All stochastic quantities defined hereafter live on a probability space denoted by $(\text{Pr}, \Omega, \mathcal{F})$ with probability measure Pr and the σ -algebra \mathcal{F} containing subsets of Ω . A random variable (vector) is a measurable map from Ω to \mathbb{R} (\mathbb{R}^d) respectively. Let $X : (\Omega, \mathcal{F}) \rightarrow (\mathbb{R}, \mathcal{B})$ be a random variable mapping into the Borel σ -algebra \mathcal{B} and the set $B \in \mathcal{B}$. We use the standard shorthand for the event $\{X \in B\} = \{\omega : X(\omega) \in B\}$.

The remainder of the article is structured as follows: in Section 2 we introduce the full mathematical model under investigation including some examples of data models. Section 3 discusses the relationship between polynomials and optimization. Our main results are then described and proven in Section 4. Section 5 details the computations involved in the average-case analysis, the proofs of which are deferred to the appendix. The article concludes on showing some numerical simulations in Section 6.

2 Problem setting

In this paper, we develop an average-case analysis for first-order methods on quadratic problems of the form

$$\arg \min_{\mathbf{x} \in \mathbb{R}^d} \left\{ f(\mathbf{x}) \stackrel{\text{def}}{=} \frac{1}{2n} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2 \right\}, \quad \text{with } \mathbf{b} \stackrel{\text{def}}{=} \mathbf{A}\tilde{\mathbf{x}} + \boldsymbol{\eta}, \quad (8)$$

where $\mathbf{A} \in \mathbb{R}^{n \times d}$ is a (possibly random) matrix (discussed in the next subsection), $\tilde{\mathbf{x}} \in \mathbb{R}^d$ is an unobserved signal vector, and $\boldsymbol{\eta} \in \mathbb{R}^n$ is a noise vector.

2.1 Data matrix, noise, signal, and initialization assumptions

Throughout the paper we make the following assumptions.

Assumption 1 (Initialization, signal, and noise). *The initial vector $\mathbf{x}_0 \in \mathbb{R}^d$, the signal $\tilde{\mathbf{x}} \in \mathbb{R}^d$, and noise vector $\boldsymbol{\eta} \in \mathbb{R}^n$ satisfy the following conditions:*

1. *The entries of $\mathbf{x}_0 - \tilde{\mathbf{x}}$ are i.i.d. random variables and there exist constants $C, R > 0$ such that for $i = 1, \dots, d$*

$$\mathbb{E}[\mathbf{x}_0 - \tilde{\mathbf{x}}] = \mathbf{0}, \quad \mathbb{E}[(\mathbf{x}_0 - \tilde{\mathbf{x}})_i^2] = \frac{1}{d}R^2, \quad \text{and} \quad \mathbb{E}[(\tilde{\mathbf{x}} - \mathbf{x}_0)_i^4] \leq \frac{1}{d^2}C. \quad (9)$$

2. *The entries of the noise vector $\boldsymbol{\eta}$ are i.i.d. random variables that verify the following for $i = 1, \dots, d$ and for some constants $\tilde{C}, \tilde{R} > 0$*

$$\mathbb{E}[\boldsymbol{\eta}] = \mathbf{0}, \quad \mathbb{E}[\eta_i^2] = \tilde{R}^2, \quad \text{and} \quad \mathbb{E}[\eta_i^4] \leq \tilde{C}. \quad (10)$$

Assumption 1 encompasses the setting where the signal is *random* and the algorithm is initialized at $\mathbf{x}_0 = \mathbf{0}$. But it is more general. Starting farther from the signal requires more iterations to converge. Hence, intuitively, equation 9 restricts the distance of the algorithm's initialization to the signal so that it remains constant across problem sizes. The unbiased initialization about the signal, put another way, says the initialization is rotationally-invariantly distributed about the signal $\tilde{\mathbf{x}}$ (see Figure 4).

Assumption 1 arises as a result of preserving a constant signal-to-noise ratio in the generative model. Such generative models with this scaling have been used in numerous works [Hastie et al., 2019, Mei and Montanari, 2019].

Tools from random matrix theory. Random matrix theory studies properties of matrices \mathbf{H} (most notably, statistics of matrix eigenvalues) whose entries H_{ij} are random variables. We refer the reader to [Bai and Silverstein, 2010, Tao, 2012] for a more thorough introduction. Many important statistics of random matrix theory can be expressed as functionals on the eigenvalues of a matrix \mathbf{H} (e.g., determinants and traces). Let $\lambda_1, \dots, \lambda_d$ be the eigenvalues of \mathbf{H} and define the *empirical spectral measure* (ESM), $\mu_{\mathbf{H}}$, as

$$\mu_{\mathbf{H}}(\lambda) \stackrel{\text{def}}{=} \frac{1}{d} \sum_{i=1}^d \delta_{\lambda_i}, \quad (11)$$

where δ_{λ_i} is a Dirac delta function, *i.e.*, a function equal to 0 except at λ_i and whose integral over the entire real line is equal to one. The empirical spectral measure puts a uniform weight on each of the eigenvalues of \mathbf{H} . When \mathbf{H} is random, this becomes a random measure. A main interest in random matrix theory is to characterize the behavior of the empirical spectral measure as the dimension of the matrix tends to infinity.

Because the ESM is a well-studied object for many random matrix ensembles, we state the following assumption on the ESM for the data matrix, \mathbf{A} . In Section 2.2, we review practical scenarios in which this is verified.




Figure 4: The pictured \mathbf{x}_0 are equiprobable. Each colored line is a different run of GD with random matrix \mathbf{A} and the shaded gray area is the set where $\|\nabla f(\mathbf{x})\|^2 < \varepsilon$. Intuitively, our result says all runs of GD starting from a random \mathbf{x}_0 take the same number of iterations to reach the shaded area.

Assumption 2 (Data matrix). *Let \mathbf{A} be a (possibly random) $n \times d$ matrix such that the number of features, d , tends to infinity proportionally to the size of the data set, n , so that $\frac{d}{n} \rightarrow r \in (0, \infty)$. Let $\mathbf{H} \stackrel{\text{def}}{=} \frac{1}{n} \mathbf{A}^T \mathbf{A}$ with eigenvalues $\lambda_1 \leq \dots \leq \lambda_d$ and let δ_{λ_i} denote the Dirac delta with mass at λ_i . We make the following assumptions on the eigenvalue distribution of this matrix:*

1. *The ESM converges weakly in probability to a deterministic measure μ with compact support,*

$$\mu_{\mathbf{H}} = \frac{1}{d} \sum_{i=1}^d \delta_{\lambda_i} \rightarrow \mu \quad \text{weakly in probability.} \quad (12)$$

2. *The largest eigenvalue of \mathbf{H} converges in probability to the largest element in the support of μ . In particular, if λ^+ denotes the top edge of the support of μ then*

$$\lambda_{\mathbf{H}}^+ \xrightarrow[d \rightarrow \infty]{\text{Pr}} \lambda^+. \quad (13)$$

3. *(Required provided the algorithm uses the smallest eigenvalue) The smallest eigenvalue of \mathbf{H} converges in probability to the smallest, non-zero element in the support of μ . In particular, if λ^- denotes the bottom edge of the support of μ then*

$$\lambda_{\mathbf{H}}^- \xrightarrow[d \rightarrow \infty]{\text{Pr}} \lambda^-. \quad (14)$$

2.2 Examples of data distributions.

In this section we review three examples of data-generating distributions that verify Assumption 2: a model with isotropic features, a correlated features model, and a one-hidden layer neural network with random weights. Numerous works studying the spectrum of the Hessian on neural networks have found that this spectrum shares many characteristics with the limiting spectral distributions discussed below including compact support, a concentration of eigenvalues near 0, and a stable top eigenvalue [Behrooz et al., 2019, Dauphin et al., 2014, Pappayan, 2018, Sagun et al., 2016]. In fact, the work of Martin and Mahoney [2018] directly compares the Hessians of deep neural networks at various stages in training with the Marčenko-Pastur density, that is, the limiting spectral density for the isotropic features model.

Isotropic features. We will now elaborate on the well developed theory surrounding the isotropic features model (see (2) and the text just above it). In particular, parts 2 and 3 of Assumption 2 on the convergence of the largest and smallest eigenvalues is well known:

Lemma 2.1 (Isotropic features). (**Bai and Silverstein [2010, Theorem 5.8]**) *Suppose the matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ is generated using the isotropic features model. The largest and smallest eigenvalue of \mathbf{H} , $\lambda_{\mathbf{H}}^+$ and $\lambda_{\mathbf{H}}^-$, resp., converge in probability to λ^+ and λ^- resp. where $\lambda^+ = \sigma^2(1 + \sqrt{r})^2$ is the top edge of the support of the Marčenko-Pastur measure and $\lambda^- = \sigma^2(1 - \sqrt{r})^2$ is the bottom edge of the support of the Marčenko-Pastur measure.*

In addition, the isotropic features model is sufficiently random that it is possible to weaken Assumption 1 and still derive for it the conclusion of Theorem 1.1. In particular, we may let \mathbf{b}_n be defined as

$$\mathbf{b}_n = \frac{1}{\sqrt{n}} R \mathbf{A} \boldsymbol{\omega}_{1,d} + \tilde{R} \boldsymbol{\omega}_{2,n} \quad (15)$$

for any deterministic sequences of vectors $\{\boldsymbol{\omega}_{1,d}\}$ and $\{\boldsymbol{\omega}_{2,n}\}$ from the d -dimensional and the n -dimensional spheres, respectively, multiplied by the signal strength R and noise \tilde{R} . Then, under the further moment assumption on \mathbf{A} that for any $k \in \mathbb{N}$

$$\sup_{i,j} \left\{ \mathbb{E} |A_{i,j}|^k \right\} < \infty, \quad (16)$$

it is a consequence of [Knowles and Yin, 2017, Theorem 3.6,3.7], that

$$\|\nabla f(\mathbf{x}_k)\|^2 \xrightarrow{d \rightarrow \infty} R^2 \int \lambda^2 P_k^2(\lambda) d\mu + \tilde{R}^2 r \int \lambda P_k^2(\lambda) d\mu = \mathcal{E} [\|\nabla f(\mathbf{x}_k)\|^2]. \quad (17)$$

This implies that for the isotropic features model under the stronger (16), but the weaker target assumption (15), we obtain the same complexity results presented in Tables 1 and 2. See also [Paquette and Trogon, 2020, Corollary 5.12] in which a central limit theorem for the gradient is derived under these same assumptions

Correlated features. In this model, one takes a random matrix $\mathbf{W} \in \mathbb{R}^{n \times d}$ generated from the isotropic features model and a symmetric positive definite correlation matrix $\Sigma_n \in \mathbb{R}^{n \times n}$. One then defines the correlated features model by

$$\mathbf{A} \stackrel{\text{def}}{=} \Sigma_n^{1/2} \mathbf{W}.$$

This makes $\mathbf{H} = \frac{1}{n} \mathbf{A}^T \mathbf{A}$ the normalized sample covariance matrix of d samples of a n -dimensional random vector with covariance structure Σ_n .

Under the assumption that the empirical spectral measure of Σ_n converges to a measure ν and that the norm of Σ_n is uniformly bounded, it is consequence of Bai and Silverstein [1998, 1999] (see also the discussions in Bai and Silverstein [2004], Hachem et al. [2016], Knowles and Yin [2017]) that Assumption 2 holds. Unlike in isotropic features, the limiting spectral measure is not known explicitly, but is instead only characterized (in general) through a fixed-point equation describing its Stieltjes transform.

One-hidden layer network with random weights. In this model, the entries of \mathbf{A} are the result of a matrix multiplication composed with a (potentially non-linear) activation function $g : \mathbb{R} \rightarrow \mathbb{R}$:

$$A_{ij} \stackrel{\text{def}}{=} g\left(\frac{[\mathbf{W}\mathbf{Y}]_{ij}}{\sqrt{m}}\right), \quad \text{where } \mathbf{W} \in \mathbb{R}^{n \times m}, \mathbf{Y} \in \mathbb{R}^{m \times d} \text{ are random matrices.} \quad (18)$$

The entries of \mathbf{W} and \mathbf{Y} are i.i.d. with zero mean, isotropic variances $\mathbb{E}[W_{ij}^2] = \sigma_w^2$ and $\mathbb{E}[Y_{ij}^2] = \sigma_y^2$, and light tails, that is, there exists constants $\theta_w, \theta_y > 0$ and $\alpha > 0$ such that for any $t > 0$

$$\Pr(|W_{11}| > t) \leq \exp(-\theta_w t^\alpha) \quad \text{and} \quad \Pr(|Y_{11}| > t) \leq \exp(-\theta_y t^\alpha). \quad (19)$$

Although stronger than bounded fourth moments, this assumption holds for any sub-Gaussian random variables (*e.g.*, Gaussian, Bernoulli, etc). As in the previous case to study the large dimensional limit, we assume that the different dimensions grow at comparable rates given by $\frac{m}{n} \rightarrow r_1 \in (0, \infty)$ and $\frac{m}{d} \rightarrow r_2 \in (0, \infty)$. This model encompasses two-layer neural networks with a squared loss, where the first layer has random weights and the second layer's weights are given by the regression coefficients \mathbf{x} . In this case, problem (8) becomes

$$\min_{\mathbf{x}} \left\{ f(\mathbf{x}) = \frac{1}{2n} \left\| g\left(\frac{1}{\sqrt{m}} \mathbf{W}\mathbf{Y}\right) \mathbf{x} - \mathbf{b} \right\|_2^2 \right\}. \quad (20)$$

The model was introduced by [Rahimi and Recht, 2008] as a randomized approach for scaling kernel methods to large datasets, and has seen a surge in interest in recent years as a way to study the generalization properties of neural networks [Hastie et al., 2019, Liao and Couillet, 2018, Louart et al., 2018, Mei and Montanari, 2019, Pennington and Worah, 2017].

The most important difference between this model and the isotropic features is the existence of a potentially non-linear activation function g . We assume g to be entire with a growth condition and have zero Gaussian-mean,

$$\text{(Gaussian mean)} \quad \int g(\sigma_w \sigma_y z) \frac{e^{-z^2/2}}{\sqrt{2\pi}} dz = 0. \quad (21)$$

Table 3: **Residual Polynomials.** Summary of the residual polynomials associated with the methods discussed in this paper. T_k is the k -th Chebyshev polynomial of the first kind, U_k is k -th Chebyshev polynomial of the second kind, and L_k is the k -th Legendre polynomial. Derivations of these polynomials can be found in Appendix A. In light of Proposition 3.1, an explicit expression for the polynomial P_k is enough to determine the polynomial Q_k .

Methods	Polynomial P_k	Parameters
Gradient Descent	$(1 - \alpha\lambda)^k$	$\alpha = 1/\lambda_{\mathbf{H}}^+$
Nesterov (cvx) ¹	$\frac{2(1-\alpha\lambda)^{(k+1)/2}}{\alpha\lambda k} (\sqrt{1-\alpha\lambda}L_k(\sqrt{1-\alpha\lambda}) - L_{k+1}(\sqrt{1-\alpha\lambda}))$	$\alpha = 1/\lambda_{\mathbf{H}}^+$
Polyak ²	$\beta^k \left[\frac{(\sqrt{\lambda_{\mathbf{H}}^+} - \sqrt{\lambda_{\mathbf{H}}^-})^2}{\lambda_{\mathbf{H}}^+ + \lambda_{\mathbf{H}}^-} \cdot T_k(\sigma(\lambda)) + \frac{2\sqrt{\lambda_{\mathbf{H}}^-} \lambda_{\mathbf{H}}^+}{\lambda_{\mathbf{H}}^+ + \lambda_{\mathbf{H}}^-} \cdot U_k(\sigma(\lambda)) \right]$	$\sigma(\lambda) = \frac{\lambda_{\mathbf{H}}^+ + \lambda_{\mathbf{H}}^- - 2\lambda}{\lambda_{\mathbf{H}}^+ - \lambda_{\mathbf{H}}^-}$ $\beta = \frac{\sqrt{\lambda_{\mathbf{H}}^+} - \sqrt{\lambda_{\mathbf{H}}^-}}{\sqrt{\lambda_{\mathbf{H}}^+} + \sqrt{\lambda_{\mathbf{H}}^-}}$
Nesterov (Strongly cvx) ³	$\frac{2\beta(\beta x)^{k/2}}{1+\beta} T_k\left(\frac{1+\beta}{2\sqrt{\beta}}\sqrt{x}\right) + \left(1 - \frac{2\beta}{1+\beta}\right) (\beta x)^{k/2} U_k\left(\frac{1+\beta}{2\sqrt{\beta}}\sqrt{x}\right)$	$x = 1 - \alpha\lambda,$ $\alpha = 1/\lambda_{\mathbf{H}}^+$ $\beta = \frac{\sqrt{\lambda_{\mathbf{H}}^+} - \sqrt{\lambda_{\mathbf{H}}^-}}{\sqrt{\lambda_{\mathbf{H}}^+} + \sqrt{\lambda_{\mathbf{H}}^-}}$

¹ [Beck and Teboulle, 2009, Nesterov, 2004] ² [Polyak, 1964] ³ [Nesterov, 2004]

The additional growth condition on the function g is precisely given as there exists positive constants $C_g, c_g, A_0 > 0$ such that for any $A \geq A_0$ and any $n \in \mathbb{N}$

$$\sup_{z \in [-A, A]} |g^{(n)}(z)| \leq C_g A^{c_g n}. \quad (22)$$

This growth condition is verified for common activation functions such as the sigmoid $g(z) = (1 + e^{-z})^{-1}$ and the softplus $g(z) = \log(1 + e^z)$, a smoothed approximation to the ReLU. The Gaussian mean assumption (21) can always be satisfied by incorporating a translation into the activation function.

Benigni and Pécché [2019] recently showed that the empirical spectral measure and largest eigenvalue of \mathbf{H} converge to a deterministic measure and largest element in the support, respectively. This implies that this model, like the isotropic features one, verifies Assumption 2. However, contrary to the isotropic features model, the limiting measure does not have an explicit expression, except for some specific instances of g in which it is known to coincide with the Marčenko-Pastur distribution.

Lemma 2.2 (One-hidden layer network). (**Benigni and Pécché [2019, Theorems 2.2 and 5.1]**) *Suppose the matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ is generated using the random features model. Then there exists a deterministic compactly supported measure μ such that $\mu_{\mathbf{H}} \xrightarrow{d \rightarrow \infty} \mu$ weakly in probability. Moreover*

$\lambda_{\mathbf{H}}^+ \xrightarrow{d \rightarrow \infty} \lambda^+$ where λ^+ is the top edge of the support of μ .

3 From optimization to polynomials

In this section, we look at the classical connection between optimization algorithms, iterative methods, and polynomials [Fischer, 1996, Flanders and Shortley, 1950, Golub and Varga, 1961, Rutishauser,

1959]. While the idea of analyzing optimization algorithms from the perspective of polynomials is well-established, many modern algorithms, such as the celebrated Nesterov accelerated gradient [Nesterov, 2004], use alternative approaches to prove convergence.

This connection between polynomials and optimization methods will be crucial to proving the average-case guarantees in Tables 1 and 2. To exploit this connection, we will construct the residual polynomials associated with the considered methods and we prove novel facts which may be of independent interest. For example, the polynomials associated with Nesterov’s method provides an alternative explanation to the ODE in [Su et al., 2016].

Throughout the paper, we consider only gradient-based methods, algorithms which can be written as a linear combination of the previous gradients and the initial iterate.

Definition 3.1 (Gradient-based method). An optimization algorithm is called a *gradient-based method* if each update of the algorithm can be written as a linear combination of the previous iterate and previous gradients. In other words, if every update is of the form

$$\mathbf{x}_{k+1} = \mathbf{x}_0 + \sum_{i=0}^k c_{ki} \nabla f(\mathbf{x}_i), \quad (23)$$

for some scalar values c_{ki} that can potentially depend continuously on $\lambda_{\mathbf{H}}^+$ and $\lambda_{\mathbf{H}}^-$.

Examples of gradient-based methods include momentum methods [Polyak, 1964], accelerated methods [Beck and Teboulle, 2009, Nesterov, 2004], and gradient descent. Now given any gradient-based method, we can associate to the method *residual polynomials* P_k and *iteration polynomials* Q_k which are polynomials of degree k , precisely as followed.

Proposition 3.1 (Polynomials and gradient-based methods). *Consider a gradient-based method with coefficients that depend continuously on $\lambda_{\mathbf{H}}^-$ and $\lambda_{\mathbf{H}}^+$. Define the sequence of polynomials $\{P_k, Q_k\}_{k=0}^{\infty}$ recursively by*

$$\begin{aligned} P_0(\mathbf{H}; \lambda_{\mathbf{H}}^{\pm}) &= \mathbf{I} \quad \text{and} \quad P_k(\mathbf{H}; \lambda_{\mathbf{H}}^{\pm}) = \mathbf{I} - \mathbf{H}Q_k(\mathbf{H}) \\ Q_0(\mathbf{H}; \lambda_{\mathbf{H}}^{\pm}) &= \mathbf{0} \quad \text{and} \quad Q_k(\mathbf{H}; \lambda_{\mathbf{H}}^{\pm}) = \sum_{i=0}^{k-1} c_{k-1,i} [\mathbf{H}Q_i(\mathbf{H}; \lambda_{\mathbf{H}}^{\pm}) - \mathbf{I}]. \end{aligned} \quad (24)$$

These polynomials P_k and Q_k are referred to as the *residual* and *iteration polynomials* respectively. We express the difference between the iterate at step k and $\tilde{\mathbf{x}}$ in terms of these polynomials:

$$\mathbf{x}_k - \tilde{\mathbf{x}} = P_k(\mathbf{H}; \lambda_{\mathbf{H}}^{\pm})(\mathbf{x}_0 - \tilde{\mathbf{x}}) + Q_k(\mathbf{H}; \lambda_{\mathbf{H}}^{\pm}) \cdot \frac{\mathbf{A}^T \boldsymbol{\eta}}{n}. \quad (25)$$

Proof. We will prove the result by induction. For $k = 0$, the claimed result holds trivially. We assume it holds up to iteration k and we will prove it holds for $k + 1$. To show this, we will use the following equivalent form of the gradient $\nabla f(\mathbf{x}) = \mathbf{H}(\mathbf{x} - \tilde{\mathbf{x}}) - \frac{\mathbf{A}^T \boldsymbol{\eta}}{n}$, which follows from the definition of \mathbf{b} . Using this and the definition of gradient-based method, we have:

$$\begin{aligned} \mathbf{x}_{k+1} - \tilde{\mathbf{x}} &= \mathbf{x}_0 - \tilde{\mathbf{x}} + \sum_{i=0}^k c_{ki} \nabla f(\mathbf{x}_i) = \mathbf{x}_0 - \tilde{\mathbf{x}} + \sum_{i=0}^k c_{ki} [\mathbf{H}(\mathbf{x}_i - \tilde{\mathbf{x}}) - \frac{\mathbf{A}^T \boldsymbol{\eta}}{n}] \\ &= \mathbf{x}_0 - \tilde{\mathbf{x}} + \sum_{i=0}^k c_{ki} \mathbf{H} [(\mathbf{I} - \mathbf{H}Q_i(\mathbf{H}; \lambda_{\mathbf{H}}^{\pm}))(\mathbf{x}_0 - \tilde{\mathbf{x}}) + Q_i(\mathbf{H}; \lambda_{\mathbf{H}}^{\pm}) \frac{\mathbf{A}^T \boldsymbol{\eta}}{n} - \frac{\mathbf{A}^T \boldsymbol{\eta}}{n}] \\ &= \mathbf{x}_0 - \tilde{\mathbf{x}} + \mathbf{H} \sum_{i=0}^k c_{ki} (\mathbf{I} - \mathbf{H}Q_i(\mathbf{H}; \lambda_{\mathbf{H}}^{\pm}))(\mathbf{x}_0 - \tilde{\mathbf{x}}) + \sum_{i=0}^k c_{ki} (\mathbf{H}Q_i(\mathbf{H}; \lambda_{\mathbf{H}}^{\pm}) - \mathbf{I}) \frac{\mathbf{A}^T \boldsymbol{\eta}}{n} \\ &= \underbrace{[\mathbf{I} - \mathbf{H}Q_{k+1}(\mathbf{H}; \lambda_{\mathbf{H}}^{\pm})]}_{=P_{k+1}(\mathbf{H}, \lambda_{\mathbf{H}}^{\pm})} (\mathbf{x}_0 - \tilde{\mathbf{x}}) + Q_{k+1}(\mathbf{H}; \lambda_{\mathbf{H}}^{\pm}) \frac{\mathbf{A}^T \boldsymbol{\eta}}{n}, \end{aligned}$$

where in the second identity we have used the induction hypothesis and in the last one the recursive definition of Q_{k+1} . \square

3.1 Examples of residual polynomials.

Motivated by the identity linking the error and the residual polynomial in Proposition 3.1, we derive the residual polynomials for some well-known optimization methods. Some of these residual polynomials are known but some, like Nesterov's accelerated methods, appear to be new.

Gradient descent. Due to the simplicity in the recurrence of the iterates for gradient descent, its residual polynomials P_k and Q_k are explicit. Take for example the typical step size $\alpha = \frac{1}{\lambda_{\mathbf{H}}^+}$. Then iterates on (8) follow the recursion

$$\mathbf{x}_k - \tilde{\mathbf{x}} = \mathbf{x}_{k-1} - \tilde{\mathbf{x}} - \alpha \nabla f(\mathbf{x}_{k-1}) = (\mathbf{I} - \alpha \mathbf{H})(\mathbf{x}_{k-1} - \tilde{\mathbf{x}}) + \alpha \frac{\mathbf{A}^T \boldsymbol{\eta}}{n}. \quad (26)$$

Applying Proposition 3.1 to this recurrence, we obtain the following polynomials:

$$P_k(\lambda; \alpha^{-1}) = (1 - \alpha \lambda)^k, \quad Q_k(\lambda; \alpha^{-1}) = \alpha \sum_{i=0}^{k-1} (1 - \alpha \lambda)^i \quad \text{with } Q_0(\lambda) = 0. \quad (27)$$

Nesterov's accelerated method. Nesterov's accelerated method [Nesterov, 2004] and its variant FISTA [Beck and Teboulle, 2009] generate iterates on (8) satisfying the recurrence

$$\begin{aligned} \mathbf{x}_{k+1} - \tilde{\mathbf{x}} &= (1 + \beta_{k-1})(\mathbf{I} - \alpha \mathbf{H})(\mathbf{x}_k - \tilde{\mathbf{x}}) - \beta_{k-1}(\mathbf{I} - \alpha \mathbf{H})(\mathbf{x}_{k-1} - \tilde{\mathbf{x}}) + \alpha \cdot \frac{\mathbf{A}^T \boldsymbol{\eta}}{n}, \\ \text{where } \alpha &= \frac{1}{\lambda_{\mathbf{H}}^+} \quad \text{and} \quad \beta_k = \begin{cases} \frac{\sqrt{\lambda_{\mathbf{H}}^+} - \sqrt{\lambda_{\mathbf{H}}^-}}{\sqrt{\lambda_{\mathbf{H}}^+} + \sqrt{\lambda_{\mathbf{H}}^-}}, & \text{if } \lambda_{\mathbf{H}}^- \neq 0 \\ \frac{k}{k+3}, & \text{if } \lambda_{\mathbf{H}}^- = 0, \end{cases} \end{aligned} \quad (28)$$

with initial vector $\mathbf{x}_0 \in \mathbb{R}^d$ and $\mathbf{x}_1 = \mathbf{x}_0 - \alpha \nabla f(\mathbf{x}_0)$. Unrolling the recurrence, we can obtain an explicit formula for the corresponding polynomials

$$\begin{aligned} P_{k+1}(\lambda; \lambda_{\mathbf{H}}^{\pm}) &= (1 + \beta_{k-1})(1 - \alpha \lambda)P_k(\lambda; \lambda_{\mathbf{H}}^{\pm}) - \beta_{k-1}(1 - \alpha \lambda)P_{k-1}(\lambda; \lambda_{\mathbf{H}}^{\pm}) \\ &\quad \text{with } P_0(\lambda; \lambda_{\mathbf{H}}^{\pm}) = 1 \quad \text{and} \quad P_1(\lambda; \lambda_{\mathbf{H}}^{\pm}) = 1 - \alpha \lambda \\ Q_{k+1}(\lambda; \lambda_{\mathbf{H}}^{\pm}) &= (1 + \beta_{k-1})(1 - \alpha \lambda)Q_k(\lambda; \lambda_{\mathbf{H}}^{\pm}) - \beta_{k-1}(1 - \alpha \lambda)Q_{k-1}(\lambda; \lambda_{\mathbf{H}}^{\pm}) + \alpha \\ &\quad \text{with } Q_0(\lambda; \lambda_{\mathbf{H}}^{\pm}) = 0 \quad \text{and} \quad Q_1(\lambda; \lambda_{\mathbf{H}}^{\pm}) = \alpha. \end{aligned} \quad (29)$$

We derive the polynomials P_k explicitly in Appendix A.1. When $\lambda_{\mathbf{H}}^- > 0$ (strongly convex), the polynomial P_k is given by

$$\begin{aligned} P_k(\lambda; \lambda_{\mathbf{H}}^{\pm}) &= \frac{2\beta}{1+\beta} (\beta(1 - \alpha \lambda))^{k/2} T_k \left(\frac{1+\beta}{2\sqrt{\beta}} \sqrt{1 - \alpha \lambda} \right) + \left(1 - \frac{2\beta}{1+\beta} \right) (\beta(1 - \alpha \lambda))^{k/2} U_k \left(\frac{1+\beta}{2\sqrt{\beta}} \sqrt{1 - \alpha \lambda} \right), \\ \text{where } \beta &= \frac{\sqrt{\lambda_{\mathbf{H}}^+} - \sqrt{\lambda_{\mathbf{H}}^-}}{\sqrt{\lambda_{\mathbf{H}}^+} + \sqrt{\lambda_{\mathbf{H}}^-}}, \quad \text{and} \quad \alpha = \frac{1}{\lambda_{\mathbf{H}}^+}, \end{aligned} \quad (30)$$

where T_k and U_k the Chebyshev polynomials of the 1st and 2nd-kind respectively. When the smallest eigenvalue of \mathbf{H} is equal to 0 (non-strongly convex setting) the polynomial P_k is given by

$$P_k(\lambda; \lambda_{\mathbf{H}}^{\pm}) = \frac{2(1 - \alpha \lambda)^{(k+1)/2}}{k\alpha\lambda} \left(\sqrt{1 - \alpha \lambda} L_k(\sqrt{1 - \alpha \lambda}) - L_{k+1}(\sqrt{1 - \alpha \lambda}) \right), \quad (31)$$

where L_k are the Legendre polynomials.

Working directly with the polynomial in (31) will prove difficult. As such, we derive an asymptotic for this polynomial. Nesterov's polynomial satisfies in a sufficiently strong sense

$$P_k(\lambda; \lambda_{\mathbf{H}}^{\pm}) \sim \frac{2J_1(k\sqrt{\alpha\lambda})}{k\sqrt{\alpha\lambda}} e^{-\alpha\lambda k/2} \quad (32)$$

where J_1 is the Bessel function of the first kind. A derivation of this can be found in Appendix A.1.2. Let $f(t, z) \stackrel{\text{def}}{=} P_{tn}(zn^{-2}; \lambda_{\mathbf{H}}^{\pm})$. Then the recurrence in (29) becomes a discrete approximation to the initial value problem

$$\partial_{tt}f + \frac{3}{t}\partial_t f + zf = 0, \quad f(t, 0) = 1 \text{ and } \partial f_t(t, 0) = 0,$$

which bears a strong resemblance to the differential equation model for Nesterov's accelerated method in [Su et al., 2016]. The solution to this initial value problem is $\frac{2J_1(k\sqrt{\alpha\lambda})}{k\sqrt{\alpha\lambda}}$. Our result in (32), not derived using this differential equation, yields an even tighter result for Nesterov's accelerated method by including the exponential.

Polyak momentum algorithm. We aim to derive the residual polynomials for the Polyak momentum algorithm (a.k.a Heavy-ball method) [Polyak, 1964]. The Polyak momentum algorithm takes as arguments the largest and smallest eigenvalues of \mathbf{H} and iterates as follows

$$\begin{aligned} \mathbf{x}_{k+1} - \tilde{\mathbf{x}} &= \mathbf{x}_k - \tilde{\mathbf{x}} + m(\mathbf{x}_{k-1} - \tilde{\mathbf{x}} - (\mathbf{x}_k - \tilde{\mathbf{x}})) + \alpha \nabla f(\mathbf{x}_k), \\ \mathbf{x}_0 &\in \mathbb{R}^d, \quad \mathbf{x}_1 - \tilde{\mathbf{x}} = \mathbf{x}_0 - \tilde{\mathbf{x}} - \frac{2}{\lambda_{\mathbf{H}}^+ + \lambda_{\mathbf{H}}^-} \nabla f(\mathbf{x}_0) \end{aligned} \quad (33)$$

where $m = -\left(\frac{\sqrt{\lambda_{\mathbf{H}}^+} - \sqrt{\lambda_{\mathbf{H}}^-}}{\sqrt{\lambda_{\mathbf{H}}^+} + \sqrt{\lambda_{\mathbf{H}}^-}}\right)^2$ and $\alpha = -\frac{4}{(\sqrt{\lambda_{\mathbf{H}}^-} + \sqrt{\lambda_{\mathbf{H}}^+})^2}$.

Using these initial conditions, the residual polynomials for Polyak momentum satisfy

$$\begin{aligned} P_{k+1}(\lambda; \lambda_{\mathbf{H}}^{\pm}) &= (1 - m + \alpha\lambda)P_k(\lambda; \lambda_{\mathbf{H}}^{\pm}) + mP_{k-1}(\lambda; \lambda_{\mathbf{H}}^{\pm}), \\ \text{with } P_0(\lambda; \lambda_{\mathbf{H}}^{\pm}) &= 1, \quad P_1(\lambda; \lambda_{\mathbf{H}}^{\pm}) = 1 - \frac{2}{\lambda_{\mathbf{H}}^+ + \lambda_{\mathbf{H}}^-} \lambda \end{aligned} \quad (34)$$

and $Q_{k+1}(\lambda; \lambda_{\mathbf{H}}^{\pm}) = (1 - m + \alpha\lambda)Q_k(\lambda; \lambda_{\mathbf{H}}^{\pm}) + mQ_{k-1}(\lambda; \lambda_{\mathbf{H}}^{\pm}) - \alpha$,
with $Q_0(\lambda; \lambda_{\mathbf{H}}^{\pm}) = 0$, $Q_1(\lambda; \lambda_{\mathbf{H}}^{\pm}) = \frac{2}{\lambda_{\mathbf{H}}^+ + \lambda_{\mathbf{H}}^-}$.

By recognizing this three-term recurrence as Chebyshev polynomials, we can derive an explicit representation for P_k namely

$$P_k(\lambda; \lambda_{\mathbf{H}}^{\pm}) = \left(\frac{\sqrt{\lambda_{\mathbf{H}}^+} - \sqrt{\lambda_{\mathbf{H}}^-}}{\sqrt{\lambda_{\mathbf{H}}^+} + \sqrt{\lambda_{\mathbf{H}}^-}}\right)^k \left[\frac{(\sqrt{\lambda_{\mathbf{H}}^+} - \sqrt{\lambda_{\mathbf{H}}^-})^2}{\lambda_{\mathbf{H}}^+ + \lambda_{\mathbf{H}}^-} \cdot T_k(\sigma(\lambda)) + \frac{2\sqrt{\lambda_{\mathbf{H}}^-} \lambda_{\mathbf{H}}^+}{\lambda_{\mathbf{H}}^+ + \lambda_{\mathbf{H}}^-} \cdot U_k(\sigma(\lambda)) \right] \quad (35)$$

where $T_k(x)$ and $U_k(x)$ are the Chebyshev polynomials of the 1st and 2nd-kind respectively

$$\text{and } \sigma(\lambda) = \frac{\lambda_{\mathbf{H}}^+ + \lambda_{\mathbf{H}}^- - 2\lambda}{\lambda_{\mathbf{H}}^+ - \lambda_{\mathbf{H}}^-}.$$




Figure 5: **Bessel approx. of Nesterov's (convex) poly.** For small λ , the Bessel approx. (blue) in (32) and Nesterov's (convex) poly. (orange) are indistinguishable. Only when λ is far from zero that one sees any, albeit minor, differences.




Figure 6: **Residual polynomials.** The oscillations in the polynomials for Nesterov’s accelerated method (convex) are pronounced near zero compared with the other methods. In fact, both Nesterov (strongly convex) and Polyak momentum polynomials decay quite rapidly to the zero polynomial. To see these oscillations in the figures one needs to have a badly conditioned matrix (condition number 40,000). The slower decay to zero in the residual polynomials for Nesterov (strongly convex) as compared with Polyak’s momentum suggest a worst rate of convergence.

3.2 Properties of residual polynomials

In the following sections, it will be convenient to know some general properties of residual polynomials. Particularly, the polynomials, $\lambda^2 P_k^2(\lambda; \lambda^\pm)$ and $\lambda P_k^2(\lambda; \lambda^\pm)$ are uniformly bounded in k and that these polynomials goes to zero on some fixed support $[\lambda^-, \lambda^+]$. The importance of these facts are twofold. First, these polynomials appear in the formula for the expected gradient, Theorem 1.1. Second, we use the boundedness and convergence properties in the proof of halting time universality, Theorem 1.2. If one *a priori* knows an explicit expression for these polynomials, then these properties are easily deduced. However when such an expression does not exist, we still can conclude these properties hold provided that the algorithm is *convergent*.

Definition 3.2 (Convergent algorithms). We say a gradient-based method is (*strongly*) *convergent* if for every matrix \mathbf{A} such that $(\mathbf{A}^T \mathbf{A} > 0)$ $\mathbf{A}^T \mathbf{A} \succeq 0$ and any vectors \mathbf{b} and \mathbf{x}_0 , we have that the sequence of iterates generated by the algorithm starting at \mathbf{x}_0 satisfies $\|\nabla f(\mathbf{x}_k)\|^2 \rightarrow 0$ as $k \rightarrow \infty$ and there exists constants C, \tilde{C} depending on λ_H^+ and λ_H^- such that

$$\begin{aligned} \|\nabla f(\mathbf{x}_k)\|^2 &\leq C(f(\mathbf{x}_0) - f(\mathbf{x}^*) + \|\mathbf{x}_0 - \mathbf{x}^*\|^2) \\ f(\mathbf{x}_k) - f(\mathbf{x}^*) &\leq \tilde{C}(f(\mathbf{x}_0) - f(\mathbf{x}^*) + \|\mathbf{x}_0 - \mathbf{x}^*\|^2) \end{aligned} \quad (36)$$

where \mathbf{x}^* in the optimum of (8).

Remark 3.1 (Minimal norm solutions). For any gradient-based method, the iterates generated by the algorithm on the least squares problem (8) satisfy $\mathbf{x}_k \in \mathbf{x}_0 + \text{span}\{\nabla f(\mathbf{x}_0), \dots, \nabla f(\mathbf{x}_{k-1})\} \subseteq \mathbf{x}_0 + \text{Null}(\mathbf{A})^\perp$. If the algorithm converges to some \mathbf{x}^* and we have that $\mathbf{A}^T \mathbf{A} \mathbf{x}^* = \mathbf{A}^T \mathbf{b}$, then the solution \mathbf{x}^* is independent of the algorithm in the following sense

$$\mathbf{x}^* = \arg \min_{\{\mathbf{x} : \mathbf{A}^T \mathbf{A} \mathbf{x} = \mathbf{A}^T \mathbf{b}\}} \|\mathbf{x}_0 - \mathbf{x}\|_2^2. \quad (37)$$

In particular when $\mathbf{x}_0 \in \text{Null}(\mathbf{A})^\perp$, the optimum \mathbf{x}^* is the minimal norm solution. See e.g., Gunasekar et al. [2018], Wilson et al. [2017] and references therein.

Remark 3.2. All the algorithms discussed in Section 3.1 are convergent.

The following lemma shows that convergent algorithms have residual polynomials which go to 0 as $k \rightarrow \infty$ on compact subsets of the positive real line. In essence if optimality measures go to zero, then so must the residual polynomial.

Lemma 3.1 (Convergent algorithms \Rightarrow Residual polynomials $\rightarrow 0$). *Suppose the algorithm \mathcal{A} is a (strongly) convergent gradient-based method. Fix positive constants $0 \leq \lambda^- < \lambda^+$ for a convergent algorithm and constants $0 < \lambda^- < \lambda^+$ if one has a strongly convergent algorithm. The residual polynomial, P_k , for the algorithm \mathcal{A} satisfies*

$$\lim_{k \rightarrow \infty} \lambda^2 P_k^2(\lambda; \lambda^\pm) = 0 \quad \text{and} \quad \lim_{k \rightarrow \infty} \lambda P_k^2(\lambda; \lambda^\pm) = 0 \quad \text{for all } \lambda \in [\lambda^-, \lambda^+].$$

Proof. Suppose we consider the noiseless setting where $\boldsymbol{\eta} = (0, 0, 0)^T$ in the generative model so that $\mathbf{A}\tilde{\mathbf{x}} = \mathbf{b}$. Fix a constant $\lambda \in [\lambda^-, \lambda^+]$ and define the following matrix and vectors

$$\mathbf{A} = \begin{bmatrix} \sqrt{3\lambda^+} & 0 & 0 \\ 0 & \sqrt{3\lambda} & 0 \\ 0 & 0 & \sqrt{3\lambda^-} \end{bmatrix}, \quad \mathbf{x}_0 - \tilde{\mathbf{x}} = (0, 1, 0)^T, \quad \text{and} \quad \boldsymbol{\eta} = (0, 0, 0)^T. \quad (38)$$

A simple computation shows that $\mathbf{H} = \frac{1}{3}\mathbf{A}^T\mathbf{A} = \text{diag}(\lambda^+, \lambda, \lambda^-)$. Because the method is (strongly) convergent, the algorithm converges for these choices of \mathbf{H} and $\mathbf{x}_0 - \tilde{\mathbf{x}}$. Moreover, we know that $\nabla f(\mathbf{x}_k) = \mathbf{H}(\mathbf{x}_k - \tilde{\mathbf{x}})$ and by Proposition 3.1, the vector $\mathbf{x}_k - \tilde{\mathbf{x}} = P_k(\mathbf{H}; \lambda^\pm)(\mathbf{x}_0 - \tilde{\mathbf{x}})$. Therefore we have that

$$\lim_{k \rightarrow \infty} \lambda^2 P_k^2(\lambda; \lambda^\pm) = \lim_{k \rightarrow \infty} (\mathbf{x}_0 - \tilde{\mathbf{x}})^T \mathbf{H}^2 P_k^2(\mathbf{H}; \lambda^\pm) (\mathbf{x}_0 - \tilde{\mathbf{x}}) = \lim_{k \rightarrow \infty} \|\nabla f(\mathbf{x}_k)\|^2 = 0. \quad (39)$$

Similarly, we consider the same matrix \mathbf{A} as in (38) but instead a pure noise setting,

$$\mathbf{x}_0 - \tilde{\mathbf{x}} = (0, 0, 0)^T, \quad \text{and} \quad \boldsymbol{\eta} = (0, \sqrt{3}, 0)^T. \quad (40)$$

As before, the matrix $\mathbf{H} = \text{diag}(\lambda^+, \lambda, \lambda^-)$. By Proposition 3.1, the iterates $\mathbf{x}_k - \tilde{\mathbf{x}} = Q_k(\mathbf{H}; \lambda^\pm) \frac{\mathbf{A}^T \boldsymbol{\eta}}{3}$ as $\mathbf{x}_0 - \tilde{\mathbf{x}} = \mathbf{0}$. With this, the gradient equals

$$\nabla f(\mathbf{x}_k) = \mathbf{H}(\mathbf{x}_k - \tilde{\mathbf{x}}) - \frac{\mathbf{A}^T \boldsymbol{\eta}}{3} = [\mathbf{H}Q_k(\mathbf{H}; \lambda^\pm) - \mathbf{I}] \frac{\mathbf{A}^T \boldsymbol{\eta}}{3} = -P_k(\mathbf{H}; \lambda^\pm) \frac{\mathbf{A}^T \boldsymbol{\eta}}{3}.$$

Here, again, we used Proposition 3.1. A (strongly) convergent method has the following

$$\lim_{k \rightarrow \infty} \lambda P_k^2(\lambda; \lambda^\pm) = \lim_{k \rightarrow \infty} \frac{\boldsymbol{\eta}^T \mathbf{A} P_k^2(\mathbf{H}; \lambda^\pm) \frac{\mathbf{A}^T \boldsymbol{\eta}}{3}}{3} = \lim_{k \rightarrow \infty} \|\nabla f(\mathbf{x}_k)\|^2 = 0. \quad (41)$$

This completes the result. \square

The following lemma shows that the residual polynomials are uniformly bounded over k on any compact subset of the positive real line.

Lemma 3.2 (Convergent algorithms \Rightarrow boundedness of P_k). *Suppose \mathcal{A} is a (strongly) convergent algorithm with residual polynomial P_k . Under the assumptions of Lemma 3.1,*

$$\max_{k, \lambda \in [\lambda^-, \lambda^+]} \lambda^2 P_k^2(\lambda; \lambda^\pm) \leq B \quad \text{and} \quad \max_{k, \lambda \in [\lambda^-, \lambda^+]} \lambda P_k^2(\lambda; \lambda^\pm) \leq \tilde{B},$$

for some constants $B, \tilde{B} > 0$.

Proof. Suppose we consider the noiseless setting $\boldsymbol{\eta} = \mathbf{0}$ in the generative model (8) so that $\mathbf{A}\tilde{\mathbf{x}} = \mathbf{b}$. It then follows that $f(\mathbf{x}^*) = 0$ where \mathbf{x}^* is the optimum. A simple computation using Proposition 3.1 shows that for all $k \geq 0$

$$\begin{aligned} f(\mathbf{x}_k) - f(\mathbf{x}^*) &= \frac{1}{2n} \|\mathbf{A}(\mathbf{x}_k - \tilde{\mathbf{x}})\|^2 = \frac{1}{2} (\mathbf{x}_k - \tilde{\mathbf{x}})^T \mathbf{H} (\mathbf{x}_k - \tilde{\mathbf{x}}) \\ &= \frac{1}{2} (\mathbf{x}_0 - \tilde{\mathbf{x}})^T \mathbf{H} P_k^2(\mathbf{H}; \lambda^\pm) (\mathbf{x}_0 - \tilde{\mathbf{x}}). \end{aligned} \quad (42)$$

Next consider the matrix \mathbf{A} and vectors $\mathbf{x}_0 - \tilde{\mathbf{x}}$ and $\boldsymbol{\eta}$ as in (38) with the initial iterate $\mathbf{x}_0 = (0, 0, 0)^T$. We consider cases: suppose $\lambda^- = 0$. Fix a constant $\lambda \in [\lambda^-, \lambda^+]$. It follows from our choice of \mathbf{A} , \mathbf{x}_0 , $\tilde{\mathbf{x}}$, and $\boldsymbol{\eta}$ that the vector $\mathbf{A}\tilde{\mathbf{x}} = \mathbf{b} = (0, -\sqrt{3}\lambda, 0)^T$ and by (42) that

$$f(\mathbf{x}_0) - f(\mathbf{x}^*) = \frac{1}{2}\lambda P_0^2(\lambda; \lambda^\pm). \quad (43)$$

The solution set $\{\mathbf{x} : \mathbf{A}^T \mathbf{A} \mathbf{x} = \mathbf{A}^T \mathbf{b}\} = \{(0, -1, a)^T : a \in \mathbb{R}\}$ if $\lambda > 0$ and otherwise it equals $\{(0, a, b)^T : a, b \in \mathbb{R}\}$ if $\lambda = 0$. From Remark 3.1, we have that $\mathbf{x}^* = \underset{\mathbf{A}^T \mathbf{A} \mathbf{x} = \mathbf{A}^T \mathbf{b}}{\arg \min} \|\mathbf{x} - \mathbf{x}_0\|^2$ and thus we deduce that

$$\mathbf{x}^* = \begin{cases} (0, -1, 0)^T, & \text{if } \lambda > 0 \\ (0, 0, 0)^T & \text{if } \lambda = 0. \end{cases}$$

In both cases, we have that $\|\mathbf{x}_0 - \mathbf{x}^*\|^2 \leq 1$. Therefore using the boundedness assumption (36) and the expression for the gradient in (39), we have that

$$\begin{aligned} \sup_{k, \lambda \in [\lambda^-, \lambda^+]} \lambda^2 P_k^2(\lambda; \lambda^\pm) &= \sup_{k, \lambda \in [\lambda^-, \lambda^+]} \|\nabla f(\mathbf{x}_k)\|^2 \leq \sup_{k, \lambda \in [\lambda^-, \lambda^+]} C(f(\mathbf{x}_0) - f(\mathbf{x}^*) + \|\mathbf{x}_0 - \mathbf{x}^*\|^2) \\ &\leq \sup_{k, \lambda \in [\lambda^-, \lambda^+]} C(\frac{1}{2}\lambda P_0^2(\lambda; \lambda^\pm) + 1) \leq B. \end{aligned}$$

Here we used that the distance to the optimum $\|\mathbf{x}_0 - \mathbf{x}^*\|^2 \leq 1$ and the polynomial in (43) is bounded on a compact set.

Now we suppose that $\lambda^- > 0$. As above, we use the same matrix \mathbf{A} and vectors $\mathbf{x}_0 - \tilde{\mathbf{x}}$ and $\boldsymbol{\eta}$ as in (38) and, in addition, we set $\mathbf{x}_0 = (0, 0, 0)^T$. In this situation, the matrix \mathbf{A} is invertible and $\mathbf{x}^* = (0, -1, 0)^T$. Hence both (43) and $\|\mathbf{x}_0 - \mathbf{x}^*\|^2 \leq 1$ holds. Using the boundedness assumption on function values (36) and the expression for the function values in (42), we deduce

$$\begin{aligned} \sup_{k, \lambda \in [\lambda^-, \lambda^+]} \lambda P_k^2(\lambda; \lambda^\pm) &= \sup_{k, \lambda \in [\lambda^-, \lambda^+]} f(\mathbf{x}_k) - f(\mathbf{x}^*) \leq \sup_{k, \lambda \in [\lambda^-, \lambda^+]} \tilde{C}(f(\mathbf{x}_0) - f(\mathbf{x}^*) + \|\mathbf{x}_0 - \mathbf{x}^*\|^2) \\ &\leq \sup_{k, \lambda \in [\lambda^-, \lambda^+]} \tilde{C}(\frac{1}{2}\lambda P_0^2(\lambda; \lambda^\pm) + 1) \leq \tilde{B}. \end{aligned}$$

The result immediately follows. \square

4 Halting time is almost deterministic

In this section we develop a framework for the average-case analysis and state a main result of this paper: the concentration of the halting time. We denote the halting time T_ε as the first iteration at which the gradient falls below some predefined ε :

$$T_\varepsilon \stackrel{\text{def}}{=} \inf \{k > 0 : \|\nabla f(\mathbf{x}_k)\|^2 \leq \varepsilon\}. \quad (44)$$

Our main result (Theorem 4.1) states that this halting time is predictable for almost all high-dimensional data, or more precisely,

$$\lim_{d \rightarrow \infty} \Pr(T_\varepsilon = \text{constant}) = 1. \quad (45)$$

Furthermore, we provide an implicit expression for this constant, otherwise known as the average complexity, and in Tables 1 and 2 an explicit expression under further assumptions. The rest of this section provides a proof of this result.

4.1 First-order methods as polynomials

Using the residual polynomials from Proposition 3.1, we get the following identity for the gradient.

Proposition 4.1 (Residual polynomials and gradients). *Suppose the iterates $\{\mathbf{x}_k\}_{k=0}^\infty$ are generated from a gradient based method. Let $\{P_k\}_{k=0}^\infty$ be a sequence of polynomials defined in (24). Then the following identity exists between the iterates and its residual polynomial,*

$$\begin{aligned} \|\nabla f(\mathbf{x}_k)\|^2 &= (\mathbf{x}_0 - \tilde{\mathbf{x}})^T \mathbf{H}^2 P_k^2(\mathbf{H}; \lambda_{\mathbf{H}}^\pm)(\mathbf{x}_0 - \tilde{\mathbf{x}}) + \frac{\boldsymbol{\eta}^T \mathbf{A}}{n} P_k^2(\mathbf{H}; \lambda_{\mathbf{H}}^\pm) \frac{\mathbf{A}^T \boldsymbol{\eta}}{n} \\ &\quad - 2(\mathbf{x}_0 - \tilde{\mathbf{x}})^T \mathbf{H} P_k^2(\mathbf{H}; \lambda_{\mathbf{H}}^\pm) \frac{\mathbf{A}^T \boldsymbol{\eta}}{n}. \end{aligned}$$

Proof. The gradient in (8) is given by the expression $\nabla f(\mathbf{x}_k) = \mathbf{H}(\mathbf{x}_k - \tilde{\mathbf{x}}) - \frac{\mathbf{A}^T \boldsymbol{\eta}}{n}$. The result follows immediately by plugging in (25) into the formula for the gradient and using the relationship that $\mathbf{H}^2 Q_k^2(\mathbf{H}; \lambda_{\mathbf{H}}^\pm) - 2\mathbf{H} Q_k(\mathbf{H}; \lambda_{\mathbf{H}}^\pm) + \mathbf{I} = (\mathbf{I} - \mathbf{H} Q_k(\mathbf{H}; \lambda_{\mathbf{H}}^\pm))^2 = P_k^2(\mathbf{H}; \lambda_{\mathbf{H}}^\pm)$. \square

This *equality* for the squared norm of the gradient is crucial for deriving average-case rates. In contrast, worst-case analysis typically uses only *bounds* on the norm. A difficulty with the polynomials P_k and Q_k is that their coefficients depend on the largest and smallest eigenvalue of \mathbf{H} , and hence are random. We can remove this randomness thanks to Assumption 2, replacing $\lambda_{\mathbf{H}}^+$ and $\lambda_{\mathbf{H}}^-$ with the top (bottom) edge of the support of μ , denoted by λ^+ and λ^- , without loss of generality.

Proposition 4.2 (Remove randomness in coefficients of polynomial). *Suppose Assumption 2 holds. Fix any k -degree polynomial \tilde{P}_k whose coefficients depend continuously on the largest and smallest eigenvalues of \mathbf{H} . Then the following hold*

$$\|\tilde{P}_k(\mathbf{H}; \lambda_{\mathbf{H}}^\pm) - \tilde{P}_k(\mathbf{H}; \lambda^\pm)\|_{\text{op}}^2 \xrightarrow{d \rightarrow \infty} 0. \quad (46)$$

Proof. Fix any $\varepsilon, \delta > 0$. Let $c_i(\cdot)$ where $i = 0, \dots, k$ be the coefficients associated with the term of degree i in $\tilde{P}_k(\mathbf{H}; \cdot)$. For each i , the continuity of $c_i(\cdot)$ implies there exists $\delta_\varepsilon > 0$ such that

$$\text{whenever } \|(\lambda_{\mathbf{H}}^+, \lambda_{\mathbf{H}}^-) - (\lambda^+, \lambda^-)\| \leq \delta_\varepsilon \quad \Rightarrow \quad |c_i(\lambda_{\mathbf{H}}^\pm) - c_i(\lambda^\pm)| \leq \frac{\varepsilon}{4(4\lambda^+)^i}. \quad (47)$$

For sufficiently large d , Assumption 2 implies $\Pr(|\lambda_{\mathbf{H}}^+ - \lambda^+| > \min\{\frac{\delta_\varepsilon}{2}, \lambda^+\}) \leq \frac{\delta}{2}$ and $\Pr(|\lambda_{\mathbf{H}}^- - \lambda^-| > \min\{\frac{\delta_\varepsilon}{2}, \lambda^+\}) \leq \frac{\delta}{2}$. With this, we define the event $\mathcal{S} = \{|\lambda_{\mathbf{H}}^+ - \lambda^+| \leq \min\{\frac{\delta_\varepsilon}{2}, \lambda^+\}\} \cap \{|\lambda_{\mathbf{H}}^- - \lambda^-| \leq \min\{\frac{\delta_\varepsilon}{2}, \lambda^+\}\}$. We have for all sufficiently large d

$$\begin{aligned} \Pr(\|\tilde{P}_k(\mathbf{H}; \lambda_{\mathbf{H}}^\pm) - \tilde{P}_k(\mathbf{H}; \lambda^\pm)\|_{\text{op}} > \varepsilon) &= \Pr(\mathcal{S} \cap \{\|\tilde{P}_k(\mathbf{H}; \lambda_{\mathbf{H}}^\pm) - \tilde{P}_k(\mathbf{H}; \lambda^\pm)\|_{\text{op}} > \varepsilon\}) \\ &\quad + \Pr(\mathcal{S}^c \cap \{\|\tilde{P}_k(\mathbf{H}; \lambda_{\mathbf{H}}^\pm) - \tilde{P}_k(\mathbf{H}; \lambda^\pm)\|_{\text{op}} > \varepsilon\}) \\ &\leq \Pr(\mathcal{S} \cap \{\|\tilde{P}_k(\mathbf{H}; \lambda_{\mathbf{H}}^\pm) - \tilde{P}_k(\mathbf{H}; \lambda^\pm)\|_{\text{op}} > \varepsilon\}) + \delta. \end{aligned} \quad (48)$$

Here we used that $\Pr(\mathcal{S}^c \cap \{\|\tilde{P}_k(\mathbf{H}; \lambda_{\mathbf{H}}^\pm) - \tilde{P}_k(\mathbf{H}; \lambda^\pm)\|_{\text{op}} > \varepsilon\}) \leq \Pr(\mathcal{S}^c) \leq \delta$ for large d . Therefore, it suffices to consider the first term in (48) and show that it is 0. By construction of the set \mathcal{S} , any element in \mathcal{S} satisfies both $\|\mathbf{H}\|_{\text{op}} \leq 2\lambda^+$ and $|c_i(\lambda_{\mathbf{H}}^\pm) - c_i(\lambda^\pm)| \leq \frac{\varepsilon}{4(4\lambda^+)^i}$. Hence on the event \mathcal{S} , we have the following

$$\|\tilde{P}_k(\mathbf{H}; \lambda_{\mathbf{H}}^\pm) - \tilde{P}_k(\mathbf{H}; \lambda^\pm)\|_{\text{op}} \leq \sum_{i=0}^k |c_i(\lambda_{\mathbf{H}}^\pm) - c_i(\lambda^\pm)| \|\mathbf{H}\|_{\text{op}}^i \leq \sum_{i=0}^k \frac{(2\lambda^+)^i \varepsilon}{4(4\lambda^+)^i} \leq \frac{\varepsilon}{2}. \quad (49)$$

From this, we deduce that $\Pr(\mathcal{S} \cap \{\|\tilde{P}_k(\mathbf{H}; \lambda_{\mathbf{H}}^\pm) - \tilde{P}_k(\mathbf{H}; \lambda^\pm)\|_{\text{op}} > \varepsilon\}) = 0$ and the result immediately follows by (48). \square

The squared norm of the gradient in (46) is a quadratic form. In Proposition 4.2, we removed the randomness in the coefficients of the polynomial and now we will relate this back to the squared norm of the gradient, and particularly, the quadratic form. The following lemmas state this precisely.

Lemma 4.1. *Suppose the sequences of non-negative random variables $X_d, Y_d \geq 0$ satisfy $\mathbb{E}[X_d] \leq \gamma < \infty$ and $Y_d \xrightarrow[d \rightarrow \infty]{\text{Pr}} 0$. Then $X_d Y_d \xrightarrow[d \rightarrow \infty]{\text{Pr}} 0$.*

Proof. Fix constants $\varepsilon, \delta > 0$ and suppose we set $\hat{\varepsilon} = \frac{\varepsilon \delta}{2\gamma}$ and $\hat{\delta} = \frac{\delta}{2}$. Because Y_d converges in probability, we have $\Pr(Y_d > \hat{\varepsilon}) \leq \hat{\delta}$ for sufficiently large d . Define the event $\mathcal{S} = \{Y_d \leq \hat{\varepsilon}\}$ and decompose the space based on this set \mathcal{S} so that for large d

$$\Pr(X_d Y_d > \varepsilon) = \Pr(\mathcal{S} \cap \{X_d Y_d > \varepsilon\}) + \Pr(\mathcal{S}^c \cap \{X_d Y_d > \varepsilon\}) \leq \Pr(\mathcal{S} \cap \{X_d Y_d > \varepsilon\}) + \frac{\delta}{2}.$$

Here we used that $\Pr(\mathcal{S}^c \cap \{X_d Y_d > \varepsilon\}) \leq \Pr(\mathcal{S}^c)$. For the other term, a direct application of Markov's inequality yields

$$\Pr(\mathcal{S} \cap \{X_d Y_d > \varepsilon\}) \leq \Pr(\mathcal{S} \cap \{\hat{\varepsilon} X_d > \varepsilon\}) \leq \frac{\hat{\varepsilon}}{\varepsilon} \cdot \mathbb{E}[X_d] \leq \frac{\delta}{2}.$$

The result immediately follows. \square

Lemma 4.2 (Remove randomness in coefficients of quadratic form). *Suppose Assumption 2 holds and let the vectors $\mathbf{w} \in \mathbb{R}^d$ and $\mathbf{v} \in \mathbb{R}^d$ be i.i.d. satisfying $\mathbb{E}[\|\mathbf{w}\|_2^2] = R^2$ and $\mathbb{E}[\|\mathbf{v}\|_2^2] = \tilde{R}^2$ for some constants $R, \tilde{R} > 0$.*

For any k degree polynomial \tilde{P}_k whose coefficients depend continuously on $\lambda_{\mathbf{H}}^+$ and $\lambda_{\mathbf{H}}^-$, the quadratic form converges in probability

$$\mathbf{w}^T \tilde{P}_k(\mathbf{H}; \lambda_{\mathbf{H}}^\pm) \mathbf{v} - \mathbf{w}^T \tilde{P}_k(\mathbf{H}; \lambda^\pm) \mathbf{v} \xrightarrow[d \rightarrow \infty]{\text{Pr}} 0.$$

Proof. Using the Cauchy-Schwarz inequality, it suffices to show that for every $\varepsilon > 0$ we have

$$\lim_{d \rightarrow \infty} \Pr\left(\|\mathbf{w}\|_2 \cdot \|\mathbf{v}\|_2 \cdot \left\| \tilde{P}_k(\mathbf{H}; \lambda_{\mathbf{H}}^\pm) - \tilde{P}_k(\mathbf{H}; \lambda^\pm) \right\|_{\text{op}} > \varepsilon\right) = 0.$$

Define $X_d = \|\mathbf{w}\|_2 \|\mathbf{v}\|_2$ and $Y_d = \left\| \tilde{P}_k(\mathbf{H}; \lambda_{\mathbf{H}}^\pm) - \tilde{P}_k(\mathbf{H}; \lambda^\pm) \right\|_{\text{op}}$. Proposition 4.2 immediately gives that $Y_d \xrightarrow[d \rightarrow \infty]{\text{Pr}} 0$. Next, Cauchy-Schwarz implies

$$\mathbb{E}[X_d] = \mathbb{E}[\|\mathbf{w}\|_2 \|\mathbf{v}\|_2] \leq \mathbb{E}[\|\mathbf{w}\|_2^2]^{1/2} \mathbb{E}[\|\mathbf{v}\|_2^2]^{1/2} = R\tilde{R}.$$

The result immediately follows after applying Lemma 4.1. \square

From Lemma 4.2 and the expression for the squared norm of the gradient in (46), we can replace the maximum (minimum) eigenvalue $\lambda_{\mathbf{H}}^+$ ($\lambda_{\mathbf{H}}^-$) in (46) with the top (bottom) edge of the support of μ , λ^+ (λ^-). This followed because the vectors $\mathbf{x}_0 - \tilde{\mathbf{x}}$ and $\frac{1}{n} \mathbf{A}^T \boldsymbol{\eta}$ satisfy \mathbf{w} and \mathbf{v} in Lemma 4.2 and the terms surrounding these vectors in (46) are polynomials in \mathbf{H} .

4.2 Concentration of the gradient

Having established the key equation linking the gradient to a polynomial in Proposition 4.1, we now show that for almost any large model the magnitude of the gradient after k iterations converges to a deterministic value which we denote $\mathcal{E} \xrightarrow[d \rightarrow \infty]{\text{Pr}} \mathbb{E}[\|\nabla f(\mathbf{x}_k)\|^2]$. We recall the statement of Theorem 1.1:

Theorem. (Concentration of the gradient) *Under Assumptions 1 and 2 the norm of the gradient concentrates around a deterministic value:*

$$\|\nabla f(\mathbf{x}_k)\|^2 \xrightarrow[d \rightarrow \infty]{\text{Pr}} \underbrace{\widetilde{R}^2}_{\text{signal}} \int \underbrace{\lambda^2 P_k^2(\lambda; \lambda^\pm)}_{\text{algorithm}} \underbrace{d\mu}_{\text{model}} + \underbrace{\widetilde{R}^2}_{\text{noise}} r \int \underbrace{\lambda P_k^2(\lambda; \lambda^\pm)}_{\text{algorithm}} \underbrace{d\mu}_{\text{model}} \stackrel{\text{def}}{=} \mathcal{E}_{d \rightarrow \infty} [\|\nabla f(\mathbf{x}_k)\|^2]. \quad (50)$$

Intuitively, the value of $\mathcal{E}_{d \rightarrow \infty} [\|\nabla f(\mathbf{x}_k)\|^2]$ is the expected gradient after first taking the model size to infinity. The above expression explicitly illustrates the effects of the model and the algorithm on the norm of the gradient: the **signal** (\widetilde{R}^2) and **noise** (\widetilde{R}^2), the optimization algorithm which enters into the formula through the polynomial P_k , and the **model used to generate \mathbf{A}** by means of the measure μ .

The main tool to prove Theorem 1.1 is the moment method which requires computing explicit expressions for the moments of the norm of the gradient. We summarize this in the following proposition. To ease notation in the next few propositions, we define the following matrices and vectors

$$\mathbf{u} \stackrel{\text{def}}{=} \mathbf{x}_0 - \tilde{\mathbf{x}}, \quad \mathbf{B} \stackrel{\text{def}}{=} \mathbf{H}^2 P_k^2(\mathbf{H}; \lambda^\pm), \quad \mathbf{C} \stackrel{\text{def}}{=} P_k^2(\mathbf{H}; \lambda^\pm), \quad (51)$$

and $\mathbf{D} \stackrel{\text{def}}{=} -2\mathbf{H}P_k^2(\mathbf{H}; \lambda^\pm)$

and let y be the quadratic form given by

$$y \stackrel{\text{def}}{=} \mathbf{u}^T \mathbf{B} \mathbf{u} + \frac{1}{n} \mathbf{u}^T \mathbf{D} \mathbf{A}^T \boldsymbol{\eta} + \frac{1}{n^2} \boldsymbol{\eta}^T \mathbf{A} \mathbf{C} \mathbf{A}^T \boldsymbol{\eta}. \quad (52)$$

Observe that the value y is simply $\|\nabla f(\mathbf{x}_k)\|^2$ in (46) with $\lambda_{\mathbf{H}}^\pm$ replaced with λ^\pm .

Proposition 4.3. *Suppose the matrix \mathbf{A} and vectors $\mathbf{x}_0, \tilde{\mathbf{x}}$, and $\boldsymbol{\eta}$ satisfy Assumptions 1 and 2. Let P_k be the k -degree polynomial defined in (24). Using the notation in (51) and (52), the following holds for any $\varepsilon > 0$*

$$\begin{aligned} \Pr(|y - [R^2 \text{tr}(\frac{\mathbf{B}}{d}) + \widetilde{R}^2 \text{tr}(\frac{\mathbf{C}\mathbf{H}}{n})]| > \varepsilon | \mathbf{H}) \\ \leq \frac{1}{\varepsilon^2} \left(\frac{C - R^4}{d} \text{tr}(\frac{\mathbf{B}^2}{d}) + \frac{\widetilde{C} - \widetilde{R}^4}{n} \text{tr}(\frac{(\mathbf{C}\mathbf{H})^2}{n}) + \frac{R^2 \widetilde{R}^2}{n} [\frac{\text{tr}(\mathbf{D}^2 \mathbf{H})}{d}] \right). \end{aligned} \quad (53)$$

Without loss of generality, we assume that the constants C and \widetilde{C} are large enough such that $C > 3R^4$ and $\widetilde{C} > 3\widetilde{R}^4$.

Proof. We can write any quadratic form as $\mathbf{w}^T \mathbf{F} \mathbf{z} = \sum_{i,j} w_i z_j F_{ij}$. Expanding the quadratic forms, the following holds

$$\mathbb{E}[y | \mathbf{H}] = \mathbb{E}[\mathbf{u}^T \mathbf{B} \mathbf{u} | \mathbf{H}] + \frac{1}{n} \mathbb{E}[\mathbf{u}^T \mathbf{D} \mathbf{A}^T \boldsymbol{\eta} | \mathbf{H}] + \frac{1}{n^2} \mathbb{E}[\boldsymbol{\eta}^T \mathbf{A} \mathbf{C} \mathbf{A}^T \boldsymbol{\eta} | \mathbf{H}] \quad (54)$$

$$\text{(ind. of } \boldsymbol{\eta} \text{ and } \mathbf{u}, \mathbb{E}[\boldsymbol{\eta}] = \mathbf{0}) \quad = \mathbb{E} \left[\sum_{i,j} u_i u_j B_{ij} | \mathbf{H} \right] + \frac{1}{n^2} \mathbb{E} \left[\sum_{i,j} \eta_i \eta_j (\mathbf{A} \mathbf{C} \mathbf{A}^T)_{ij} | \mathbf{H} \right] \quad (55)$$

$$\text{(isotropic prop. of } \boldsymbol{\eta} \text{ and } \mathbf{u}) \quad = R^2 \cdot \sum_i \frac{B_{ii}}{d} + \widetilde{R}^2 \cdot \sum_i \frac{(\mathbf{A} \mathbf{C} \mathbf{A}^T)_{ii}}{n^2} \quad (56)$$

$$= R^2 \cdot \frac{\text{tr}(\mathbf{B})}{d} + \widetilde{R}^2 \cdot \frac{\text{tr}(\mathbf{C}\mathbf{H})}{n}. \quad (57)$$

In the last equality, we used that $\text{tr}(\mathbf{A} \mathbf{C} \mathbf{A}^T) = \text{tr}(\mathbf{C} \mathbf{A}^T \mathbf{A}) = n \cdot \text{tr}(\mathbf{C}\mathbf{H})$.

To prove (53), we will use Chebyshev's inequality; hence we need to compute the $\text{Var}(y | \mathbf{H}) = \mathbb{E}[y^2 | \mathbf{H}] - (\mathbb{E}[y | \mathbf{H}])^2$. First, a simple computation yields that

$$\mathbb{E}[y^2 | \mathbf{H}]^2 = \underbrace{\left[\frac{R^2 \text{tr}(\mathbf{B})}{d} \right]^2}_{(i)} + \underbrace{\left[\frac{\widetilde{R}^2 \text{tr}(\mathbf{C}\mathbf{H})}{n} \right]^2}_{(ii)} + 2 \underbrace{\left[\frac{R^2 \text{tr}(\mathbf{B})}{d} \right] \left[\frac{\widetilde{R}^2 \text{tr}(\mathbf{C}\mathbf{H})}{n} \right]}_{(iii)}. \quad (58)$$

Next, we compute $\mathbb{E}[y^2|\mathbf{H}]$. By expanding out the terms in (52), we get the following

$$\begin{aligned} \mathbb{E}[y^2|\mathbf{H}] &= \underbrace{\mathbb{E}[(\mathbf{u}^T \mathbf{B} \mathbf{u})^2|\mathbf{H}]}_{(a)} + \underbrace{\mathbb{E}\left[\left(\frac{\boldsymbol{\eta}^T \mathbf{A} \mathbf{C} \mathbf{A}^T \boldsymbol{\eta}}{n^2}\right)^2|\mathbf{H}\right]}_{(b)} + \underbrace{\mathbb{E}\left[\frac{2\mathbf{u}^T \mathbf{B} \mathbf{u} \boldsymbol{\eta}^T \mathbf{A} \mathbf{C} \mathbf{A}^T \boldsymbol{\eta}}{n^2}|\mathbf{H}\right]}_{(c)} \\ &\quad + \underbrace{\mathbb{E}\left[\left(\frac{\mathbf{u}^T \mathbf{D} \mathbf{A}^T \boldsymbol{\eta}}{n}\right)^2|\mathbf{H}\right]}_{(d)} + \underbrace{\mathbb{E}\left[2\left(\mathbf{u}^T \mathbf{B} \mathbf{u} + \frac{\boldsymbol{\eta}^T \mathbf{A} \mathbf{C} \mathbf{A}^T \boldsymbol{\eta}}{n^2}\right) \cdot \frac{\mathbf{u}^T \mathbf{D} \mathbf{A}^T \boldsymbol{\eta}}{n}|\mathbf{H}\right]}_{(e)}. \end{aligned} \quad (59)$$

To compute the variance of y , we take (59) and subtract (58). Since this is quite a long expression, we will match up terms and compute these terms individually. First consider the terms (a) and (i) in equations (59) and (58) respectively. By expanding out the square, we get

$$\text{Var}(\mathbf{u}^T \mathbf{B} \mathbf{u} | \mathbf{H}) = \mathbb{E}[(\mathbf{u}^T \mathbf{B} \mathbf{u})^2 | \mathbf{H}] - \left[\frac{R^2 \text{tr}(\mathbf{B})}{d}\right]^2 = \sum_{i,j,k,\ell} \mathbb{E}[u_i u_j u_k u_\ell] B_{ij} B_{k\ell} - \left[\frac{R^2 \text{tr}(\mathbf{B})}{d}\right]^2.$$

We need each index to appear exactly twice in the above for its contribution to be non-negligible since $\mathbb{E}[u_i^2] = \frac{R^2}{d}$ and $\mathbb{E}[\mathbf{u}] = \mathbf{0}$. There are four possible ways in which this can happen: $\{i = j = k = \ell\}$, $\{i = j, k = \ell, k \neq i\}$, $\{i = k, j = \ell, i \neq j\}$, or $\{i = \ell, j = k, i \neq j\}$. By the symmetry of the \mathbf{B} matrix, the last two cases are identical. Noting that $\mathbb{E}[u_i^4] \leq \frac{C}{d^2}$ and $\mathbb{E}[u_i^2] = \frac{R^2}{d}$, we, consequently, get the following expression for the variance

$$\begin{aligned} \text{Var}(\mathbf{u}^T \mathbf{B} \mathbf{u} | \mathbf{H}) &= \sum_i \mathbb{E}[u_i^4] \cdot B_{ii}^2 + \sum_{i \neq j} \mathbb{E}[u_i^2] \cdot \mathbb{E}[u_j^2] \cdot (B_{ii} B_{jj} + 2B_{ij}^2) - \frac{R^4}{d^2} [\text{tr}(\mathbf{B})]^2 \\ &\leq \frac{C - R^4}{d^2} \cdot \sum_i B_{ii}^2 + \frac{2R^4}{d^2} \sum_{i \neq j} B_{ij}^2 + \frac{R^4}{d^2} \left(\sum_i B_{ii}^2 + \sum_{i \neq j} B_{ii} B_{jj} - [\text{tr}(\mathbf{B})]^2 \right) \\ &= \frac{C - R^4}{d^2} \cdot \sum_i B_{ii}^2 + \frac{2R^4}{d^2} \sum_{i \neq j} B_{ij}^2 \\ &\leq \frac{C - R^4}{d^2} \cdot \left(\sum_i B_{ii}^2 + \sum_{i \neq j} B_{ij}^2 \right) = \frac{C - R^4}{d} \cdot \left[\frac{\text{tr}(\mathbf{B}^2)}{d} \right]. \end{aligned} \quad (60)$$

In the second equality, we used that $\sum_i B_{ii}^2 + \sum_{i \neq j} B_{ii} B_{jj} = [\text{tr}(\mathbf{B})]^2$ and in the second inequality we can without loss of generality choose C so that $C > 3R^4$. Finally, we used that $\sum_i B_{ii}^2 + \sum_{i \neq j} B_{ij}^2 = \text{tr}(\mathbf{B}^2)$.

Next, we consider the terms (b) and (ii) in equations (59) and (58) respectively. Similar to the previous case, by expanding out the square, we get the following

$$\text{Var}\left(\frac{\boldsymbol{\eta}^T \mathbf{A} \mathbf{C} \mathbf{A}^T \boldsymbol{\eta}}{n^2} | \mathbf{H}\right) = \mathbb{E}\left[\frac{1}{n^4} \sum_{i,j,k,\ell} \eta_i \eta_j \eta_k \eta_\ell (\mathbf{A} \mathbf{C} \mathbf{A}^T)_{ij} (\mathbf{A} \mathbf{C} \mathbf{A}^T)_{k\ell} | \mathbf{H}\right] - \left[\frac{\tilde{R}^2 \text{tr}(\mathbf{C} \mathbf{H})}{n}\right]^2.$$

Because of independence, isotropic variance $\mathbb{E}[\eta_i^2] = \tilde{R}^2$, and mean $\mathbb{E}[\boldsymbol{\eta}] = \mathbf{0}$, we need each index to appear exactly twice in the above expression in order for its contribution to be non-negligible. There are four possible ways in which this can happen: $\{i = j = k = \ell\}$, $\{i = j, k = \ell, k \neq i\}$, $\{i = k, j =$

$\ell, i \neq j\}$, or $\{i = \ell, j = k, i \neq j\}$. As before, we have the following expression for the variance

$$\begin{aligned}
\text{Var}\left(\frac{1}{n^2} \cdot \boldsymbol{\eta}^T \mathbf{A} \mathbf{C} \mathbf{A}^T \boldsymbol{\eta} \mid \mathbf{H}\right) &\leq \frac{\tilde{C}}{n^4} \sum_i (\mathbf{A} \mathbf{C} \mathbf{A}^T)_{ii}^2 + \frac{\tilde{R}^4}{n^4} \sum_{i \neq j} (\mathbf{A} \mathbf{C} \mathbf{A}^T)_{ii} (\mathbf{A} \mathbf{C} \mathbf{A}^T)_{jj} \\
&\quad + \frac{2\tilde{R}^4}{n^4} \sum_{i \neq j} (\mathbf{A} \mathbf{C} \mathbf{A}^T)_{ij}^2 - \frac{\tilde{R}^4}{n^2} [\text{tr}(\mathbf{C} \mathbf{H})]^2 \\
&= \frac{\tilde{C} - \tilde{R}^4}{n^4} \sum_i (\mathbf{A} \mathbf{C} \mathbf{A}^T)_{ii}^2 + \frac{\tilde{R}^4}{n^4} \left[\left(\sum_i (\mathbf{A} \mathbf{C} \mathbf{A}^T)_{ii}^2 + \sum_{i \neq j} (\mathbf{A} \mathbf{C} \mathbf{A}^T)_{ii} (\mathbf{A} \mathbf{C} \mathbf{A}^T)_{jj} \right) \right] \\
&\quad + \frac{2\tilde{R}^4}{n^4} \sum_{i \neq j} (\mathbf{A} \mathbf{C} \mathbf{A}^T)_{ij}^2 - \frac{\tilde{R}^4}{n^2} [\text{tr}(\mathbf{C} \mathbf{H})]^2 \\
&\leq \frac{\tilde{C} - \tilde{R}^4}{n^4} \left[\sum_i (\mathbf{A} \mathbf{C} \mathbf{A}^T)_{ii}^2 + \sum_{i \neq j} (\mathbf{A} \mathbf{C} \mathbf{A}^T)_{ij}^2 \right] = \frac{\tilde{C} - \tilde{R}^4}{n} \cdot \left[\frac{\text{tr}((\mathbf{C} \mathbf{H})^2)}{n} \right].
\end{aligned} \tag{61}$$

Here we can without loss of generality choose \tilde{C} so that $\tilde{C} > 3\tilde{R}^4$. Next, we compare (c) and (iii) in equation (59) and (58), respectively. We begin by expanding out (c) in equation (59) which yields

$$\mathbb{E} \left[\frac{2}{n^2} \cdot \mathbf{u}^T \mathbf{B} \mathbf{u} \cdot \boldsymbol{\eta}^T \mathbf{A} \mathbf{C} \mathbf{A}^T \boldsymbol{\eta} \mid \mathbf{H} \right] = \mathbb{E} \left[\frac{2}{n^2} \left(\sum_{i,j} u_i B_{ij} u_j \right) \left(\sum_{k,\ell} \eta_k (\mathbf{A} \mathbf{C} \mathbf{A}^T)_{k\ell} \eta_\ell \right) \mid \mathbf{H} \right].$$

The only terms which contribute are when $i = j$ and $k = \ell$. Therefore, we deduce the following

$$\begin{aligned}
\frac{2}{n^2} \cdot \mathbb{E} \left[\mathbf{u}^T \mathbf{B} \mathbf{u} \cdot \boldsymbol{\eta}^T \mathbf{A} \mathbf{C} \mathbf{A}^T \boldsymbol{\eta} \mid \mathbf{H} \right] &- 2 \left[\frac{R^2 \text{tr}(\mathbf{B})}{d} \right] \cdot \left[\frac{\tilde{R}^2 \text{tr}(\mathbf{C} \mathbf{H})}{n} \right] \\
&= \frac{2\tilde{R}^2 R^2}{n^2 d} \sum_{i,j} B_{ii} (\mathbf{A} \mathbf{C} \mathbf{A}^T)_{jj} - 2 \left[\frac{R^2 \text{tr}(\mathbf{B})}{d} \right] \cdot \left[\frac{\tilde{R}^2 \text{tr}(\mathbf{C} \mathbf{H})}{n} \right] \\
&= \frac{2\tilde{R}^2 R^2}{nd} [\text{tr}(\mathbf{B}) \text{tr}(\mathbf{C} \mathbf{H})] - 2 \left[\frac{R^2 \text{tr}(\mathbf{B})}{d} \right] \cdot \left[\frac{\tilde{R}^2 \text{tr}(\mathbf{C} \mathbf{H})}{n} \right] = 0.
\end{aligned} \tag{62}$$

We have now used up all the terms in (58) so the remaining terms, (d) and (e), in (59) we will show are themselves already going to 0 as $d \rightarrow \infty$. Again expanding the term (d), we get

$$\mathbb{E} \left[\frac{1}{n^2} (\mathbf{u}^T \mathbf{D} \mathbf{A}^T \boldsymbol{\eta})^2 \mid \mathbf{H} \right] = \mathbb{E} \left[\frac{1}{n^2} \left(\sum_{i,j} u_i (\mathbf{D} \mathbf{A}^T)_{ij} \eta_j \right)^2 \mid \mathbf{H} \right]. \tag{63}$$

By independence and isotropic variance of \mathbf{u} and $\boldsymbol{\eta}$, the only terms which remain after taking expectations are the ones with u_i^2 and η_j^2 terms. Therefore, we deduce

$$\begin{aligned}
\frac{1}{n^2} \mathbb{E} \left[(\mathbf{u}^T \mathbf{D} \mathbf{A}^T \boldsymbol{\eta})^2 \mid \mathbf{H} \right] &= \frac{1}{n^2} \sum_{i,j} \mathbb{E} [u_i^2] \cdot \mathbb{E} [\eta_j^2] \cdot (\mathbf{D} \mathbf{A}^T)_{ij}^2 = \frac{R^2 \tilde{R}^2}{n^2 d} \sum_{i,j} (\mathbf{D} \mathbf{A}^T)_{ij}^2 \\
&= \frac{R^2 \tilde{R}^2}{n} \cdot \left[\frac{\text{tr}(\mathbf{D}^2 \mathbf{H})}{d} \right].
\end{aligned} \tag{64}$$

The only term which remains in (59) is (e). Since $\mathbb{E}[\boldsymbol{\eta}] = \mathbf{0}$, the term $\mathbf{u}^T \mathbf{B} \mathbf{u} \cdot \mathbf{u}^T \frac{\mathbf{D} \mathbf{A}^T}{n} \boldsymbol{\eta}$ contributes nothing to the expectation. Similarly since $\mathbb{E}[\mathbf{u}] = \mathbf{0}$, the term $\frac{1}{n^3} \cdot \boldsymbol{\eta}^T \mathbf{A} \mathbf{C} \mathbf{A}^T \boldsymbol{\eta} \cdot \mathbf{u}^T \mathbf{C} \mathbf{A}^T \boldsymbol{\eta}$ is also zero in expectation.

Putting all the quantities (60), (61), (62), (64) together with (58) and (59), a straight forward application of Chebyshev's inequality yields

$$\begin{aligned}
\Pr \left(\left| y - \left[R^2 \text{tr}\left(\frac{\mathbf{B}}{d}\right) + \tilde{R}^2 \text{tr}\left(\frac{\mathbf{C} \mathbf{H}}{n}\right) \right] \right| > \varepsilon \mid \mathbf{H} \right) &\leq \frac{\text{Var}\left(y - \left[R^2 \text{tr}\left(\frac{\mathbf{B}}{d}\right) + \tilde{R}^2 \text{tr}\left(\frac{\mathbf{C} \mathbf{H}}{n}\right) \right] \mid \mathbf{H}\right)}{\varepsilon^2} \\
&\leq \frac{1}{\varepsilon^2} \left(\frac{C - R^4}{d} \text{tr}\left(\frac{\mathbf{B}^2}{d}\right) + \frac{\tilde{C} - \tilde{R}^4}{n} \text{tr}\left(\frac{(\mathbf{C} \mathbf{H})^2}{n}\right) + \frac{R^2 \tilde{R}^2}{n} \left[\frac{\text{tr}(\mathbf{D}^2 \mathbf{H})}{d} \right] \right).
\end{aligned}$$

□

Because $\|\nabla f(\mathbf{x}_k)\|^2$ and y only differ in the coefficients of the polynomial and these polynomials converge in probability (see Section 4.1), then intuitively, the proof of Proposition 4.3 says showed that conditioned on \mathbf{H} , the $\text{Var}(\|\nabla f(\mathbf{x}_k)\|^2|\mathbf{H})$ is $\mathcal{O}(\frac{1}{d})$ and

$$\mathbb{E}[\|\nabla f(\mathbf{x}_k)\|^2|\mathbf{H}] = R^2\text{tr}(\frac{\mathbf{B}}{d}) + \tilde{R}^2\text{tr}(\frac{\mathbf{C}\mathbf{H}}{n}). \quad (65)$$

Consequently, conditioned on \mathbf{H} , the squared norm of the gradient is roughly (65). So in view of this, it suffices to understand the expected traces of polynomials in \mathbf{H} . Random matrix theory studies convergence properties of the limiting distribution of high dimensional matrices, particularly the empirical spectral measure. An important tool derived from using Assumption 2 linking the ESM and the expected trace to the moments of the measure μ is given below.

Proposition 4.4 (Convergence of ESM). *Let \tilde{P}_k be any k -degree polynomial. Under Assumption 2, the following is true*

$$\frac{1}{d}\text{tr} \tilde{P}_k(\mathbf{H}) = \int \tilde{P}_k(\lambda; \lambda^\pm) d\mu_{\mathbf{H}} \xrightarrow[d \rightarrow \infty]{\text{Pr}} \int \tilde{P}_k(\lambda; \lambda^\pm) d\mu.$$

Proof. For sufficiently large d , Assumption 2 says $\Pr(\lambda_{\mathbf{H}}^+ > \lambda^+ + \hat{\varepsilon}) \leq \frac{\delta}{2}$. Define the event $\mathcal{S} = \{\lambda_{\mathbf{H}}^+ \leq \lambda^+ + \hat{\varepsilon}\}$. We construct a bounded, continuous function h by

$$h(\lambda) = \begin{cases} \tilde{P}_k(0; \lambda^\pm), & \text{if } \lambda < 0 \\ \tilde{P}_k(\lambda; \lambda^\pm), & \text{if } 0 \leq \lambda \leq \lambda^+ + \hat{\varepsilon} \\ \tilde{P}_k(\lambda^+ + \hat{\varepsilon}; \lambda^\pm), & \text{otherwise.} \end{cases}$$

Because the function h is bounded and continuous, Assumption 2 guarantees that

$$\Pr\left(\left|\int h(\lambda) d\mu_{\mathbf{H}} - \int h(\lambda) d\mu\right| > \varepsilon\right) \leq \delta. \quad (66)$$

Depending on whether \mathcal{S} has occurred, we have for all sufficiently large d

$$\begin{aligned} \Pr\left(\left|\int \tilde{P}_k(\lambda; \lambda^\pm) d\mu_{\mathbf{H}} - \int \tilde{P}_k(\lambda; \lambda^\pm) d\mu\right| > \varepsilon\right) \\ &= \Pr(\mathcal{S} \cap \{|\int \tilde{P}_k(\lambda; \lambda^\pm) d\mu_{\mathbf{H}} - \int \tilde{P}_k(\lambda; \lambda^\pm) d\mu| > \varepsilon\}) \\ &\quad + \Pr(\mathcal{S}^c \cap \{|\int \tilde{P}_k(\lambda; \lambda^\pm) d\mu_{\mathbf{H}} - \int \tilde{P}_k(\lambda; \lambda^\pm) d\mu| > \varepsilon\}) \\ &\leq \Pr(\mathcal{S} \cap \{|\int \tilde{P}_k(\lambda; \lambda^\pm) d\mu_{\mathbf{H}} - \int \tilde{P}_k(\lambda; \lambda^\pm) d\mu| > \varepsilon\}) + \frac{\delta}{2}. \end{aligned} \quad (67)$$

In the last line, the probability $\Pr(\mathcal{S}^c \cap \{|\int \tilde{P}_k(\lambda; \lambda^\pm) d\mu_{\mathbf{H}} - \int \tilde{P}_k(\lambda; \lambda^\pm) d\mu| > \varepsilon\}) \leq \Pr(\mathcal{S}^c) \leq \frac{\delta}{2}$ for large d . Hence, we consider only the first term in (67). By construction, for any element in \mathcal{S} it is clear that $h(\lambda) = \tilde{P}_k(\lambda)$. For sufficiently large d , equation (66) yields

$$\Pr(\mathcal{S} \cap \{|\int \tilde{P}_k(\lambda; \lambda^\pm) d\mu_{\mathbf{H}} - \int \tilde{P}_k(\lambda; \lambda^\pm) d\mu| > \varepsilon\}) \leq \Pr\left(\left|\int h(\lambda) d\mu_{\mathbf{H}} - \int h(\lambda) d\mu\right| > \varepsilon\right) \leq \frac{\delta}{2}.$$

The result follows after combining with (67). \square

Now that we have described the main components of our argument, we present a preliminary concentration result for the gradient.

Proposition 4.5. *Suppose the vectors $\mathbf{x}_0, \tilde{\mathbf{x}}$, and $\boldsymbol{\eta}$ and the matrix \mathbf{A} satisfy Assumptions 1 and 2 resp. The following holds*

$$\left| \|\nabla f(\mathbf{x}_k)\|^2 - \underbrace{\left(R^2 \frac{1}{d} \text{tr}(\mathbf{H}^2 P_k^2(\mathbf{H}; \lambda^\pm)) \right)}_{\text{signal}} + \underbrace{\left(\tilde{R}^2 \frac{1}{n} \text{tr}(\mathbf{H} P_k^2(\mathbf{H}; \lambda^\pm)) \right)}_{\text{noise}} \right| \xrightarrow{d \rightarrow \infty} 0. \quad (68)$$

Proof. Recall the definitions in (51) and (52) and equation (46). We note that the only difference between $\|\nabla f(\mathbf{x}_k)\|^2$ and y is that the coefficients of the polynomials in $\|\nabla f(\mathbf{x}_k)\|^2$ continuously depend on $\lambda_{\mathbf{H}}^\pm$ while the coefficients in y depend on λ^\pm . The polynomials P_k and Q_k together with Assumptions 1 and 2 ensure that all the conditions of Lemma 4.2 hold by setting \mathbf{w} and \mathbf{v} to combinations of \mathbf{u} and $\frac{1}{n} \mathbf{A}^T \boldsymbol{\eta}$ and the polynomials to \mathbf{B} , \mathbf{C} , and \mathbf{D} . Therefore we have $\|\nabla f(\mathbf{x}_k)\|^2 - y \xrightarrow{d \rightarrow \infty} 0$ so it suffices to prove (68) with $\|\nabla f(\mathbf{x}_k)\|^2$ replaced by y .

Fix constants $\varepsilon, \delta > 0$. Proposition 4.4 guarantees convergence in probability of any expected trace to a constant which depends on the polynomial and the deterministic measure μ . This together with the definitions of \mathbf{B} , \mathbf{C} , and \mathbf{D} yield for sufficiently large d

$$\begin{aligned} \Pr \left(\left| \frac{\text{tr}(\mathbf{B}^2)}{d} \right| > M_1 \stackrel{\text{def}}{=} \varepsilon + \int \lambda^4 P_k^4(\lambda; \lambda^\pm) d\mu \right) &\leq \frac{\delta}{6}, \\ \Pr \left(\left| \frac{\text{tr}((\mathbf{C}\mathbf{H})^2)}{n} \right| > M_2 \stackrel{\text{def}}{=} \varepsilon + r \int \lambda^2 P_k^4(\lambda; \lambda^\pm) d\mu \right) &\leq \frac{\delta}{6}, \\ \text{and } \Pr \left(\left| \frac{\text{tr}(\mathbf{D}^2 \mathbf{H})}{d} \right| > M_3 \stackrel{\text{def}}{=} \varepsilon + 4 \int \lambda^3 P_k^4(\lambda; \lambda^\pm) d\mu \right) &\leq \frac{\delta}{6}. \end{aligned} \quad (69)$$

We define the set \mathcal{S} for which the expected traces of the random matrices are bounded, namely,

$$\mathcal{S} = \left\{ \left| \frac{\text{tr}(\mathbf{B}^2)}{d} \right| \leq M_1 \right\} \cap \left\{ \left| \frac{\text{tr}((\mathbf{C}\mathbf{H})^2)}{n} \right| \leq M_2 \right\} \cap \left\{ \left| \frac{\text{tr}(\mathbf{D}^2 \mathbf{H})}{d} \right| \leq M_3 \right\},$$

and we observe because of (69) that the probability $\Pr(\mathcal{S}^c) \leq \frac{\delta}{2}$. The total law of probability yields the following

$$\begin{aligned} \Pr \left(|y - [R^2 \text{tr}(\frac{\mathbf{B}}{d}) + \tilde{R}^2 \text{tr}(\frac{\mathbf{C}\mathbf{H}}{n})]| > \varepsilon \right) &= \Pr \left(\mathcal{S} \cap \{|y - [R^2 \text{tr}(\frac{\mathbf{B}}{d}) + \tilde{R}^2 \text{tr}(\frac{\mathbf{C}\mathbf{H}}{n})]| > \varepsilon\} \right) \\ &\quad + \Pr \left(\mathcal{S}^c \cap \{|y - [R^2 \text{tr}(\frac{\mathbf{B}}{d}) + \tilde{R}^2 \text{tr}(\frac{\mathbf{C}\mathbf{H}}{n})]| > \varepsilon\} \right) \\ &\leq \Pr \left(\mathcal{S} \cap \{|y - [R^2 \text{tr}(\frac{\mathbf{B}}{d}) + \tilde{R}^2 \text{tr}(\frac{\mathbf{C}\mathbf{H}}{n})]| > \varepsilon\} \right) + \frac{\delta}{2}. \end{aligned} \quad (70)$$

Hence it suffices to bound the first term in (70). The idea is to condition on the matrix \mathbf{H} and apply Proposition 4.3. The law of total expectation yields

$$\begin{aligned} &\Pr \left(\mathcal{S} \cap \{|y - [\text{tr}(\frac{R^2 \mathbf{B}}{d}) + \text{tr}(\frac{\tilde{R}^2 \mathbf{C}\mathbf{H}}{n})]| > \varepsilon\} \right) \\ &(\text{conditioned on } \mathbf{H}) = \mathbb{E} \left[1_{\mathcal{S}} \Pr \left(|y - [\text{tr}(\frac{R^2 \mathbf{B}}{d}) + \text{tr}(\frac{\tilde{R}^2 \mathbf{C}\mathbf{H}}{n})]| > \varepsilon \mid \mathbf{H} \right) \right] \\ &(\text{Proposition 4.3}) \leq \frac{1}{\varepsilon^2} \mathbb{E} \left[1_{\mathcal{S}} \left(\frac{C-R^4}{d} \text{tr}(\frac{\mathbf{B}^2}{d}) + \frac{\tilde{C}-\tilde{R}^4}{n} \text{tr}(\frac{(\mathbf{C}\mathbf{H})^2}{n}) + \frac{R^2 \tilde{R}^2}{n} [\text{tr}(\frac{\mathbf{D}^2 \mathbf{H}}{d})] \right) \right]. \end{aligned} \quad (71)$$

Here for the indicator of the event \mathcal{S} we use the notation $1_{\mathcal{S}}(\omega)$ where the indicator is 1 if $\omega \in \mathcal{S}$ and 0 otherwise. By construction of the event \mathcal{S} , each of the expected traces in (71) are bounded and therefore, we deduce that

$$\Pr \left(\mathcal{S} \cap \{|y - [\text{tr}(\frac{R^2 \mathbf{B}}{d}) + \text{tr}(\frac{\tilde{R}^2 \mathbf{C}\mathbf{H}}{n})]| > \varepsilon\} \right) = \frac{1}{\varepsilon^2} \cdot \mathcal{O}\left(\frac{1}{d}\right).$$

By choosing d sufficiently large, we can make the right hand side smaller than $\frac{\delta}{2}$. The result immediately follows from (70). \square

Proposition 4.5 reveals that for high-dimensional data the squared norm of the gradient $\|\nabla f(\mathbf{x}_k)\|^2$ is a polynomial in the eigenvalues of the matrix \mathbf{H} . Every eigenvalue, not just the largest or smallest, appears in this formula (68). This means that first-order methods indeed see all of the eigenvalues of the matrix \mathbf{H} , not just the top or bottom one. However, the expected trace is still a random quantity due to its dependency on the random matrix. We remove this randomness and complete the proof of Theorem 1.1 after noting that the moments of the empirical spectral measure converge in probability to a deterministic quantity, denoted by $\mathcal{E}_{d \rightarrow \infty} [\|\nabla f(\mathbf{x}_k)\|^2]$.

Proof of Theorem 1.1. Propositions 4.4 and 4.5 yield the result. \square

4.3 Halting time converges to a constant

The concentration of the norm of the gradient in (50) gives a candidate for the limiting value of the halting time T_ε . More precisely, we define this candidate for the halting time τ_ε from $\mathcal{E}_{d \rightarrow \infty} [\|\nabla f(\mathbf{x}_k)\|^2]$ and we recall the halting time, T_ε , as

$$\tau_\varepsilon \stackrel{\text{def}}{=} \inf \{k > 0 : \mathcal{E}_{d \rightarrow \infty} [\|\nabla f(\mathbf{x}_k)\|^2] \leq \varepsilon\} \quad \text{and} \quad T_\varepsilon \stackrel{\text{def}}{=} \inf \{k > 0 : \|\nabla f(\mathbf{x}_k)\|^2 \leq \varepsilon\}. \quad (72)$$

We note that the deterministic value τ_ε is, by definition, the average complexity of GD whereas T_ε is a random variable depending on randomness from the data, noise, signal, and initialization. This leads to our main result that states the almost sure convergence of the halting time to a constant value. We begin by showing that τ_ε is well-defined.

Lemma 4.3 (τ_ε is well-defined). *Under the assumptions of Proposition 4.5, the iterates of a convergent algorithm satisfy $\mathcal{E}_{d \rightarrow \infty} [\|\nabla f(\mathbf{x}_k)\|^2] \xrightarrow{k \rightarrow \infty} 0$.*

Proof. Because the iterates are generated from a convergent algorithm, both $\lambda^2 P_k^2(\lambda; \lambda^\pm) \rightarrow 0$ and $\lambda P_k^2(\lambda; \lambda^\pm) \rightarrow 0$ and these polynomials are uniformly bounded in k for each $\lambda \in [\lambda^-, \lambda^+]$ (see Lemma 3.1 and 3.2). By dominated convergence theorem, the result follows. \square

With our candidate for the limiting halting time τ_ε well-defined, we show that number of iterations until $\|\nabla f(\mathbf{x}_k)\|^2 \leq \varepsilon$ equals τ_ε for high-dimensional data. We state a more general result of Theorem 1.2.

Theorem 4.1 (Halting time universality). *Suppose that $\mathcal{E}_{d \rightarrow \infty} [\|\nabla f(\mathbf{x}_k)\|^2] > \varepsilon$ and $\|\nabla f(\mathbf{x}_0)\|^2 > \varepsilon$. Provided that $\mathcal{E}_{d \rightarrow \infty} [\|\nabla f(\mathbf{x}_k)\|^2] \neq \varepsilon$ for all k , the probability of reaching ε in a pre-determined number of steps satisfies*

$$\lim_{d \rightarrow \infty} \Pr(T_\varepsilon = \tau_\varepsilon) = 1.$$

If the constant $\varepsilon = \mathcal{E}_{d \rightarrow \infty} [\|\nabla f(\mathbf{x}_k)\|^2]$ for some k , then the following holds

$$\lim_{d \rightarrow \infty} \Pr(T_\varepsilon \in [\tau_\varepsilon, \tau_\varepsilon + M_\varepsilon]) = 1, \quad \text{where } M_\varepsilon \stackrel{\text{def}}{=} \inf \{k - \tau_\varepsilon > 0 \mid \xi_k < \varepsilon\}.$$

Proof. To simplify notation, we define $\xi_k \stackrel{\text{def}}{=} \mathcal{E}_{d \rightarrow \infty} [\|\nabla f(\mathbf{x}_k)\|^2]$. First, we consider the case where $\varepsilon \neq \xi_k$ for all k . We are interested in bounding the following probabilities

$$\Pr(T_\varepsilon \neq \tau_\varepsilon) = \Pr(T_\varepsilon < \tau_\varepsilon) + \Pr(T_\varepsilon > \tau_\varepsilon). \quad (73)$$

We bound each of these probabilities independently; first consider $\Pr(T_\varepsilon < \tau_\varepsilon)$ in (73). For $\tau_\varepsilon = 1$, we note that $\Pr(T_\varepsilon < \tau_\varepsilon) = 0$ since $T_\varepsilon \geq 1$. So we can assume that $\tau_\varepsilon > 1$. Since $T_\varepsilon \leq \tau_\varepsilon - 1$, we obtain

$$\Pr(T_\varepsilon < \tau_\varepsilon) = \Pr\left(\bigcup_{k=0}^{\tau_\varepsilon-1} \{T_\varepsilon = k\}\right) \leq \sum_{k=0}^{\tau_\varepsilon-1} \Pr(T_\varepsilon = k) \leq \sum_{k=0}^{\tau_\varepsilon-1} \Pr(\|\nabla f(\mathbf{x}_k)\|^2 \leq \varepsilon). \quad (74)$$

Now we bound the probabilities $\Pr(\|\nabla f(\mathbf{x}_k)\|^2 \leq \varepsilon)$. As τ_ε is the first time ξ falls below ε , we conclude that $\xi_{\tau_\varepsilon} < \varepsilon < \xi_{\tau_\varepsilon-1}, \xi_{\tau_\varepsilon-2}, \dots, \xi_0$ where we used that $\varepsilon \neq \xi_k$ for any k . Next we define the constant $0 < \delta \stackrel{\text{def}}{=} \min_{0 \leq k \leq \tau_\varepsilon} \{|\varepsilon - \xi_k|\}$ and we observe that $\delta < |\varepsilon - \xi_k| = \xi_k - \varepsilon$ for all $k < \tau_\varepsilon$. Fix a constant $\hat{\varepsilon} > 0$ and index k . Theorem 1.1 says that by making $d(k)$ sufficiently large

$$\Pr(\|\nabla f(\mathbf{x}_k)\|^2 \leq \varepsilon) \leq \Pr(\|\nabla f(\mathbf{x}_k)\|^2 < \xi_k - \delta) \leq \frac{\hat{\varepsilon}}{\tau_\varepsilon - 1}.$$

Here we used that τ_ε is finite for every $\varepsilon > 0$ (Lemma 4.3). Set $D \stackrel{\text{def}}{=} \max\{d(1), d(2), \dots, d(\tau_\varepsilon - 1)\}$. Then for all $d > D$, we have from (74) the following

$$\Pr(T_\varepsilon < \tau_\varepsilon) \leq \sum_{k=0}^{\tau_\varepsilon-1} \Pr(\|\nabla f(\mathbf{x}_k)\|^2 \leq \varepsilon) \leq \sum_{k=0}^{\tau_\varepsilon-1} \frac{\hat{\varepsilon}}{\tau_\varepsilon - 1} = \hat{\varepsilon}.$$

Lastly, we bound $\Pr(T_\varepsilon > \tau_\varepsilon)$. The idea is similar to the other direction. Let δ be defined as above. Therefore, again by Theorem 1.1, we conclude for sufficiently large d

$$\Pr(T_\varepsilon > \tau_\varepsilon) \leq \Pr(\|\nabla f(\mathbf{x}_{\tau_\varepsilon})\|^2 > \varepsilon) \leq \Pr(\|\nabla f(\mathbf{x}_{\tau_\varepsilon})\|^2 - \xi_{\tau_\varepsilon} > \delta) \rightarrow 0.$$

Indeed, we used that $\xi_{\tau_\varepsilon} < \varepsilon$ and $\delta < |\varepsilon - \xi_{\tau_\varepsilon}| = \varepsilon - \xi_{\tau_\varepsilon}$. This completes the proof when $\varepsilon \neq \xi_k$.

Next, we consider the second case where $\xi_k = \varepsilon$. Note that $M_\varepsilon < \infty$ for all ε because $\lim_{k \rightarrow \infty} \xi_k = 0$. In this setting, we are interested in bounding

$$\Pr(T_\varepsilon \notin [\tau_\varepsilon, \tau_\varepsilon + M_\varepsilon]) = \Pr(T_\varepsilon < \tau_\varepsilon) + \Pr(T_\varepsilon > \tau_\varepsilon + M_\varepsilon).$$

The arguments will be similar to the previous setting. Replacing the definition of δ above with $\delta \stackrel{\text{def}}{=} \min_{0 \leq k \leq \tau_\varepsilon-1} \{|\varepsilon - \xi_k|\}$ yields that $\delta > 0$ since $\varepsilon < \xi_{\tau_\varepsilon-1}, \xi_{\tau_\varepsilon-2}, \dots, \xi_0$. With this choice of δ , the previous argument holds and we deduce that $\Pr(T_\varepsilon < \tau_\varepsilon) \rightarrow 0$. Next we show that $\Pr(T_\varepsilon > \tau_\varepsilon + M_\varepsilon)$. As before, we know that $\Pr(T_\varepsilon > \tau_\varepsilon + M_\varepsilon) \leq \Pr(\|\nabla f(\mathbf{x}_{\tau_\varepsilon+M_\varepsilon})\|^2 > \varepsilon)$. By definition of M_ε , we have that $\varepsilon > \xi_{\tau_\varepsilon+M_\varepsilon}$. Now define $\delta \stackrel{\text{def}}{=} \varepsilon - \xi_{\tau_\varepsilon+M_\varepsilon} > 0$. The previous argument holds with this choice of δ ; therefore, one has that $\Pr(T_\varepsilon > \tau_\varepsilon + M_\varepsilon) \rightarrow 0$. \square

For large models the number of iterations to reach a nearly optimal point equals its average complexity which loosely says $T_\varepsilon = \tau_\varepsilon$. The variability in the halting time goes to zero. Since the dependence in τ_ε on the distribution of the data is limited to only the first two moments, almost all instances of high-dimensional data have the same limit. In Tables 1 and 2, we compute the value of $\mathcal{E}_{d \rightarrow \infty} [\|\nabla f(\mathbf{x}_k)\|^2]$ for various models.

5 Derivation of the worst and average-case complexity

In this section, we derive an expression for the average-case complexity in the isotropic features model. Here the empirical spectral measure $\mu_{\mathbf{H}}$ converges to the Marčenko-Pastur measure μ_{MP} (2). The average-case complexity, τ_ε , is controlled by the value of the expected gradient norm in (50). Hence to analyze the average-case rate, it suffices to derive an expression for this value, $\mathcal{E}_{d \rightarrow \infty} [\|\nabla f(\mathbf{x}_k)\|^2]$.

In light of (50), we must integrate the residual polynomials in Table 3 against the Marčenko-Pastur measure. By combining Theorem 1.1 with the integrals derived in Appendix B, we obtain the average-case complexities. Apart from Nesterov's accelerated method (convex), an *exact* formula for the average-case rates are obtained. In the convex setting for Nesterov, the integral is difficult to directly compute so instead we use the asymptotic polynomial in (32). Hence for Nesterov's

accelerated method (convex), we only get an asymptotic average-case rate for sufficiently large k (see Appendix B). Tables 1 and 2 summarize the asymptotic rates where both iteration and problem size are large.

We now turn to the worst-case guarantees. We discuss below how to make the worst-case complexity comparable.

5.1 Traditional worst-case complexity

Recall, the prior discussion on the dimension-dependent constants in the typical worst-case complexity bounds. We now make this precise below.

Worst-case complexity: strongly convex and noiseless non-strongly convex regimes. Consider GD and note the other methods will have similar analysis. Recall, the standard analytical worst-case bound for the strongly convex regime and the exact worst-case bound for the non-strongly convex setting [Taylor et al., 2017], respectively,

$$\begin{aligned} \|\nabla f(\mathbf{x}_k)\|^2 &\leq (\lambda_{\mathbf{H}}^+)^2 \|\mathbf{x}_0 - \mathbf{x}^*\|^2 \left(1 - \frac{\lambda_{\mathbf{H}}^-}{\lambda_{\mathbf{H}}^+}\right)^{2k} \quad (\text{strongly convex}) \\ \text{and } \|\nabla f(\mathbf{x}_k)\|^2 &\leq \frac{(\lambda_{\mathbf{H}}^+)^2 \|\mathbf{x}_0 - \mathbf{x}^*\|^2}{(k+1)^2} \quad (\text{convex}) \end{aligned}$$

where \mathbf{x}^* is the minimal norm solution of (8). For sufficiently large d , the largest $\lambda_{\mathbf{H}}^+$ and smallest eigenvalues $\lambda_{\mathbf{H}}^-$ of \mathbf{H} converge in probability to $\sigma^2(1 + \sqrt{r})^2$ and $\sigma^2(1 - \sqrt{r})^2$ respectively. These are the top and bottom edge of the Marčenko-Pastur distribution. We also note in the noiseless setting $\|\mathbf{x}_0 - \mathbf{x}^*\|^2 = \|\mathbf{x}_0 - \tilde{\mathbf{x}}\|^2$. Hence by Assumption 1 and $\tilde{R}^2 = 0$, on average, $\|\mathbf{x}_0 - \tilde{\mathbf{x}}\|^2 = R^2$. Moreover when the matrix \mathbf{H} is nonsingular, as in the strongly convex setting, the optimum $\|\mathbf{x}^*\|^2$ does not grow as dimension increases despite the noise. As such, $\|\mathbf{x}_0 - \mathbf{x}^*\|^2$, on average, equals \tilde{R}^2 . From these observations we derive the worst-case complexities.

Worst-case complexity: noisy non-strongly convex regime. While discussing the worst-case complexity in Section 1.2, we noted a discrepancy in the noisy, non-strongly convex regime between the average rate and the exact worst complexity. For instance, the exact worst complexity for gradient descent (GD) [Taylor et al., 2017] is

$$\|\nabla f(\mathbf{x}_k)\|^2 \leq \frac{(\lambda_{\mathbf{H}}^+)^2 \|\mathbf{x}_0 - \mathbf{x}^*\|^2}{(k+1)^2} \quad \text{where } \mathbf{x}^* \text{ is the optimum of (8).} \quad (75)$$

For sufficiently large d , the largest eigenvalue of $\lambda_{\mathbf{H}}^+$ converges a.s. to $\lambda^+ = 16\sigma^2$, the top edge of the support of μ_{MP} . Hence, to derive worst-case complexity bounds, it suffices to understand the behavior of the distance to the optimum.

The vectors \mathbf{x}^* and $\tilde{\mathbf{x}}$ are different when noise is added to the signal. For simplicity, we consider the setting where the matrix \mathbf{A} is invertible. Intuitively, the optimum $\mathbf{x}^* \approx \mathbf{A}^{-1}\mathbf{b} = \tilde{\mathbf{x}} + \mathbf{A}^{-1}\boldsymbol{\eta}$ where $\tilde{\mathbf{x}}$ is the underlying random signal. Because the signal $\tilde{\mathbf{x}}$ is scaled, Assumption 1 says $\mathbb{E}[\|\mathbf{x}_0 - \tilde{\mathbf{x}}\|^2] = R^2$. Therefore, the distance to the optimum $\mathbf{x}_0 - \mathbf{x}^*$ is controlled by the noise which in turn is bounded by the reciprocal of the minimum eigenvalue of $\mathbf{A}^T\mathbf{A}$, namely

$$\|\mathbf{x}_0 - \mathbf{x}^*\|^2 \approx \|\mathbf{x}_0 - \tilde{\mathbf{x}}\|^2 + \|\mathbf{A}^{-1}\boldsymbol{\eta}\|^2 \geq \frac{|\mathbf{u}_{\min}^T \boldsymbol{\eta}|^2}{\lambda_{\min}(\mathbf{A}^T\mathbf{A})},$$

where $(\lambda_{\min}(\mathbf{A}^T\mathbf{A}), \mathbf{u}_{\min})$ is an eigenvalue-eigenvector pair corresponding to the minimum eigenvalue of $\mathbf{A}^T\mathbf{A}$. Unfortunately, the smallest eigenvalue is not well-behaving. Particularly there does not exist any scaling so that expectation of $\lambda_{\min}(\mathbf{A}^T\mathbf{A})^{-1}$ is finite and the distribution is heavy-tailed.

Instead we show that this quantity $\frac{|\mathbf{u}_{\min}^T \boldsymbol{\eta}|^2}{\lambda_{\min}(\mathbf{A}^T \mathbf{A})}$ grows faster than $\tilde{R}^2 d$. To do so, we assume that all moments of the entries of the matrix \mathbf{A} are bounded, namely,

$$\mathbb{E}[|A_{ij}|^k] < \infty \quad \text{for all } k \leq 10^4. \quad (76)$$

This bounded moment assumption is a mild assumption on the entries. For instance it includes any sub-exponential random variables. It should be noted here that under the simple isotropic features model it is clear that $\|\mathbf{x}_0 - \mathbf{x}^*\|^2$ is *dimension-dependent*, but the exact dependence is more complicated. Under this condition (76), we can prove a bound, which gives the dependence on the problem size, for the growth rate of the distance to the optimum.

Lemma 5.1 (Growth of $\|\mathbf{x}_0 - \mathbf{x}^*\|^2$). *Suppose Assumptions 1 and 2 hold such that the noise vector $\boldsymbol{\eta} \in \mathbb{R}^d$ and the entries of the data matrix $\mathbf{A} \in \mathbb{R}^{d \times d}$ satisfy bounded moments (76). Let \mathbf{x}^* be the minimal norm solution to (8). For any $\delta > 0$ there exists a constant $M_\delta > 0$ such that*

$$\liminf_{n \rightarrow \infty} \Pr(\|\mathbf{x}_0 - \mathbf{x}^*\|^2 \geq d \cdot \tilde{R}^2 M_\delta) \geq 1 - \delta. \quad (77)$$

Proof. We begin by defining the constant $M_\delta > 0$. The $n \times n$ matrix \mathbf{A} is invertible a.s. so without loss of generality the smallest eigenvalue of $\mathbf{A}^T \mathbf{A}$ is non-zero. Here the dimensions are equal, $d = n$. From [Edelman, 1988, Corollary 3.1] and [Tao and Vu, 2010, Theorem 1.3], we know that $n\lambda_{\min}(\mathbf{A}^T \mathbf{A})$ converges in distribution where we denote $\lambda_{\min}(\mathbf{A}^T \mathbf{A})$ to be the smallest eigenvalue of $\mathbf{A}^T \mathbf{A}$. It is immediately clear that $\log(n\lambda_{\min}(\mathbf{A}^T \mathbf{A}))$ also converges in distribution. By Theorem 3.2.7 in Durrett [2010], the sequence of distribution functions $\{F_n(x) = \Pr(\log(n\lambda_{\min}(\mathbf{A}^T \mathbf{A})) \leq x)\}$ is tight, that is, there exists an $C_\delta > 0$ such that

$$\limsup_{n \rightarrow \infty} \Pr(n\lambda_{\min}(\mathbf{A}^T \mathbf{A}) \notin (e^{-C_\delta}, e^{C_\delta})) = \limsup_{n \rightarrow \infty} 1 - F_n(C_\delta) + F_n(-C_\delta) \leq \frac{\delta}{2}.$$

In particular, we know that

$$\limsup_{n \rightarrow \infty} \Pr(n^{-1}(\lambda_{\min}(\mathbf{A}^T \mathbf{A}))^{-1} < e^{-C_\delta}) \leq \frac{\delta}{2}. \quad (78)$$

For any χ_1^2 -squared with 1-degree of freedom random variable of X , there exists a constant $\hat{C}_\delta > 0$ such that

$$\Pr(X \leq \hat{C}_\delta) \leq \frac{\delta}{2}. \quad (79)$$

Let $M_\delta \stackrel{\text{def}}{=} \frac{1}{4} \min\{e^{-2C_\delta}, \hat{C}_\delta^2\}$. With M_δ defined, we are now ready to prove (77). The matrix \mathbf{A} is a.s. invertible so gradient descent converges to $\mathbf{x}^* = \mathbf{A}^{-1} \mathbf{b}$. Next we observe that (77) is equivalent to solving

$$\limsup_{n \rightarrow \infty} \Pr(\|\mathbf{x}_0 - \mathbf{A}^{-1} \mathbf{b}\| < \tilde{R} \sqrt{nM_\delta}) \leq \delta. \quad (80)$$

Plugging in the value of \mathbf{b} and using the reverse triangle inequality, we obtain

$$\Pr(\|\mathbf{x}_0 - \mathbf{A}^{-1} \mathbf{b}\| < \tilde{R} \sqrt{nM_\delta}) \leq \Pr(\|\mathbf{A}^{-1} \boldsymbol{\eta}\| < \tilde{R} \sqrt{nM_\delta} + \|\mathbf{x}_0 - \tilde{\mathbf{x}}\|).$$

Using Markov's inequality, we can obtain a bound on $\|\mathbf{x}_0 - \tilde{\mathbf{x}}\|$:

$$\Pr(\|\mathbf{x}_0 - \tilde{\mathbf{x}}\| \geq \tilde{R} \sqrt{nM_\delta}) \leq \frac{R^2}{nM_\delta \tilde{R}^2}.$$

Consider now the event given by $\mathcal{S} \stackrel{\text{def}}{=} \{\|\mathbf{x}_0 - \tilde{\mathbf{x}}\| < \tilde{R} \sqrt{nM_\delta}\}$. The total law of probability yields

$$\begin{aligned} & \Pr(\|\mathbf{x}_0 - \mathbf{A}^{-1} \mathbf{b}\| < \tilde{R} \sqrt{nM_\delta}) \\ & \leq \Pr(\mathcal{S}^c) + \Pr(\mathcal{S} \cap \{\|\mathbf{A}^{-1} \boldsymbol{\eta}\| < \tilde{R} \sqrt{nM_\delta} + \|\mathbf{x}_0 - \tilde{\mathbf{x}}\|\}) \\ & \leq \frac{R^2}{nM_\delta \tilde{R}^2} + \Pr(\|\mathbf{A}^{-1} \boldsymbol{\eta}\| < 2\tilde{R} \sqrt{nM_\delta}) = \frac{R^2}{nM_\delta \tilde{R}^2} + \Pr(\|\mathbf{A}^{-1} \boldsymbol{\eta}\|^2 < 4n\tilde{R}^2 M_\delta). \end{aligned} \quad (81)$$

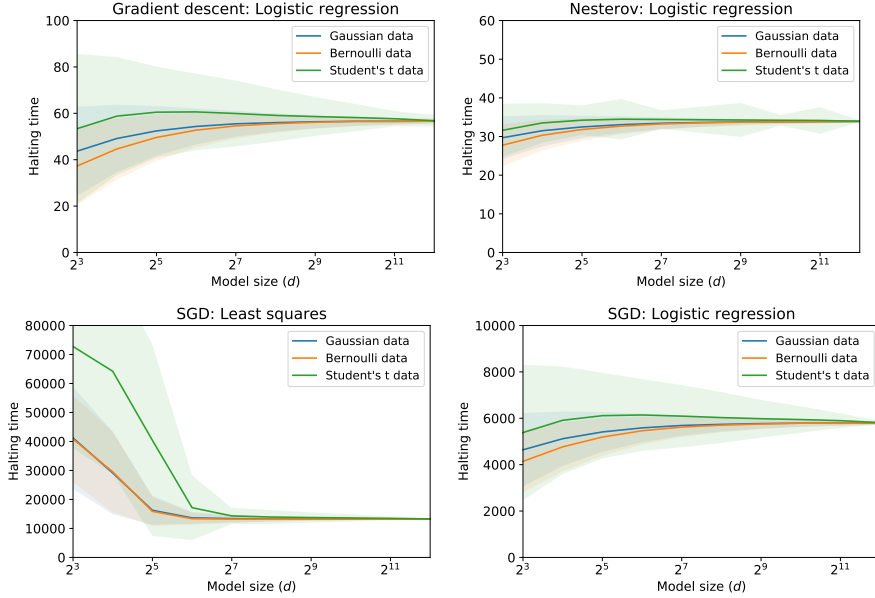


Figure 7: **Halting time universality beyond least squares.** We compute the halting time on algorithms and models not covered by our theory and note that the convergence to a deterministic and universal halting time is also empirically observed in these settings. For different model size d (x -axis) we sample the vectors $\tilde{\mathbf{x}}, \mathbf{x}_0$ and the matrix \mathbf{A} ($\tilde{R}^2 = 0.01$ and $r = 0.5$) and report the halting time (y -axis) and its standard deviation (shaded area) for GD and Nesterov (convex) on logistic regression and SGD on both least squares and logistic regression.

A simple calculation gives that $n^{-1}\tilde{R}^{-2}\|\mathbf{A}^{-1}\boldsymbol{\eta}\|^2 \geq n^{-1}(\lambda_{\min}(\mathbf{A}^T\mathbf{A}))^{-1}\tilde{R}^{-2}(\mathbf{u}_{\min}^T\boldsymbol{\eta})^2$ where the orthonormal vector \mathbf{u}_{\min} is the eigenvector associated with the eigenvalue $(\lambda_{\min}(\mathbf{A}^T\mathbf{A}))^{-1}$. From this, we deduce the following inequalities

$$\begin{aligned} \Pr(\|\mathbf{A}^{-1}\boldsymbol{\eta}\|^2 < 4n\tilde{R}^2M_\delta) &\leq \Pr(n^{-1}\lambda_{\min}(\mathbf{A}^T\mathbf{A})^{-1} \cdot \tilde{R}^{-2}(\mathbf{u}_{\min}^T\boldsymbol{\eta})^2 < \min\{e^{-2C_\delta}, \hat{C}_\delta^2\}) \\ &\leq \Pr(n^{-1}\lambda_{\min}(\mathbf{A}^T\mathbf{A})^{-1} < \min\{e^{-C_\delta}, \hat{C}_\delta\}) + \Pr(\tilde{R}^{-2}(\mathbf{u}_{\min}^T\boldsymbol{\eta})^2 < \min\{e^{-C_\delta}, \hat{C}_\delta\}) \\ &\leq \Pr(n^{-1}\lambda_{\min}(\mathbf{A}^T\mathbf{A})^{-1} < e^{-C_\delta}) + \Pr(\tilde{R}^{-2}(\mathbf{u}_{\min}^T\boldsymbol{\eta})^2 < \hat{C}_\delta). \end{aligned}$$

Since $\boldsymbol{\eta}$ is Gaussian and \mathbf{u}_{\min} is orthonormal, we know that $\tilde{R}^{-2}(\mathbf{u}_{\min}^T\boldsymbol{\eta})^2 \sim \chi_1^2$, a chi-squared distribution, so (79) holds and we already showed that $n^{-1}(\lambda_{\min}(\mathbf{A}^T\mathbf{A}))^{-1}$ satisfies (78). By taking limsup, we have

$$\limsup_{n \rightarrow \infty} \Pr(\|\mathbf{A}^{-1}\boldsymbol{\eta}\|^2 < 4n\tilde{R}^2M_\delta) \leq \delta.$$

The inequality in (80) immediately follows after taking the limsup of (81). \square

Combining this lemma with the equation (75), we get with high probability that

$$\|\nabla f(\mathbf{x}_k)\|^2 \leq \frac{(\lambda_{\mathbf{H}}^+)^2 \|\mathbf{x}_0 - \mathbf{x}^*\|^2}{(k+1)^2} \approx \frac{16\sigma^2\tilde{R}^2d}{(k+1)^2}.$$

By setting the right-hand side equal to ε , we get the worst-case complexity result.

5.2 Adversarial Model

Next we recall the adversarial model. Here we assume a noisy generative model for \mathbf{b} (Assumption 1). Then our adversary chooses the matrix \mathbf{A} without knowledge of \mathbf{b} in such a way that *maximizes the*

norm of the gradient subject to the constraint that the convex hull of the eigenvalues of $\mathbf{H} = \frac{1}{n} \mathbf{A}^T \mathbf{A}$ equals $[\lambda^-, \lambda^+]$. For comparison to the average-case analysis with isotropic features, we would choose λ^\pm to be the endpoints of the Marčenko-Pastur law. In light of Proposition 4.3, the adversarial model seeks to solve the constrained optimization problem

$$\begin{aligned} \max_{\mathbf{H}} \left\{ \mathbb{E} [\|\nabla f(\mathbf{x}_k)\|^2] = \frac{R^2}{d} \text{tr}(\mathbf{H}^2 P_k^2(\mathbf{H}; \lambda_{\mathbf{H}}^\pm)) + \frac{\tilde{R}^2}{n} \text{tr}(\mathbf{H} P_k^2(\mathbf{H}; \lambda_{\mathbf{H}}^\pm)) \right\} \\ \text{subject to } \lambda_{\mathbf{H}}^+ = \lambda^+ \text{ and } \lambda_{\mathbf{H}}^- = \lambda^-, \end{aligned} \quad (82)$$

where the largest (smallest) eigenvalue of \mathbf{H} is restricted to the upper (lower) edge of Marčenko-Pastur measure. The optimal \mathbf{H} of (82), \mathbf{H}_{\max} , has all but two of its eigenvalues at

$$\lambda_k^* \stackrel{\text{def}}{=} \arg \max_{\lambda \in [\lambda^-, \lambda^+]} \left\{ R^2 \lambda^2 P_k^2(\lambda; \lambda^\pm) + \tilde{R}^2 \lambda P_k^2(\lambda; \lambda^\pm) \right\}. \quad (83)$$

The other two eigenvalues must live at λ^+ and λ^- in order to satisfy the constraints. The empirical spectral measure for this \mathbf{H}_{\max} is exactly

$$\mu_{\mathbf{H}_{\max}} = \frac{1}{d} \sum_{i=1}^d \delta_{\lambda_i} = \frac{1}{d} \cdot \delta_{\lambda^+} + \frac{1}{d} \cdot \delta_{\lambda^-} + \left(1 - \frac{2}{d}\right) \cdot \delta_{\lambda_k^*}.$$

Since this empirical spectral measure weakly converges to $\delta_{\lambda_k^*}$, we satisfy the conditions of Assumption 1 for these \mathbf{H}_{\max} and spectral measure $\mu_{\mathbf{H}_{\max}}$. Hence, Theorem 1.1 holds and the maximum expected squared norm of the gradient as the model size goes to infinity equals

$$\begin{aligned} \lim_{d \rightarrow \infty} \max_{\mathbf{H}} \mathbb{E} [\|\nabla f(\mathbf{x}_k)\|^2] &= \int [R^2 \lambda^2 P_k^2(\lambda; \lambda^\pm) + \tilde{R}^2 \lambda P_k^2(\lambda; \lambda^\pm)] \delta_{\lambda_k^*} \\ &= \max_{\lambda \in [\lambda^-, \lambda^+]} R^2 \lambda^2 P_k^2(\lambda; \lambda^\pm) + \tilde{R}^2 \lambda P_k^2(\lambda; \lambda^\pm). \end{aligned} \quad (84)$$

We called the above expression the *adversarial average-case complexity*. Table 1 shows these convergence guarantees. We defer the derivations to Appendix C.

Remark 5.1. *In the strongly convex setting, we omitted the adversarial average-case guarantees out of brevity. For all the algorithms, the value of λ_k^* occurs near or at the minimum eigenvalue λ^+ . As such there is (almost) no distinction between the traditional worst-case guarantees and the adversarial guarantees.*

6 Numerical Simulations

To illustrate our theoretical results we report simulations using gradient descent (GD) and Nesterov's accelerated method (convex) [Beck and Teboulle, 2009, Nesterov, 2004] on the least squares problem under the isotropic features model. We further investigate the halting time in logistic regression as well as least squares with mini-batch stochastic gradient descent (SGD). See Appendix D for details.

Setup. The vectors \mathbf{x}_0 and $\tilde{\mathbf{x}}$ are sampled i.i.d. from the Gaussian $N(0, \frac{1}{d} \mathbf{I})$ whereas the entries of \mathbf{A} are sampled either from a standardized Gaussian, a Bernoulli distribution, or a Student's t -distribution with 5 degrees of freedom, normalized so that they all have the same mean and variance. We train the following models:

- **Least squares.** The least squares problem minimizes the objective function $f(\mathbf{x}) = \frac{1}{2n} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2$. The targets, $\mathbf{b} = \mathbf{A}\tilde{\mathbf{x}} + \boldsymbol{\eta}$, are generated by adding a noise vector $\boldsymbol{\eta}$ to our signal, $\mathbf{A}\tilde{\mathbf{x}}$. The entries of $\boldsymbol{\eta}$ are sampled from a normal, $N(0, \tilde{R}^2)$, for different values of \tilde{R}^2 .

- **Logistic regression.** For the logistic regression model we generate targets in the domain $(0, 1)$ using $\mathbf{b} = \sigma(\mathbf{A}\tilde{\mathbf{x}} + \boldsymbol{\eta})$ where σ is the logistic function. The output of our model is $\mathbf{y} = \sigma(\mathbf{A}\mathbf{x})$, and the objective function is the standard cross-entropy loss:

$$f(\mathbf{x}) = -\frac{1}{n} \sum_{i=1}^n b_i \cdot \log(y_i) + (1 - b_i) \cdot \log(1 - y_i).$$




Figure 8: **Effect of the ratio $r = d/n$ on the halting time for various levels of noise \tilde{R}^2 .** The left plot shows the average halting time of gradient descent as a function of the ratio parameter r . As predicted by the theory, the halting time increases as \mathbf{A} approaches a square matrix ($r \rightarrow 1$), and the difference between the linear rates ($r \neq 1$) and the sublinear rates ($r = 1$) grows as the noise level increases. A total of 104,960 models were trained, keeping fixed the number of entries in the matrix $dn = 2^{24}$. In the right plot we show the same curve but for SGD instead, with a batch-size of $\frac{n}{8}$. We plot the curves for all values $r \neq 1$ with the value for $r = 1$ as a single point due its large value.

Parameter settings. In all simulations, the halting criterion is the number of steps until the gradient falls below ε , $\|\nabla f(\mathbf{x})\|^2 < \varepsilon$, where ε is chosen to be 10^{-6} for GD and Nesterov and ε is 10^{-4} for SGD. The step size for GD and Nesterov’s accelerated method is fixed to be $1/L$ where L is the Lipschitz constant of the gradient. For least squares, $L = \lambda_{\mathbf{H}}^+$. We approximate $\lambda_{\mathbf{H}}^+$ by performing 64 steps of the power iteration method on the matrix \mathbf{H} , initialized with a constant vector of norm 1. For logistic regression, we set the step size to be $4/\lambda_{\mathbf{H}}^+$.

In SGD, we sample rows from the matrix \mathbf{A} . The mini-batch size parameter is a fixed fraction of the data set size $\frac{n}{16}$, so that the comparison of halting times across model sizes is consistent. When the models are over-parametrized ($n < d$), a strong growth condition [Schmidt and Le Roux, 2013] holds. This means a scaling of the GD step size can be used to ensure convergence. In the under-parametrized setting, SGD does not converge to the optimum. In this case we chose a step size such that the expected squared gradient norm at the stationary point equals the halting criterion. See Appendix D.1 for derivations.

Results and conclusions. Figure 1 confirms our theoretical results: variability in the halting time decreases and the halting time converges to a deterministic quantity independent of the distribution of the data. Experimentally, the standard deviation decreased at a rate of $d^{-1/2}$, consistent with results in random matrix theory. For medium sized problems ($d = 2^5$), the heavy-tailed Student’s t distribution occasionally produces ill-conditioned matrices resulting in large halting times. These ill-conditioned matrices disappear as the model size grows in large part because the maximum eigenvalue becomes stable.

More interestingly, our results extend to non-quadratic functions, such as logistic regression, as well as SGD (see Figure 7). Surprisingly, we see different behaviors between logistic and least square

models for smaller matrices when using SGD. Moreover, we note that the large halting times seen in the Student's t distribution for GD on medium sized problems disappear when we instead run SGD.

Secondly, Figure 8 evaluates the halting times dependency on the ratio r . As predicted by the theory, the halting time takes its maximum value (*i.e.*, algorithm is slowest) precisely when $r = 1$. For SGD different step sizes are used for the over-parametrized and under-parametrized regime resulting in an asymmetric curve and a clear discontinuity at $r = 1$. We leave the study of these phenomena as future work.

Acknowledgements

The authors would like to thank our colleagues Nicolas Le Roux, Ross Goroshin, Zaid Harchaoui, Damien Scieur, and Dmitriy Drusvyatskiy for their feedback on this manuscript, and Henrik Ueber-schaer for providing useful random matrix theory references.

A Derivation of polynomials

In this section, we construct the residual polynomials for various popular first-order methods, including Nesterov's accelerated gradient and Polyak momentum.

A.1 Nesterov's accelerated methods

Nesterov accelerated methods generate iterates using the relation

$$\begin{aligned} \mathbf{x}_{k+1} &= \mathbf{y}_k - \alpha \nabla f(\mathbf{y}_k), \quad \text{where} \quad \alpha = \frac{1}{\lambda_{\mathbf{H}}^+} \\ \mathbf{y}_{k+1} &= \mathbf{x}_{k+1} + \beta_k (\mathbf{x}_{k+1} - \mathbf{x}_k), \quad \text{where} \quad \beta_k = \begin{cases} \frac{\sqrt{\lambda_{\mathbf{H}}^+} - \sqrt{\lambda_{\mathbf{H}}^-}}{\sqrt{\lambda_{\mathbf{H}}^+} + \sqrt{\lambda_{\mathbf{H}}^-}}, & \text{if } \lambda_{\mathbf{H}}^- \neq 0 \\ \frac{k}{k+3}, & \text{if } \lambda_{\mathbf{H}}^- = 0. \end{cases} \end{aligned}$$

By developing the recurrence of the iterates on the least squares problem (8), we get the following three-term recurrence

$$\mathbf{x}_{k+1} - \tilde{\mathbf{x}} = (1 + \beta_{k-1})(\mathbf{I} - \alpha \mathbf{H})(\mathbf{x}_k - \tilde{\mathbf{x}}) - \beta_{k-1}(\mathbf{I} - \alpha \mathbf{H})(\mathbf{x}_{k-1} - \tilde{\mathbf{x}}) + \alpha \cdot \frac{\mathbf{A}^T \boldsymbol{\eta}}{n},$$

with the initial vector $\mathbf{x}_0 \in \mathbb{R}^d$ and $\mathbf{x}_1 = \mathbf{x}_0 - \alpha \nabla f(\mathbf{x}_0)$. Using these standard initial conditions, we deduce from Proposition 3.1 the following

$$\begin{aligned} P_{k+1}(\mathbf{H}; \lambda_{\mathbf{H}}^{\pm})(\mathbf{x}_0 - \tilde{\mathbf{x}}) &+ Q_{k+1}(\mathbf{H}; \lambda_{\mathbf{H}}^{\pm}) \frac{\mathbf{A}^T \boldsymbol{\eta}}{n} \\ &= [(1 + \beta_{k-1})(\mathbf{I} - \alpha \mathbf{H})P_k(\mathbf{H}; \lambda_{\mathbf{H}}^{\pm}) - \beta_{k-1}(\mathbf{I} - \alpha \mathbf{H})P_{k-1}(\mathbf{H}; \lambda_{\mathbf{H}}^{\pm})](\mathbf{x}_0 - \tilde{\mathbf{x}}) \\ &+ [(1 + \beta_{k-1})(\mathbf{I} - \alpha \mathbf{H})Q_k(\mathbf{H}; \lambda_{\mathbf{H}}^{\pm}) - \beta_{k-1}(\mathbf{I} - \alpha \mathbf{H})Q_{k-1}(\mathbf{H}; \lambda_{\mathbf{H}}^{\pm}) + \alpha \mathbf{I}] \frac{\mathbf{A}^T \boldsymbol{\eta}}{n}. \end{aligned}$$

It immediately follows that the residual polynomials satisfy the same three-term recurrence, namely,

$$\begin{aligned} P_{k+1}(\lambda; \lambda_{\mathbf{H}}^{\pm}) &= (1 + \beta_{k-1})(1 - \alpha \lambda)P_k(\lambda; \lambda_{\mathbf{H}}^{\pm}) - \beta_{k-1}(1 - \alpha \lambda)P_{k-1}(\lambda; \lambda_{\mathbf{H}}^{\pm}) \\ &\quad \text{with } P_0(\lambda; \lambda_{\mathbf{H}}^{\pm}) = 1, \quad P_1(\lambda; \lambda_{\mathbf{H}}^{\pm}) = 1 - \alpha \lambda \\ Q_{k+1}(\lambda; \lambda_{\mathbf{H}}^{\pm}) &= (1 + \beta_{k-1})(1 - \alpha \lambda)Q_k(\lambda; \lambda_{\mathbf{H}}^{\pm}) - \beta_{k-1}(1 - \alpha \lambda)Q_{k-1}(\lambda; \lambda_{\mathbf{H}}^{\pm}) + \alpha \\ &\quad \text{with } Q_0(\lambda; \lambda_{\mathbf{H}}^{\pm}) = 0, \quad Q_1(\lambda; \lambda_{\mathbf{H}}^{\pm}) = \alpha. \end{aligned} \tag{85}$$

By Proposition 3.1, we only need to derive an explicit expression for the P_k polynomials.

A.1.1 Strongly-convex setting

The polynomial recurrence relationship for Nesterov's accelerated method in the strongly-convex setting is given by

$$P_{k+1}(\lambda; \lambda_{\mathbf{H}}^{\pm}) = (1 + \beta)(1 - \alpha\lambda)P_k(\lambda; \lambda_{\mathbf{H}}^{\pm}) - \beta(1 - \alpha\lambda)P_{k-1}(\lambda; \lambda_{\mathbf{H}}^{\pm})$$

where $P_0(\lambda; \lambda_{\mathbf{H}}^{\pm}) = 1$, $P_1(\lambda; \lambda_{\mathbf{H}}^{\pm}) = 1 - \alpha\lambda$, $\alpha = \frac{1}{\lambda_{\mathbf{H}}^{\pm}}$ and $\beta = \frac{\sqrt{\lambda_{\mathbf{H}}^+} - \sqrt{\lambda_{\mathbf{H}}^-}}{\sqrt{\lambda_{\mathbf{H}}^+} + \sqrt{\lambda_{\mathbf{H}}^-}}$. (86)

We generate an explicit representation for the polynomial by constructing the generating function for the polynomials P_k , namely

$$\begin{aligned} \mathfrak{G}(\lambda, t) &\stackrel{\text{def}}{=} \sum_{k=0}^{\infty} t^k P_k(\lambda; \lambda_{\mathbf{H}}^{\pm}) \\ \text{(recurrence in (86))} &= 1 + \frac{1}{t(1 + \beta)(1 - \alpha\lambda)} \sum_{k=2}^{\infty} t^k P_k(\lambda; \lambda_{\mathbf{H}}^{\pm}) + \frac{t\beta}{1 + \beta} \sum_{k=0}^{\infty} t^k P_k(\lambda; \lambda_{\mathbf{H}}^{\pm}) \\ \text{(initial conditions)} &= 1 + \frac{(\mathfrak{G}(\lambda, t) - (1 + t(1 - \alpha\lambda)))}{t(1 + \beta)(1 - \alpha\lambda)} + \frac{t\beta}{1 + \beta} \mathfrak{G}(\lambda, t). \end{aligned}$$

We solve this expression for \mathfrak{G} , which gives

$$\mathfrak{G}(\lambda, t) = \frac{1 - t\beta(1 - \alpha\lambda)}{1 + \beta(1 - \alpha\lambda)t^2 - t(1 + \beta)(1 - \alpha\lambda)}. \quad (87)$$

Ultimately, we want to relate the generating function for the polynomials P_k to a generating function for known polynomials. Notably, in this case, the Chebyshev polynomials of the 1st and 2nd kind – denoted $(T_k(x))$ and $(U_k(x))$ respectively – resemble the generating function for the residual polynomials of Nesterov accelerated method. The generating function for Chebyshev polynomials is given as

$$\sum_{k=0}^{\infty} (T_k(x) + \delta U_k(x))t^k = \frac{1 - tx + \delta}{1 - 2tx + t^2}. \quad (88)$$

To give the explicit relationship between (87) and (88), we make the substitution $t \mapsto \frac{t}{(\beta(1 - \alpha\lambda))^{1/2}}$. A simple calculation yields the following

$$\begin{aligned} \sum_{k=1}^{\infty} \frac{t^k P_k(\lambda; \lambda_{\mathbf{H}}^{\pm})}{(\beta(1 - \alpha\lambda))^{k/2}} &= \frac{1 - \frac{\beta\sqrt{1 - \alpha\lambda}}{\sqrt{\beta}}t}{1 - \frac{(1 + \beta)\sqrt{1 - \alpha\lambda}}{2\sqrt{\beta}} \cdot 2t + t^2} = \frac{1 - \frac{2\beta}{1 + \beta}tx}{1 - 2tx + t^2} = \frac{\frac{2\beta}{1 + \beta}(1 - tx) + \left(1 - \frac{2\beta}{1 + \beta}\right)}{1 - 2tx + t^2} \\ \text{where } x &= \frac{(1 + \beta)\sqrt{1 - \alpha\lambda}}{2\sqrt{\beta}}. \end{aligned} \quad (89)$$

We can compare (88) with (89) to derive an expression for the polynomials P_k

$$\begin{aligned} P_k(\lambda; \lambda_{\mathbf{H}}^{\pm}) &= (\beta(1 - \alpha\lambda))^{k/2} \left[\frac{2\beta}{1 + \beta} T_k \left(\frac{(1 + \beta)\sqrt{1 - \alpha\lambda}}{2\sqrt{\beta}} \right) \right. \\ &\quad \left. + \left(1 - \frac{2\beta}{1 + \beta} \right) U_k \left(\frac{(1 + \beta)\sqrt{1 - \alpha\lambda}}{2\sqrt{\beta}} \right) \right], \end{aligned} \quad (90)$$

where T_k is the Chebyshev polynomial of the first kind and U_k is the Chebyshev polynomial of the second kind.

A.1.2 Convex setting: Legendre polynomials and Bessel asymptotic

When the objective function is convex (*i.e.* $\lambda_{\mathbf{H}}^- = 0$), the recurrence for the residual polynomial associated with Nesterov's accelerated method residual reduces to

$$P_{k+1}(\lambda; \lambda_{\mathbf{H}}^{\pm}) = (1 + \beta_{k-1})(1 - \alpha\lambda)P_k(\lambda; \lambda_{\mathbf{H}}^{\pm}) - \beta_{k-1}(1 - \alpha\lambda)P_{k-1}(\lambda; \lambda_{\mathbf{H}}^{\pm})$$

$$\text{with } P_0(\lambda; \lambda_{\mathbf{H}}^{\pm}) = 1, \quad P_1(\lambda; \lambda_{\mathbf{H}}^{\pm}) = 1 - \alpha\lambda, \quad \alpha = \frac{1}{\lambda_{\mathbf{H}}^+}, \quad \text{and } \beta_k = \frac{k}{k+3}. \quad (91)$$

We now seek to solve this recurrence.

Nesterov's polynomials as Legendre polynomials. First we observe that these polynomials are also polynomials in $u = \alpha\lambda$, so we can define new polynomials $\tilde{P}_k(u)$ such that $\tilde{P}_k(\alpha\lambda) = P_k(\lambda; \lambda_{\mathbf{H}}^{\pm})$. Let us define new polynomials $\tilde{R}_k(u) = \tilde{P}_k(u)(1-u)^{-k/2}$. Then the recurrence in (91) can be reformulated as

$$\tilde{R}_{k+1}(u) = (1 + \beta_{k-1})(1-u)^{1/2}\tilde{R}_k(u) - \beta_{k-1}\tilde{R}_{k-1}(u), \quad \tilde{R}_0(u) = 1, \quad \tilde{R}_1(u) = (1-u)^{1/2}. \quad (92)$$

A simple computation shows that the polynomials $\{\tilde{R}_k\}$ are polynomials in $v = (1-u)^{1/2}$. Because of this observation, we define new polynomials $R_k(v)$ where $R_k((1-u)^{1/2}) = \tilde{R}_k(u)$. Now we will find a formula for the polynomials R_k by constructing its generating function,

$$\mathfrak{G}(v, t) \stackrel{\text{def}}{=} \sum_{k=0}^{\infty} R_k(v)t^k$$

The recurrence in (92) together with the definition of R_k yields the following differential identity

$$\begin{aligned} 2\partial_t(vt^{1/2}\mathfrak{G}(v, t)) &= \sum_{k=0}^{\infty} (2k+1)R_k(v)vt^{k-1/2} = vt^{-1/2} + \sum_{k=1}^{\infty} (2k+1)R_k(v)vt^{k-1/2} \\ &= vt^{-1/2} + \sum_{k=1}^{\infty} (k+2) \cdot \frac{2k+1}{k+2} \cdot v \cdot R_k(v)t^{k-1/2} \\ \text{(recurrence in (92))} \quad &= vt^{-1/2} + \sum_{k=1}^{\infty} (k-1)R_{k-1}(v)t^{k-1/2} + \sum_{k=1}^{\infty} (k+2)R_{k+1}(v)t^{k-1/2} \\ &= vt^{-1/2} + t^{3/2} \sum_{k=0}^{\infty} kR_k(v)t^{k-1} + \sum_{k=2}^{\infty} (k+1)t^{k-3/2}R_k(v) \\ (\partial_t(t\mathfrak{G})) &= \sum_{k=0}^{\infty} (k+1)t^k R_k(v) = vt^{-1/2} + t^{3/2}\partial_t(\mathfrak{G}) + t^{-3/2}\partial_t(t(\mathfrak{G} - (1+vt))). \end{aligned}$$

One can see this is a first-order linear ODE with initial conditions given by

$$\partial_t(\mathfrak{G}) + \frac{1-tv}{t^3-2vt^2+t}\mathfrak{G} = \frac{1+tv}{t^3-2vt^2+t}, \quad \text{with } \mathfrak{G}(v, 0) = 1, \quad \partial_t\mathfrak{G}(v, 0) = v.$$

Using an integrating factor of $\mu(t) = \frac{t}{\sqrt{t^2-2tv+1}}$, the solution to this initial value problem is

$$\mathfrak{G}(v, t) = \frac{2v\sqrt{t^2-2vt+1} + tv^2 + t - 2v}{t(1-v^2)}.$$

At first glance, this does not seem related to any known generating function for a polynomial; however if we differentiate this function we get that

$$\sum_{k=1}^{\infty} kR_k(v)t^k = t\partial_t(\mathfrak{G}) = \frac{2v(vt-1 + \sqrt{t^2-2tv+1})}{t(1-v^2)\sqrt{t^2-2tv+1}} = \frac{2v(vt-1)}{t(1-v^2)\sqrt{t^2-2tv+1}} + \frac{2v}{t(1-v^2)},$$

and it is known that the generating function for the Legendre Polynomials $\{L_k\}$ is exactly

$$\sum_{k=0}^{\infty} L_k(v)t^k = \frac{1}{\sqrt{t^2 - 2vt + 1}}.$$

Hence it follows that

$$\begin{aligned} \frac{2v(vt - 1)}{t(1 - v^2)\sqrt{t^2 - 2vt + 1}} &= \left(\frac{2v^2}{1 - v^2} - \frac{2v}{t(1 - v^2)} \right) \frac{1}{\sqrt{t^2 - 2vt + 1}} \\ &= \frac{2v^2}{1 - v^2} \sum_{k=0}^{\infty} L_k(v)t^k - \frac{2v}{1 - v^2} \sum_{k=0}^{\infty} L_k(v)t^{k-1}, \end{aligned}$$

and so for $k \geq 1$, by comparing coefficients we deduce that

$$kR_k(v) = \frac{2v^2}{1 - v^2} L_k(v) - \frac{2v}{1 - v^2} L_{k+1}(v).$$

Replacing all of our substitutions back in, we get the following representation for the Nesterov polynomials for $k \geq 1$

$$P_k(\lambda; \lambda_{\mathbf{H}}^{\pm}) = \frac{2(1 - \alpha\lambda)^{(k+1)/2}}{k\alpha\lambda} \left(\sqrt{1 - \alpha\lambda} \cdot L_k(\sqrt{1 - \alpha\lambda}) - L_{k+1}(\sqrt{1 - \alpha\lambda}) \right), \quad (93)$$

where $\{L_k\}$ are the Legendre polynomials.

Bessel asymptotic for Nesterov's residual polynomials. In this section, we derive an asymptotic for the residual polynomials of Nesterov's accelerated method in the convex setting. We will show that the polynomials P_k in (93) satisfy in a sufficiently strong sense

$$P_k(\lambda; \lambda_{\mathbf{H}}^{\pm}) \sim \frac{2J_1(k\sqrt{\alpha\lambda})}{k\sqrt{\alpha\lambda}} e^{-\alpha\lambda k/2}, \quad (94)$$

where J_1 is the Bessel function of the first kind. We will show that the Bessel asymptotic (94) follows directly from the Legendre polynomials in (93). To see this, recall the integral representation of a Legendre polynomial is given below by

$$L_k(\sqrt{1 - u}) = \frac{1}{\pi} \int_0^{\pi} (\sqrt{1 - u} + i\sqrt{u} \cos(\phi))^k d\phi,$$

and so we have

$$\begin{aligned} \frac{1}{\sqrt{u}} (\sqrt{1 - u} L_k(\sqrt{1 - u}) - L_{k+1}(\sqrt{1 - u})) &= \frac{-i}{\pi} \int_0^{\pi} (\sqrt{1 - u} + i\sqrt{u} \cos(\phi))^k \cos(\phi) d\phi \\ &= \frac{1}{\pi} \int_0^{\pi} \text{Im}\{(\sqrt{1 - u} + i\sqrt{u} \cos(\phi))^k\} \cos(\phi) d\phi \\ \text{(symmetry about } \frac{\pi}{2}) &= \frac{2}{\pi} \int_0^{\pi/2} \text{Im}\{(\sqrt{1 - u} + i\sqrt{u} \cos(\phi))^k\} \cos(\phi) d\phi. \end{aligned}$$

Now define the polynomial $\tilde{P}_k(u) = P_k(\lambda^+ u; \lambda_{\mathbf{H}}^{\pm})$ where the polynomials P_k satisfies Nesterov's recurrence (85). Using the derivation of Nesterov's polynomial from the previous section (93), we obtain the following expression

$$\begin{aligned} \tilde{P}_k(u) &= \frac{2(1 - u)^{(k+1)/2}}{ku} (\sqrt{1 - u} \cdot L_k(\sqrt{1 - u}) - L_{k+1}(\sqrt{1 - u})) \\ &= \frac{4(1 - u)^{(k+1)/2}}{k\pi\sqrt{u}} \int_0^{\pi/2} \text{Im}\{(\sqrt{1 - u} + i\sqrt{u} \cos(\phi))^k\} \cos(\phi) d\phi. \end{aligned} \quad (95)$$

We can get an explicit expression for the imaginary part of the k -th power in (95) by expressing $\sqrt{1-u} + i\sqrt{u}\cos(\phi)$ in terms of its polar form. In particular, we have that

$$\theta(u, \phi) \stackrel{\text{def}}{=} \tan^{-1} \left(\sqrt{\frac{u}{1-u}} \cos(\phi) \right) \quad \text{and} \quad R(u, \phi) \stackrel{\text{def}}{=} \sqrt{1-u+u\cos^2(\phi)} = \sqrt{1-u\sin^2(\phi)}.$$

Hence, we have the following

$$\begin{aligned} \tilde{P}_k(u) &= \frac{4(1-u)^{(k+1)/2}}{k\pi\sqrt{u}} \int_0^{\pi/2} \text{Im}\{(\sqrt{1-u} + i\sqrt{u}\cos(\phi))^k\} \cos(\phi) d\phi \\ &= \frac{4(1-u)^{(k+1)/2}}{k\pi\sqrt{u}} \int_0^{\pi/2} R(u, \phi)^k \sin(k\theta(u, \phi)) \cos(\phi) d\phi. \end{aligned} \quad (96)$$

Define the following integral

$$I_k(u) \stackrel{\text{def}}{=} \frac{2}{\pi} \int_0^{\pi/2} R(u, \phi)^k \sin(k\theta(u, \phi)) \cos(\phi) d\phi, \quad (97)$$

and note the similarity of this integral with the Bessel function, namely

$$J_1(k\sqrt{u}) = \frac{2}{\pi} \int_0^{\pi/2} \sin(k\sqrt{u}\cos(\phi)) \cos(\phi) d\phi.$$

Using this definition, the polynomial can be written as $\tilde{P}_k(u) = \frac{2(1-u)^{(k+1)/2}}{k\sqrt{u}} I_k(u)$. Since I_k is always bounded, then for $u \geq \log^2(k)/k$, the magnitude of $|\tilde{P}_k(u)|$ is smaller than any power of k . This follows by using the bound that $(1-x)^k \leq \exp(-kx)$ and noting that $\exp(-\log^2(k))$ decays faster than any polynomial in k . So the interesting asymptotic is for $u \leq \log^2(k)/k$, and for this we show the following.

Lemma A.1. *There is an absolute constant C so that for all $k \geq 1$ and $0 \leq u \leq \log^2(k)/k$*

$$|I_k(u) - J_1(k\sqrt{u})| \leq \begin{cases} Ck^{1/3}\sqrt{u}, & \text{if } u \leq k^{-4/3} \\ Ck^{-1/3}, & \text{if } u > k^{-4/3}. \end{cases}$$

Corollary A.1 (Nesterov's polynomial asymptotic). *There is an absolute constant C so that for all $k \geq 1$ and all $0 \leq u \leq \frac{\log^2(k)}{k}$, the following holds*

$$\left| \tilde{P}_k(u) - \frac{2e^{-uk/2}}{k\sqrt{u}} J_1(k\sqrt{u}) \right| \leq \begin{cases} Ce^{-uk/2} k^{-2/3}, & \text{if } u \leq k^{-4/3} \\ Ce^{-uk/2} u^{-1/2} k^{-4/3}, & \text{if } u > k^{-4/3}. \end{cases} \quad (98)$$

In particular, the following result holds for all $0 \leq u \leq \frac{\log^2(k)}{k}$

$$\left| \tilde{P}_k^2(u) - \frac{4e^{-uk} J_1^2(k\sqrt{u})}{k^2 u} \right| \leq \begin{cases} C(k^{-4/3} + k^{-13/6} u^{-3/4}), & \text{if } u \leq k^{-4/3} \\ Ce^{-uk}(u^{-1} k^{-8/3} + u^{-5/4} k^{-17/6}), & \text{if } u \geq k^{-4/3}. \end{cases} \quad (99)$$

Proof of Corollary A.1. First, we have that $\tilde{P}_k(u) = \frac{2(1-u)^{(k+1)/2}}{k\sqrt{u}} I_k(u)$. A simple triangle inequality shows that

$$\left| \tilde{P}_k(u) - \frac{2e^{-uk/2}}{k\sqrt{u}} J_1(k\sqrt{u}) \right| \leq \left| \tilde{P}_k(u) - \frac{2e^{-uk/2}}{k\sqrt{u}} I_k(u) \right| + \left| \frac{2e^{-uk/2}}{k\sqrt{u}} I_k(u) - \frac{2e^{-uk/2}}{k\sqrt{u}} J_1(k\sqrt{u}) \right|.$$

The first difference is small because $|(1-u)^{(k+1)/2} - e^{-uk/2}| \leq Ce^{-uk/2}(u + ku^2)$ for some absolute constant C and I_k is bounded. The second difference follows directly from Lemma A.1. The second inequality (99) follows from $|a^2 - b^2| \leq |a - b|(|a - b| + 2|b|)$ and $J_1(x) \leq \frac{C}{\sqrt{x}}$. \square

Proof of Lemma A.1. First, we observe that z^k is k -Lipschitz on the interval $[0, 1]$. Since e^{-u} and $1 - u$ lie in the interval $[0, 1]$ for any $u \in [0, 1]$, the Lipschitz property of z^k and the 2nd-order Taylor approximation of e^{-uk} imply that there exists an $\xi \in [0, 1]$ such that

$$\begin{aligned} |e^{-ku \sin^2(\phi)/2} - R(u, \phi)^k| &\leq \frac{k}{2} |e^{-u \sin^2(\phi)} - (1 - u \sin^2(\phi))| \\ &= \frac{k}{2} |1 - u \sin^2(\phi) + \frac{(u \sin^2(\phi))^2}{2} - \frac{e^{-\xi}(u \sin^2(\phi))^3}{3!} - (1 - u \sin^2(\phi))| \quad (100) \\ &= \frac{k}{2} \left| \frac{(u \sin^2(\phi))^2}{2} - \frac{e^{-\xi}}{6} (u \sin^2(\phi))^3 \right| \leq \frac{ku^2}{4} \leq \frac{C \log^4(k)}{k}. \end{aligned}$$

Here we used that $u \leq \frac{\log^2(k)}{k}$. Similarly, we have that

$$|R(u, \phi)^k - 1| = |(1 - u \sin^2(\phi))^{k/2} - 1| \leq \frac{k}{2} |1 - u \sin^2(\phi) - 1| = \frac{k}{2} u \sin^2(\phi). \quad (101)$$

We also know that $\sin(kx)$ is k -Lipschitz. Therefore again by Taylor approximation on $\tan^{-1}(x)$, we deduce the following bound for some $\xi_{u, \phi} \in [0, \sqrt{\frac{u}{1-u}}]$

$$\begin{aligned} |\sin(k\theta(u, \phi)) - \sin\left(k\sqrt{\frac{u}{1-u}} \cos(\phi)\right)| &\leq k \left| \tan^{-1}\left(\sqrt{\frac{u}{1-u}} \cos(\phi)\right) - \sqrt{\frac{u}{1-u}} \cos(\phi) \right| \\ &\leq k \left| \frac{6\xi_{u, \phi}^2 - 2}{(\xi_{u, \phi}^2 + 1)^3} \right| \left(\sqrt{\frac{u}{1-u}}\right)^3. \quad (102) \end{aligned}$$

Moreover, suppose $v = \sqrt{u}$ and consider the function $\sqrt{\frac{u}{1-u}} = \frac{v}{\sqrt{1-v^2}}$. Using a Taylor approximation at $v = 0$ and the k -Lip. of $\sin(kx)$, we obtain that

$$|\sin\left(k\sqrt{\frac{u}{1-u}} \cos(\phi)\right) - \sin(k\sqrt{u} \cos(\phi))| \leq k \left| \sqrt{\frac{u}{1-u}} - \sqrt{u} \right| \leq ku^{3/2}. \quad (103)$$

Since $0 \leq u \leq \frac{\log^2(k)}{k}$, we have that $\sqrt{\frac{u}{1-u}}$ is bounded by some constant C independent of k and hence, the constant $\xi_{u, \phi}$ is bounded. This implies that $\frac{6\xi_{u, \phi}^2 - 2}{(\xi_{u, \phi}^2 + 1)^3}$ is bounded by some absolute constant \tilde{C} . By putting together (102) and (103), we deduce that

$$|\sin(k\theta(u, \phi)) - \sin(k\sqrt{u} \cos(\phi))| \leq k \left| \frac{6\xi_{u, \phi}^2 - 2}{(\xi_{u, \phi}^2 + 1)^3} \right| \left(\sqrt{\frac{u}{1-u}}\right)^3 + ku^{3/2} \leq Cku^{3/2} \quad (104)$$

where C is some absolute constant. Here we used that $u \leq \frac{\log^2(k)}{k}$ and $\log^2(k) \leq Ck^{1/3}$. We now consider three cases. Suppose $u \leq k^{-4/3}$. Using $u \leq k^{-4/3}$ and (101) we deduce that

$$|R(u, \phi)^k - 1| \leq \frac{k}{2} u \leq \frac{k^{1/3} \sqrt{u}}{2}. \quad (105)$$

By putting together (104) and (105), we get the following bound for all $u \leq k^{-4/3}$.

$$\begin{aligned} |I_k(u) - J_1(k\sqrt{u})| &\leq \frac{2}{\pi} \left| \int_0^{\pi/2} R(u, \phi)^k \sin(k\theta(u, \phi)) \cos(\phi) d\phi - \int_0^{\pi/2} \sin(k\theta(u, \phi)) \cos(\phi) d\phi \right| \\ &\quad + \frac{2}{\pi} \left| \int_0^{\pi/2} \sin(k\theta(u, \phi)) \cos(\phi) d\phi - \int_0^{\pi/2} \sin(k\sqrt{u} \cos(\phi)) \cos(\phi) d\phi \right| \\ &\leq \frac{k^{1/3} \sqrt{u}}{2} + Ck^{1/3} \sqrt{u}. \end{aligned}$$

The result immediately follows. On the other hand for $k^{-4/3} \leq u \leq \log^2(k)/k$, we cut the range of ϕ . Let ϕ_0 be such that $\phi_0 = k^{-2/3}u^{-1/2}$. We know from $u \geq k^{-4/3}$ that

$$\phi_0 = k^{-2/3}u^{-1/2} \leq k^{-2/3}k^{2/3} = 1. \quad (106)$$

Now for $\phi \leq \phi_0$, we have in this range that

$$\begin{aligned} & \left| \int_0^{\phi_0} R(u, \phi)^k \sin(k\theta(u, \phi)) \cos(\phi) d\phi - \int_0^{\phi_0} \sin(k\sqrt{u} \cos(\phi)) \cos(\phi) d\phi \right| \\ & \leq \int_0^{\phi_0} |R(u, \phi)^k - 1| d\phi + \int_0^{\phi_0} |\sin(k\theta(u, \phi)) - \sin(\sqrt{u}k \cos(\phi))| d\phi \\ & \leq \frac{k}{2}u\phi_0^2 + Cku^{3/2} \leq \frac{1}{2k^{1/3}} + \frac{C \log^3(k)}{k^{1/2}}. \end{aligned} \quad (107)$$

In the last inequality we used that $\phi_0 = k^{-2/3}u^{-1/2}$ and $u \leq \log^2(k)/k$. Since $\frac{\log^3(k)}{k^{1/2}} \leq Ck^{-1/3}$, the result immediately follows.

For larger ϕ , we use integration by parts with $F(\phi) = \cos(k\sqrt{u} \cos(\phi))$ and $G(\phi) = e^{-k \sin^2(\phi)u/2} \cot(\phi)$ to express

$$\begin{aligned} I_1 & \stackrel{\text{def}}{=} \int_{\phi_0}^{\pi/2} \sin(k\sqrt{\frac{u}{1-u}} \cos(\phi)) e^{-k \sin^2 \phi^2 u/2} \cos(\phi) d\phi = \frac{1}{k\sqrt{\frac{u}{1-u}}} \int_{\phi_0}^{\pi/2} F'(\phi)G(\phi) d\phi \\ & = \frac{\sqrt{1-u}}{k\sqrt{u}} F(\phi)G(\phi) \Big|_{\phi_0}^{\pi/2} - \frac{\sqrt{1-u}}{k\sqrt{u}} \int_{\phi_0}^{\pi/2} G'(\phi)F(\phi) d\phi \\ (\text{def. of } G'(\phi)) & = \frac{\sqrt{1-u}}{k\sqrt{u}} \left[F(\phi)G(\phi) \Big|_{\phi_0}^{\pi/2} - \int_{\phi_0}^{\pi/2} F(\phi) e^{-k \sin^2(\phi)u/2} (\cot(\phi)(-k \sin(2\phi)u) - \csc^2(\phi)) d\phi \right] \\ & = \frac{\sqrt{1-u}}{k\sqrt{u}} F(\phi)G(\phi) \Big|_{\phi_0}^{\pi/2} + \sqrt{u(1-u)} \int_{\phi_0}^{\pi/2} F(\phi) e^{-k \sin^2(\phi)u/2} \cot(\phi) \sin(2\phi) d\phi \\ & \quad + \frac{\sqrt{1-u}}{k\sqrt{u}} \int_{\phi_0}^{\pi/2} F(\phi) e^{-k \sin^2(\phi)u/2} \csc^2(\phi) d\phi. \end{aligned}$$

Since $u \in [0, 1]$, we get the following bound

$$\begin{aligned} |I_1| & \leq \underbrace{\frac{C}{k\sqrt{u}} |F(\phi_0)G(\phi_0)|}_{(a)} + \underbrace{\left| C\sqrt{u} \int_{\phi_0}^{\pi/2} F(\phi) e^{-k \sin^2 \phi u/2} \cot(\phi) \sin(2\phi) d\phi \right|}_{(b)} \\ & \quad + \underbrace{\left| \frac{C}{k\sqrt{u}} \int_{\phi_0}^{\pi/2} F(\phi) e^{-k \sin^2 \phi u/2} \csc^2(\phi) d\phi \right|}_{(c)} \end{aligned} \quad (108)$$

for some $C > 0$. We will bound each of the terms in (108) independently. For (a), Taylor's approximation yields that $|\cot(\phi_0) - \frac{1}{\phi_0}| \leq C$ which implies that $|\cot(\phi_0)| \leq \frac{C}{\phi_0}$ for some positive constants. Therefore, we deduce that the quantity (a) is bounded by $\frac{C}{k\sqrt{u}\phi_0}$. For (b) since $|F(\phi)| \leq 1$, $|\cot(\phi) \sin(2\phi)| = |2 \cos^2(\phi)| \leq 2$ and, of course, $|e^{-k \sin^2 \phi u/2}| \leq 1$, we have that the quantity (b) is bounded by $C\sqrt{u}$. As for the quantity (c), we use the following approximation $|\csc^2(\phi) - \frac{1}{\phi^2}| \leq C$ so that $|\csc^2(\phi)| \leq \frac{C}{\phi^2}$. Hence the integral (c) is bounded by $\frac{C}{k\sqrt{u}\phi_0}$. Therefore we conclude that

$$|I_1| \leq \frac{C}{k\sqrt{u}\phi_0} + C\sqrt{u} + \frac{C}{k\sqrt{u}\phi_0} \leq Ck^{-1/3}. \quad (109)$$

Here we used that $k\sqrt{u}\phi_0 = k^{1/3}$ and $\sqrt{u} \leq \log(k)/\sqrt{k} \leq Ck^{-1/3}$.

Now let's repeat this process repeating $G(\phi) = e^{-k\sin^2(\phi)u/2} \cot(\phi)$ with $G(\phi) = \cot(\phi)$. This time we have that

$$I_2 \stackrel{\text{def}}{=} \int_{\phi_0}^{\pi/2} \sin(k\sqrt{\frac{u}{1-u}} \cos(\phi)) \cos(\phi) = \frac{1}{k\sqrt{\frac{u}{1-u}}} \int_{\phi_0}^{\pi/2} F'(\phi)G(\phi) d\phi.$$

Using the same bounds as before, we deduce the following

$$|I_2| \leq \frac{C}{k\sqrt{u}\phi_0} + \frac{C}{k\sqrt{u}\phi_0} \leq Ck^{-1/3}. \quad (110)$$

For $u \geq k^{-4/3}$, we have the following result

$$\begin{aligned} |I_k(u) - J_1(k\sqrt{u})| &\leq \frac{2}{\pi} \left| \int_0^{\phi_0} R(u, \phi)^k \sin(k\theta(u, \phi)) \cos(\phi) d\phi - \int_0^{\phi_0} \sin(k\sqrt{u} \cos(\phi)) \cos(\phi) d\phi \right| \\ &\quad + \frac{2}{\pi} \left| \int_{\phi_0}^{\pi/2} R(u, \phi)^k \sin(k\theta(u, \phi)) \cos(\phi) d\phi - \int_{\phi_0}^{\pi/2} \sin(k\sqrt{u} \cos(\phi)) \cos(\phi) d\phi \right| \\ \text{(by (107))} \quad &\leq Ck^{-1/3} + \frac{2}{\pi} \int_{\phi_0}^{\pi/2} |(R(u, \phi)^k - e^{-k\sin^2 \phi u/2}) \sin(k\theta(u, \phi)) \cos(\phi)| d\phi + \frac{2}{\pi} |I_1| \\ &\quad + \frac{2}{\pi} |I_2| + \frac{2}{\pi} \int_{\phi_0}^{\pi/2} |[\sin(k\theta(u, \phi)) - \sin(k\sqrt{u} \cos(\phi))] \cos(\phi)| d\phi \\ \text{(by (109) and (110))} \quad &\leq Ck^{-1/3} + \frac{2}{\pi} \int_{\phi_0}^{\pi/2} |(R(u, \phi)^k - e^{-k\sin^2 \phi u/2}) \sin(k\theta(u, \phi)) \cos(\phi)| d\phi \\ &\quad + \frac{2}{\pi} \int_{\phi_0}^{\pi/2} |[\sin(k\theta(u, \phi)) - \sin(k\sqrt{u} \cos(\phi))] \cos(\phi)| d\phi \\ \text{(by (100) and (104))} \quad &\leq Ck^{-1/3}. \end{aligned}$$

This finishes the proof for the lemma. \square

A.2 Polyak Momentum (Heavy-ball) Method

The polynomials that generated Polyak's heavy ball method satisfies the following three-term recursion

$$\begin{aligned} P_{k+1}(\lambda; \lambda_{\mathbf{H}}^{\pm}) &= (1 - m + \alpha\lambda)P_k(\lambda; \lambda_{\mathbf{H}}^{\pm}) + mP_{k-1}(\lambda; \lambda_{\mathbf{H}}^{\pm}), \quad P_0 = 1, \text{ and } P_1(\lambda; \lambda_{\mathbf{H}}^{\pm}) = 1 - \beta\lambda \\ \text{where } m &= -\left(\frac{\sqrt{\lambda_{\mathbf{H}}^+} - \sqrt{\lambda_{\mathbf{H}}^-}}{\sqrt{\lambda_{\mathbf{H}}^+} + \sqrt{\lambda_{\mathbf{H}}^-}}\right)^2, \quad \alpha = \frac{-4}{(\sqrt{\lambda_{\mathbf{H}}^+} + \sqrt{\lambda_{\mathbf{H}}^-})^2}, \quad \text{and } \beta = \frac{2}{\lambda_{\mathbf{H}}^+ + \lambda_{\mathbf{H}}^-}. \end{aligned} \quad (111)$$

As in the previous examples, we will construct the generating function for the polynomials P_k using the recurrence in (111)

$$\begin{aligned} \mathfrak{G}(\lambda, t) &\stackrel{\text{def}}{=} \sum_{k=0}^{\infty} t^k P_k(\lambda; \lambda_{\mathbf{H}}^{\pm}) = 1 + \frac{1}{t(1 - m + \alpha\lambda)} \sum_{k=2}^{\infty} t^k P_k(\lambda; \lambda_{\mathbf{H}}^{\pm}) - \frac{mt}{1 - m + \alpha\lambda} \sum_{k=0}^{\infty} t^k P_k(\lambda; \lambda_{\mathbf{H}}^{\pm}) \\ &= 1 + \frac{1}{t(1 - m + \alpha\lambda)} [\mathfrak{G}(\lambda, t) - 1 - t(1 - \beta\lambda)] - \frac{mt}{1 - m + \alpha\lambda} \mathfrak{G}(\lambda, t). \end{aligned}$$

We solve for the generating function

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

This generating function for Polyak resembles the generating function for Chebyshev polynomials of the first and second kind (88). First, we set $t \mapsto \frac{t}{\sqrt{-m}}$ (note that $m < 0$ by definition in (111)). Under this transformation, we have the following

$$\begin{aligned} \sum_{k=0}^{\infty} \frac{t^k P_k(\lambda; \lambda_{\mathbf{H}}^{\pm})}{(-m)^{k/2}} &= \frac{1 - \frac{t}{\sqrt{-m}}(-m + (\alpha + \beta)\lambda)}{1 - 2t \left(\frac{1-m+\alpha\lambda}{2\sqrt{-m}} \right) + t^2} = \frac{1 - \sigma(\lambda)t \cdot \frac{-m+(\alpha+\beta)\lambda}{\sqrt{-m}\cdot\sigma(\lambda)}}{1 - 2\sigma(\lambda)t + t^2} \\ &= \frac{\frac{-m+(\alpha+\beta)\lambda}{\sqrt{-m}\cdot\sigma(\lambda)}(1 - \sigma(\lambda)t) + 1 - \frac{-m+(\alpha+\beta)\lambda}{\sqrt{-m}\cdot\sigma(\lambda)}}{1 - 2\sigma(\lambda)t + t^2}, \end{aligned} \quad (112)$$

where $\sigma(\lambda) = \frac{\lambda_{\mathbf{H}}^+ + \lambda_{\mathbf{H}}^- - 2\lambda}{\lambda_{\mathbf{H}}^+ - \lambda_{\mathbf{H}}^-}$. A simple computation shows that

$$\frac{-m + (\alpha + \beta)\lambda}{\sqrt{-m}\sigma(\lambda)} = \frac{(\sqrt{\lambda_{\mathbf{H}}^+} - \sqrt{\lambda_{\mathbf{H}}^-})^2}{\lambda_{\mathbf{H}}^+ + \lambda_{\mathbf{H}}^-} \quad \text{and} \quad 1 - \frac{-m + (\alpha + \beta)\lambda}{\sqrt{-m}\sigma(\lambda)} = \frac{2\sqrt{\lambda_{\mathbf{H}}^+ \lambda_{\mathbf{H}}^-}}{\lambda_{\mathbf{H}}^+ + \lambda_{\mathbf{H}}^-}.$$

By matching terms in the generating function for Chebyshev polynomials (88) and Polyak's generating function (112), we derive an expression for the polynomials P_k in Polyak's momentum

$$P_k(\lambda; \lambda_{\mathbf{H}}^{\pm}) = \left(\frac{\sqrt{\lambda_{\mathbf{H}}^+} - \sqrt{\lambda_{\mathbf{H}}^-}}{\sqrt{\lambda_{\mathbf{H}}^+} + \sqrt{\lambda_{\mathbf{H}}^-}} \right)^k \left[\frac{(\sqrt{\lambda_{\mathbf{H}}^+} - \sqrt{\lambda_{\mathbf{H}}^-})^2}{\lambda_{\mathbf{H}}^+ + \lambda_{\mathbf{H}}^-} \cdot T_k(\sigma(\lambda)) + \frac{2\sqrt{\lambda_{\mathbf{H}}^+ \lambda_{\mathbf{H}}^-}}{\lambda_{\mathbf{H}}^+ + \lambda_{\mathbf{H}}^-} \cdot U_k(\sigma(\lambda)) \right] \quad (113)$$

where $T_k(x)$ ($U_k(x)$) is the Chebyshev polynomial of the 1st (2nd) kind respectively

$$\text{and} \quad \sigma(\lambda) = \frac{\lambda_{\mathbf{H}}^+ + \lambda_{\mathbf{H}}^- - 2\lambda}{\lambda_{\mathbf{H}}^+ - \lambda_{\mathbf{H}}^-}.$$

B Average-case complexity

In this section, we compute the average-case complexity for various first-order methods. To do so, we integrate the residual polynomials found in Table 3 against the Marčenko-Pastur density.

Lemma B.1 (Average-case: Gradient descent). *Let $d\mu_{\text{MP}}$ be the Marčenko-Pastur law defined in (2) and P_k, Q_k be the residual polynomials for gradient descent.*

1. For $r = 1$ and $\ell = \{1, 2\}$, the following holds

$$\int \lambda^\ell P_k^2(\lambda; \lambda^\pm) d\mu_{\text{MP}} = \frac{(\lambda^+)^{\ell+1}}{2\pi\sigma^2} \cdot \frac{\Gamma(2k + \frac{3}{2})\Gamma(\ell + \frac{1}{2})}{\Gamma(2k + \ell + 2)} \sim \frac{(\lambda^+)^{\ell+1}}{2\pi\sigma^2} \cdot \frac{\Gamma(\ell + \frac{1}{2})}{(2k + 3/2)^{\ell+1/2}}.$$

2. For $r \neq 1$, the following holds

$$\begin{aligned} \int \lambda P_k^2(\lambda; \lambda^\pm) d\mu_{\text{MP}} &= \frac{(\lambda^+ - \lambda^-)^2}{2\pi\sigma^2 r} \left(1 - \frac{\lambda^-}{\lambda^+} \right)^{2k} \frac{\Gamma(2k + \frac{3}{2})\Gamma(\frac{3}{2})}{\Gamma(2k + 3)} \\ &\sim \frac{(\lambda^+ - \lambda^-)^2}{2\pi\sigma^2 r} \left(1 - \frac{\lambda^-}{\lambda^+} \right)^{2k} \cdot \frac{\Gamma(\frac{3}{2})}{(2k + \frac{3}{2})^{3/2}}. \end{aligned}$$

3. For $r \neq 1$, the following holds

$$\begin{aligned} &\int \lambda^2 P_k^2(\lambda; \lambda^\pm) d\mu_{\text{MP}} \\ &= \frac{(\lambda^+ - \lambda^-)^2}{2\pi\sigma^2 r} \left(1 - \frac{\lambda^-}{\lambda^+} \right)^{2k} \left(\frac{\lambda^- \cdot \Gamma(2k + \frac{3}{2})\Gamma(\frac{3}{2})}{\Gamma(2k + 3)} + \frac{(\lambda^+ - \lambda^-) \cdot \Gamma(2k + \frac{3}{2})\Gamma(\frac{5}{2})}{\Gamma(2k + 4)} \right) \\ &\sim \frac{(\lambda^+ - \lambda^-)^2}{2\pi\sigma^2 r} \left(1 - \frac{\lambda^-}{\lambda^+} \right)^{2k} \left(\frac{\lambda^- \cdot \Gamma(\frac{3}{2})}{(2k + \frac{3}{2})^{3/2}} + \frac{(\lambda^+ - \lambda^-)\Gamma(\frac{5}{2})}{(2k + \frac{3}{2})^{5/2}} \right). \end{aligned}$$

Proof. The proof relies on writing the integrals in terms of β -functions. Let $\ell = \{1, 2\}$. Using a change of variables $\lambda = \lambda^- + (\lambda^+ - \lambda^-)w$, we deduce the following expression

$$\begin{aligned} \int \lambda^\ell P_k^2(\lambda; \lambda^\pm) d\mu_{\text{MP}} &= \frac{1}{2\pi\sigma^2 r} \int_{\lambda^-}^{\lambda^+} \lambda^{\ell-1} \left(1 - \frac{\lambda}{\lambda^+}\right)^{2k} \sqrt{(\lambda - \lambda^-)(\lambda^+ - \lambda)} d\lambda \\ &= \frac{(\lambda^+ - \lambda^-)^2}{2\pi\sigma^2 r} \left(1 - \frac{\lambda^-}{\lambda^+}\right)^{2k} \int_0^1 (1-w)^{2k} (\lambda^- + (\lambda^+ - \lambda^-)w)^{\ell-1} \sqrt{w(1-w)} dw. \end{aligned} \quad (114)$$

We consider cases depending on whether $\lambda^- = 0$ or not (i.e. $r = 1$). First suppose $\lambda^- = 0$ so by equation (114) we have

$$\frac{1}{2\pi\sigma^2 r} \int_{\lambda^-}^{\lambda^+} \lambda^{\ell-1} \left(1 - \frac{\lambda}{\lambda^+}\right)^{2k} \sqrt{(\lambda - \lambda^-)(\lambda^+ - \lambda)} d\lambda = \frac{(\lambda^+)^{\ell+1}}{2\pi\sigma^2 r} \int_0^1 (1-w)^{2k+1/2} w^{\ell-1/2} dw.$$

The result follows after noting that the integral is a β -function with parameters $2k + 3/2$ and $\ell + 1/2$ as well as the asymptotics of β -functions, $\beta(x, y) = \Gamma(y)x^{-y}$ for x large and y fixed.

Next consider when $r \neq 1$ and $\ell = 1$. Using (114), we have that

$$\int \lambda P_k^2(\lambda; \lambda^\pm) d\mu_{\text{MP}} = \frac{(\lambda^+ - \lambda^-)^2}{2\pi\sigma^2 r} \left(1 - \frac{\lambda^-}{\lambda^+}\right)^{2k} \int_0^1 (1-w)^{2k+1/2} w^{1/2} dw.$$

The integral is a β -function with parameters $2k + 3/2$ and $3/2$. Applying the asymptotics of β -functions, finishes this case.

Lastly consider when $r \neq 1$ and $\ell = 2$. Similar to the previous case, using (114), the following holds

$$\begin{aligned} \int \lambda^2 P_k^2(\lambda; \lambda^\pm) d\mu_{\text{MP}} &= \frac{(\lambda^+ - \lambda^-)^2}{2\pi\sigma^2 r} \left(1 - \frac{\lambda^-}{\lambda^+}\right)^{2k} \left(\lambda^- \int_0^1 (1-w)^{2k+1/2} w^{1/2} dw \right. \\ &\quad \left. + (\lambda^+ - \lambda^-) \int_0^1 (1-w)^{2k+1/2} w^{3/2} dw \right). \end{aligned}$$

The first integral is a β -function with parameters $2k + 3/2$ and $3/2$ and the second term is a β -function with parameters $2k + 3/2$ and $5/2$. Again using the asymptotics for β -functions yields the result. \square

Lemma B.2 (Average-case: Nesterov accelerated method (strongly convex)). *Let $d\mu_{\text{MP}}$ be the Marčenko-Pastur law defined in (2) and P_k be the residual polynomial for Nesterov accelerated method on a strongly convex objective function (90). Then the following holds*

$$\begin{aligned} \int \lambda P_k^2(\lambda; \lambda^\pm) d\mu_{\text{MP}} &= \frac{(\lambda^+ - \lambda^-)^2}{2^5 4^k \sigma^2 r} \left(\beta \left(1 - \frac{\lambda^-}{\lambda^+}\right)\right)^k \left[\frac{4\beta^2}{(1+\beta)^2} \left(-k^2 + \frac{k}{2} + 1 - \binom{2k+2}{k} + \binom{2k+2}{k+1}\right) \right. \\ &\quad \left. + \frac{4\beta}{1+\beta} \left(1 - \frac{2\beta}{1+\beta}\right) \left(2k+1 - \binom{2k+2}{k} + \binom{2k+2}{k+1}\right) + 2 \left(1 - \frac{2\beta}{1+\beta}\right)^2 \left(\binom{2k+2}{k+1} - 1\right) \right] \\ &\sim \frac{(\lambda^+ - \lambda^-)^2}{4\sigma^2 r \sqrt{\pi}} \left(1 - \frac{2\beta}{1+\beta}\right)^2 \left(\beta \left(1 - \frac{\lambda^-}{\lambda^+}\right)\right)^k \frac{1}{k^{1/2}}, \end{aligned} \quad (115)$$

and the integral equals

$$\begin{aligned}
\int \lambda^2 P_k^2(\lambda; \lambda^\pm) d\mu_{\text{MP}} &= \lambda^- \int \lambda P_k^2(\lambda; \lambda^\pm) d\mu_{\text{MP}} + \frac{(\lambda^+ - \lambda^-)^3}{2^7 4^k \sigma^2 r} \left(\beta \left(1 - \frac{\lambda^-}{\lambda^+} \right) \right)^k \left[\right. \\
&\quad \frac{4\beta^2}{(1+\beta)^2} \left(\frac{1}{3}(2k^3 - 9k^2 + k + 6) - 4 \binom{2k+2}{k} + \binom{2k+2}{k-1} + 3 \binom{2k+2}{k+1} \right) \\
&\quad + \frac{4\beta}{1+\beta} \left(1 - \frac{2\beta}{1+\beta} \right) \left(-2k^2 + 3k + 2 - 4 \binom{2k+2}{k} + \binom{2k+2}{k-1} + 3 \binom{2k+2}{k+1} \right) \\
&\quad \left. + 4 \left(1 - \frac{2\beta}{1+\beta} \right)^2 \left(k - \binom{2k+2}{k} + \binom{2k+2}{k+1} \right) \right] \\
&\sim \frac{\lambda^- (\lambda^+ - \lambda^-)^2}{4\sigma^2 r \sqrt{\pi}} \left(1 - \frac{2\beta}{1+\beta} \right)^2 \left(\beta \left(1 - \frac{\lambda^-}{\lambda^+} \right) \right)^k \frac{1}{k^{1/2}},
\end{aligned} \tag{116}$$

where $\beta = \frac{\sqrt{\lambda^+} - \sqrt{\lambda^-}}{\sqrt{\lambda^+} + \sqrt{\lambda^-}}$ and $\alpha = \frac{1}{\lambda^+}$.

Proof. Throughout this proof, we define $P_k(\lambda)$ to be $P_k(\lambda; \lambda^\pm)$ in order to simplify the notation. In order to integrate the Chebyshev polynomials, we reduce our integral to trig functions via a series of change of variables. Under the change of variables that sends $\lambda = \lambda^- + (\lambda^+ - \lambda^-)w$, we have that for any $\ell \geq 1$

$$\int \lambda^\ell P_k^2(\lambda) d\mu_{\text{MP}} = \frac{(\lambda^+ - \lambda^-)^2}{2\pi\sigma^2 r} \int_0^1 P_k^2(\lambda^- + (\lambda^+ - \lambda^-)w) (\lambda^- + (\lambda^+ - \lambda^-)w)^{\ell-1} \sqrt{w(1-w)} dw. \tag{117}$$

We note that under this transformation $1 - \alpha\lambda = (1 - \frac{\lambda^-}{\lambda^+})(1 - w)$ and $\frac{1+\beta}{2\sqrt{\beta}}\sqrt{1 - \alpha\lambda} = (1 - w)^{1/2}$. Moreover by expanding out Nesterov's polynomial (90), we deduce the following

$$\begin{aligned}
P_k^2(\lambda) &= (\beta x)^k \left(\frac{4\beta^2}{(1+\beta)^2} T_k^2(y) + \frac{2\beta}{1+\beta} \left(1 - \frac{2\beta}{1+\beta} \right) T_k(y) U_k(y) + \left(1 - \frac{4\beta}{1+\beta} \right)^2 U_k^2(y) \right), \\
\text{where } y &= \frac{1+\beta}{2\sqrt{\beta}} \sqrt{x} \quad \text{and } x = 1 - \alpha\lambda.
\end{aligned} \tag{118}$$

First, we consider the setting where $\ell = 1$ in (117) and hence we deduce that

$$\begin{aligned}
\int \lambda P_k^2(\lambda) d\mu_{\text{MP}} &= \frac{(\lambda^+ - \lambda^-)^2}{2\pi\sigma^2 r} \left(\beta \left(1 - \frac{\lambda^-}{\lambda^+} \right) \right)^k \int_0^1 (1-w)^k \sqrt{w(1-w)} \left[\frac{4\beta^2}{(1+\beta)^2} T_k^2((1-w)^{1/2}) \right. \\
&\quad \left. + \frac{4\beta}{1+\beta} \left(1 - \frac{2\beta}{1+\beta} \right) T_k((1-w)^{1/2}) U_k((1-w)^{1/2}) + \left(1 - \frac{2\beta}{1+\beta} \right)^2 U_k^2((1-w)^{1/2}) \right] dw \\
(1-w = \cos^2(\theta)) &= \frac{2(\lambda^+ - \lambda^-)^2}{2\pi\sigma^2 r} \left(\beta \left(1 - \frac{\lambda^-}{\lambda^+} \right) \right)^k \int_0^{\pi/2} \cos^{2k+2}(\theta) \sin^2(\theta) \left[\frac{4\beta^2}{(1+\beta)^2} \cos^2(k\theta) \right. \\
&\quad \left. + \frac{4\beta}{1+\beta} \left(1 - \frac{2\beta}{1+\beta} \right) \cos(k\theta) \cdot \frac{\sin((k+1)\theta)}{\sin(\theta)} + \left(1 - \frac{2\beta}{1+\beta} \right)^2 \frac{\sin^2((k+1)\theta)}{\sin^2(\theta)} \right] d\theta \\
(\text{by symmetry}) &= \frac{2(\lambda^+ - \lambda^-)^2}{8\pi\sigma^2 r} \left(\beta \left(1 - \frac{\lambda^-}{\lambda^+} \right) \right)^k \int_0^{2\pi} \cos^{2k+2}(\theta) \sin^2(\theta) \left[\frac{4\beta^2}{(1+\beta)^2} \cos^2(k\theta) \right. \\
&\quad \left. + \frac{4\beta}{1+\beta} \left(1 - \frac{2\beta}{1+\beta} \right) \cos(k\theta) \cdot \frac{\sin((k+1)\theta)}{\sin(\theta)} + \left(1 - \frac{2\beta}{1+\beta} \right)^2 \frac{\sin^2((k+1)\theta)}{\sin^2(\theta)} \right] d\theta.
\end{aligned} \tag{119}$$

We will treat each term in the summand separately. Because the integral of $\int_0^{2\pi} e^{ik\theta} = 0$ for any $k \in \mathbb{N}$, we only need to keep track of the constant terms. From this observation, we get the following

$$\begin{aligned}
\cos^{2k+2}(\theta) \sin^2(\theta) \cos^2(k\theta) &= \frac{1}{2^4 4^k} \left(-k^2 + \frac{k}{2} + 1 - \binom{2k+2}{k} + \binom{2k+2}{k+1} \right) + \text{terms} \\
\cos^{2k+2}(\theta) \sin(\theta) \cos(k\theta) \sin((k+1)\theta) &= \frac{1}{2^4 4^k} \left(2k+1 - \binom{2k+2}{k} + \binom{2k+2}{k+1} \right) + \text{terms} \\
\cos^{2k+2}(\theta) \sin^2((k+1)\theta) &= \frac{1}{2^4 4^k} \left(-2 + 2 \binom{2k+2}{k+1} \right) + \text{terms}
\end{aligned} \tag{120}$$

We note that $\frac{1}{4^k} \left(\binom{2k+2}{k+1} - \binom{2k+2}{k} \right) \sim \frac{4}{\sqrt{\pi} k^{3/2}}$ and $\frac{1}{4^k} \binom{2k+2}{k+1} \sim \frac{4}{\sqrt{\pi} k^{1/2}}$.

Next, we consider the setting where $\ell = 2$ and we observe from (117) that

$$\int (\lambda^2 - \lambda^- \lambda) P_k^2(\lambda) d\mu_{\text{MP}} = \frac{(\lambda^+ - \lambda^-)^3}{2\pi\sigma^2 r} \int_0^1 P_k^2(\lambda^- + (\lambda^+ - \lambda^-)w) w \sqrt{w(1-w)} dw.$$

Since we know how to evaluate $\int \lambda P_k^2(\lambda) d\mu_{\text{MP}}$, we only need to analyze the RHS of this integral. A similar analysis as in (119) applies

$$\begin{aligned} \int (\lambda^2 - \lambda^- \lambda) P_k^2(\lambda) d\mu_{\text{MP}} &= \frac{2(\lambda^+ - \lambda^-)^3}{8\pi\sigma^2 r} \left(\beta \left(1 - \frac{\lambda^-}{\lambda^+} \right) \right)^k \int_0^{2\pi} \cos^{2k+2}(\theta) \sin^4(\theta) \left[\frac{4\beta^2}{(1+\beta)^2} \cos^2(k\theta) \right. \\ &\quad \left. + \frac{4\beta}{1+\beta} \left(1 - \frac{2\beta}{1+\beta} \right) \cos(k\theta) \cdot \frac{\sin((k+1)\theta)}{\sin(\theta)} + \left(1 - \frac{2\beta}{1+\beta} \right)^2 \frac{\sin^2((k+1)\theta)}{\sin^2(\theta)} \right] d\theta. \end{aligned}$$

As before, we will treat each term in the summand separately and use that $\int_0^{2\pi} e^{ik\theta} = 0$ to only keep track of the constant terms:

$$\begin{aligned} \cos^{2k+2}(\theta) \sin^4(\theta) \cos^2(k\theta) &= \frac{1}{2^6 4^k} \left(\frac{1}{3} (2k^3 - 9k^2 + k + 6) \right. \\ &\quad \left. - 4 \binom{2k+2}{k} + \binom{2k+2}{k-1} + 3 \binom{2k+2}{k+1} \right) + \text{terms} \\ \cos^{2k+2}(\theta) \sin^3(\theta) \cos(k\theta) \sin((k+1)\theta) &= \frac{1}{2^7 4^k} \left(-4k^2 + 6k + 4 \right. \\ &\quad \left. - 8 \binom{2k+2}{k} + 2 \binom{2k+2}{k-1} + 6 \binom{2k+2}{k+1} \right) + \text{terms} \\ \cos^{2k+2}(\theta) \sin^2(\theta) \sin^2((k+1)\theta) &= \frac{1}{2^4 4^k} \left(k - \binom{2k+2}{k} + \binom{2k+2}{k+1} \right) + \text{terms}. \end{aligned} \tag{121}$$

The result immediately follows. \square

Lemma B.3 (Average-case: Polyak Momentum). *Let $d\mu_{\text{MP}}$ be the Marčenko-Pastur law defined in (2) and P_k be the residual polynomial for Polyak's (heavy-ball) method (113). Then the following holds*

$$\begin{aligned} \int \lambda P_k^2(\lambda; \lambda^\pm) d\mu_{\text{MP}} &= \frac{(\lambda^+ - \lambda^-)^2}{32r\sigma^2} \left(\frac{\sqrt{\lambda^+} - \sqrt{\lambda^-}}{\sqrt{\lambda^+} + \sqrt{\lambda^-}} \right)^{2k} \\ &\quad \times \left[\left(\frac{(\sqrt{\lambda^+} - \sqrt{\lambda^-})^2}{\lambda^+ + \lambda^-} \right)^2 + 2 \frac{(\sqrt{\lambda^+} - \sqrt{\lambda^-})^2}{\lambda^+ + \lambda^-} \frac{2\sqrt{\lambda^-} \lambda^+}{\lambda^+ + \lambda^-} + 2 \left(\frac{2\sqrt{\lambda^-} \lambda^+}{\lambda^+ + \lambda^-} \right)^2 \right] \end{aligned} \tag{122}$$

$$\text{and } \int \lambda^2 P_k^2(\lambda) d\mu_{\text{MP}} = \frac{\lambda^+ - \lambda^-}{2} \int \lambda P_k^2(\lambda; \lambda^\pm) d\mu_{\text{MP}}.$$

Proof. In order to simplify notation, we define the following

$$\beta = \frac{\sqrt{\lambda^+} - \sqrt{\lambda^-}}{\sqrt{\lambda^+} + \sqrt{\lambda^-}}, \quad c = \frac{(\sqrt{\lambda^+} - \sqrt{\lambda^-})^2}{\lambda^+ + \lambda^-}, \quad d = \frac{2\sqrt{\lambda^-} \lambda^+}{\lambda^+ + \lambda^-}, \quad \text{and } \sigma(\lambda) = \frac{\lambda^+ + \lambda^- - 2\lambda}{\lambda^+ - \lambda^-} \tag{123}$$

$$\text{with } \tilde{P}_k^2(x) \stackrel{\text{def}}{=} \beta^{2k} [c^2 T_k^2(x) + 2cd \cdot T_k(x) U_k(x) + d^2 U_k^2(x)] \quad \text{and } \tilde{P}_k(\sigma(\lambda)) = P_k(\lambda; \lambda^\pm).$$

Under the change of variables, $u = \sigma(\lambda)$, we deduce the following for any $\ell \geq 1$

$$\begin{aligned} \int \lambda^\ell P_k^2(\lambda; \lambda^\pm) d\mu_{\text{MP}} &= \frac{1}{2\pi\sigma^2 r} \int_{\lambda^-}^{\lambda^+} \lambda^{\ell-1} \tilde{P}_k^2(\sigma(\lambda)) \sqrt{(\lambda^+ - \lambda)(\lambda - \lambda^-)} d\lambda \\ &= \frac{(\lambda^+ - \lambda^-)^2}{8\pi\sigma^2 r} \int_{-1}^1 \left(\frac{\lambda^+ + \lambda^-}{2} - \frac{\lambda^+ - \lambda^-}{2} u \right)^{\ell-1} \tilde{P}_k^2(u) \sqrt{1-u^2} du. \end{aligned} \tag{124}$$

First, we consider when $\ell = 1$. We convert this into a trig. integral using the substitution $u = \cos(\theta)$ and its nice relationship with the Chebyshev polynomials. In particular, we deduce the following

$$\int \lambda P_k^2(\lambda; \lambda^\pm) d\mu_{\text{MP}} = \frac{(\lambda^+ - \lambda^-)^2}{8\pi\sigma^2 r} \beta^{2k} \int_0^\pi [c^2 \cos^2(k\theta) + 2cd \frac{\cos(k\theta) \sin((k+1)\theta)}{\sin(\theta)} + d^2 \frac{\sin^2((k+1)\theta)}{\sin^2(\theta)}] \sin^2(\theta) d\theta.$$

Treating each term in the summand separately, we can evaluate each integral

$$\begin{aligned} \int_0^\pi \cos^2(k\theta) \sin^2(\theta) d\theta &= \frac{\pi}{4}, \quad \int_0^\pi \sin^2((k+1)\theta) d\theta = \frac{\pi}{2}, \\ \text{and } \int_0^\pi \cos(k\theta) \sin((k+1)\theta) \sin(\theta) d\theta &= \frac{\pi}{4}. \end{aligned} \quad (125)$$

The result follows. Next we consider when $\ell = 2$. A quick calculation using (124) shows that

$$\int \lambda^2 P_k^2(\lambda; \lambda^\pm) d\mu_{\text{MP}} = \frac{\lambda^+ + \lambda^-}{2} \int \lambda P_k^2(\lambda) d\mu_{\text{MP}} - \frac{(\lambda^+ - \lambda^-)^3}{16\pi\sigma^2 r} \int_{-1}^1 u \tilde{P}_k^2(u) \sqrt{1-u^2} du.$$

Since we evaluated the first integral, it suffices to analyze the second one. Again, we use a trig. substitution $u = \cos(\theta)$ and deduce the following

$$\int_{-1}^1 u \tilde{P}_k^2(u) \sqrt{1-u^2} du = \beta^{2k} \int_0^\pi \cos(\theta) \sin^2(\theta) [c^2 \cos^2(k\theta) + 2cd \frac{\cos(k\theta) \sin((k+1)\theta)}{\sin(\theta)} + d^2 \frac{\sin^2((k+1)\theta)}{\sin^2(\theta)}] d\theta.$$

Treating each term in the summand separately, we can evaluate each integral

$$\begin{aligned} \int_0^\pi \cos^2(k\theta) \sin^2(\theta) \cos(\theta) d\theta &= \int_0^\pi \sin^2((k+1)\theta) \cos(\theta) d\theta = 0, \\ \text{and } \int_0^\pi \cos(k\theta) \sin((k+1)\theta) \sin(\theta) \cos(\theta) d\theta &= 0. \end{aligned} \quad (126)$$

□

Lemma B.4 (Average-case: Nesterov accelerated method (convex)). *Let $d\mu_{\text{MP}}$ be the Marčenko-Pastur law defined in (2). Suppose the polynomials P_k are the residual polynomials for Nesterov's accelerated gradient (93). If the ratio $r = 1$, the following asymptotic holds*

$$\int \lambda P_k^2(\lambda; \lambda^\pm) d\mu_{\text{MP}} \sim \frac{(\lambda^+)^2 \log(k)}{\pi^2 \sigma^2 k^3} \quad \text{and} \quad \int \lambda^2 P_k^2(\lambda; \lambda^\pm) d\mu_{\text{MP}} \sim \frac{2(\lambda^+)^3}{\pi^2 \sigma^2} \frac{1}{k^4}. \quad (127)$$

Proof. Define the polynomial $\tilde{P}_k(u) = P_k(\lambda^+ u; \lambda^\pm)$ where the polynomial P_k satisfies Nesterov's recurrence (93). Using the change of variables $u = \frac{\lambda}{\lambda^+}$, we get the following relationship

$$\begin{aligned} \int_0^{\lambda^+} \lambda^\ell P_k^2(\lambda; \lambda^\pm) d\mu_{\text{MP}} &= \frac{(\lambda^+)^{\ell+1}}{2\pi\sigma^2} \int_0^1 u^{\ell-1} \tilde{P}_k^2(u) \sqrt{u(1-u)} du \\ &= \frac{(\lambda^+)^{\ell+1}}{2\pi\sigma^2} \int_0^1 u^{\ell-1} \frac{4J_1^2(k\sqrt{u})}{k^2 u} e^{-uk} \sqrt{u(1-u)} du \\ &\quad + \frac{(\lambda^+)^{\ell+1}}{2\pi\sigma^2} \int_0^1 u^{\ell-1} [\tilde{P}_k^2(u) - \frac{4J_1^2(k\sqrt{u})}{k^2 u} e^{-uk}] \sqrt{u(1-u)} du. \end{aligned} \quad (128)$$

In the equality above, the first integral will become the asymptotic and the second integral we bound using Corollary A.1. We start by bounding the second integral. We break this integral into three components based on the value of u

$$\left(\int_0^1 = \underbrace{\int_0^{k^{-4/3}}}_{(i)} + \underbrace{\int_{k^{-4/3}}^{\log^2(k)/k}}_{(ii)} + \underbrace{\int_{\log^2(k)/k}^1}_{(iii)} \right) u^{\ell-1} \left[\tilde{P}_k^2(u) - \frac{4J_1^2(\sqrt{uk})}{uk^2} e^{-uk} \right] \sqrt{u(1-u)} du. \quad (129)$$

For (i) in (129), we bound the integrand using Corollary A.1 such that for all $u \leq k^{-4/3}$

$$u^{\ell-1/2} \left| \tilde{P}_k^2(u) - \frac{4J_1^2(k\sqrt{u})}{uk^2} e^{-uk} \right| \leq C(u^{\ell-1/2}k^{-4/3} + k^{-13/6}u^{\ell-5/4}). \quad (130)$$

Therefore, we get that the integral (i) is bounded by

$$\begin{aligned} \int_0^{k^{-4/3}} u^{\ell-1} \left| \tilde{P}_k^2(u) - \frac{4J_1^2(\sqrt{uk})}{uk^2} e^{-uk} \right| \sqrt{u(1-u)} du &\leq C \int_0^{k^{-4/3}} u^{\ell-1/2} k^{-4/3} + k^{-13/6} u^{\ell-5/4} du \\ &= C(k^{-2-4/3\ell} + k^{-11/6-4/3\ell}) \\ &\leq C \begin{cases} k^{-19/6}, & \text{if } \ell = 1 \\ k^{-9/2}, & \text{if } \ell = 2, \end{cases} \end{aligned} \quad (131)$$

for sufficiently large k and absolute constant C .

For (ii) in (129), we bound the integrand using Corollary A.1 to get for all $k^{-4/3} \leq u \leq \log^2(k)/k$ we have

$$u^{\ell-1/2} \left| \tilde{P}_k^2(u) - \frac{4J_1^2(k\sqrt{u})}{k^2u} e^{-uk} \right| \leq C e^{-uk} (u^{\ell-3/2}k^{-8/3} + u^{\ell-7/4}k^{-17/6}). \quad (132)$$

Therefore, we get that the integral (ii) is bounded by

$$\begin{aligned} \int_{k^{-4/3}}^{\log^2(k)/k} u^{\ell-1} \left| \tilde{P}_k^2(u) - \frac{4J_1^2(k\sqrt{u})}{k^2u} e^{-uk} \right| \sqrt{u(1-u)} du \\ \leq C \int_{k^{-4/3}}^{\log^2(k)/k} e^{-uk} (u^{\ell-3/2}k^{-8/3} + u^{\ell-7/4}k^{-17/6}) du \\ (v = uk) \leq C \int_0^\infty e^{-v} v^{\ell-3/2} k^{-(\ell+13/6)} dv + C \int_0^\infty e^{-v} v^{\ell-7/4} k^{-(\ell+25/12)} dv \\ = C(k^{-(\ell+13/6)} + k^{-(\ell+25/12)}) \\ \leq C \begin{cases} k^{-37/12}, & \text{if } \ell = 1, \\ k^{-49/12}, & \text{if } \ell = 2. \end{cases} \end{aligned} \quad (133)$$

For (iii) in (129), we use a simple bound on the functions $\tilde{P}_k(u) = \frac{2e^{-ku/2}}{k\sqrt{u}} I_k(u)$ where $I_k(u)$ is defined in (97) and $J_1(k\sqrt{u})$ for $u \geq \log^2(k)/k$

$$\left| \tilde{P}_k^2(u) - \frac{4J_1^2(k\sqrt{u})}{k^2u} e^{-uk} \right| \leq \frac{4(1-u)^{k+1}}{k^2u} I_k^2(u) + e^{-uk} \frac{4J_1^2(k\sqrt{u})}{k^2u} \leq \frac{Ce^{-\log^2(k)}}{k \log^2(k)}. \quad (134)$$

In the last inequality, we used that the functions $I_k(u)$ and $J_1(k\sqrt{u})$ are bounded and $(1-u)^{k+1} \leq e^{-uk}$. Since $e^{-\log^2(k)}$ decays faster than any polynomial, we have that

$$\int_{\log^2(k)/k}^1 \left| P_k^2(u) - \frac{4J_1^2(k\sqrt{u})}{k^2u} e^{-ku} \right| u^{\ell-1} \sqrt{u(1-u)} du \leq C e^{-\log^2(k)} \quad (135)$$

for sufficiently large k and some absolute constant C .

It follows by combining (131), (133), and (135) into (129) we have for $\ell \geq 1$ the following

$$\left| \int_0^1 \left[\tilde{P}_k^2(u) - \frac{4J_1(k\sqrt{u})}{k^2u} e^{-ku} \right] u^{\ell-1} \sqrt{u(1-u)} du \right| \leq C k^{-(\ell+25/12)}. \quad (136)$$

All that remains is to integrate the Bessel part in (128) to derive the asymptotic. Here we must consider cases when $\ell = 1$ and $\ell = 2$ separately. For $\ell = 1$ using the change of variables $v = k\sqrt{u}$ we

have that

$$\begin{aligned}
\frac{(\lambda^+)^2}{2\pi\sigma^2} \int_0^1 \left(\frac{2J_1(k\sqrt{u})}{k\sqrt{u}} \right)^2 e^{-uk} \sqrt{u(1-u)} du &= \frac{2 \cdot 4(\lambda^+)^2}{2\pi\sigma^2} \frac{1}{k^3} \int_0^k J_1^2(v) e^{-v^2/k} \sqrt{1-v^2/k^2} dv \\
&\stackrel{(\sqrt{1-x} \approx 1-x \text{ for } x \text{ small})}{\sim} \frac{2 \cdot 4 \cdot (\lambda^+)^2}{2\pi\sigma^2} \frac{1}{k^3} \int_1^\infty J_1^2(v) e^{-v^2/k} dv \\
&\stackrel{(\text{Bessel asymptotic, } J_1^2(v) \sim \frac{1}{\pi v})}{\sim} \frac{2 \cdot 4 \cdot (\lambda^+)^2}{2\pi\sigma^2} \frac{1}{k^3} \int_1^\infty \frac{1}{\pi v} e^{-v^2/k} dv \\
&= \frac{2 \cdot 4 \cdot (\lambda^+)^2}{2\pi\sigma^2} \cdot \frac{1}{k^3} \cdot \frac{-1}{2\pi} \mathcal{E}_i(-\sqrt{k}),
\end{aligned}$$

where \mathcal{E}_i is the exponential integral. It is known that the exponential integral $\frac{-1}{2\pi} \mathcal{E}_i(-\sqrt{k}) \sim \frac{\log(k)}{4\pi}$. For $\ell = 2$ using the change of variables $v = uk$ we have the following

$$\begin{aligned}
\frac{(\lambda^+)^3}{2\pi\sigma^2} \int_0^1 \frac{4J_1^2(k\sqrt{u})}{k^2u} e^{-uk} u \sqrt{u(1-u)} du &= \frac{4 \cdot (\lambda^+)^3}{2\pi\sigma^2} \cdot \frac{1}{k^{7/2}} \int_0^k e^{-v} J_1^2(\sqrt{vk}) v^{1/2} \sqrt{1-\frac{v}{k}} dv \\
&\sim \frac{2 \cdot 4 \cdot (\lambda^+)^3}{2\pi^2\sigma^2} \cdot \frac{1}{k^4} \int_0^\infty \cos^2(\sqrt{vk} + C) e^{-v} dv \\
&\sim \frac{4 \cdot (\lambda^+)^3}{2\pi^2\sigma^2} \cdot \frac{1}{k^4} \int_0^\infty e^{-v} dv = \frac{4 \cdot (\lambda^+)^3}{2\pi^2\sigma^2} \cdot \frac{1}{k^4}.
\end{aligned}$$

The results follow. \square

C Adversarial model computations

In this section, we derive the adversarial guarantees for gradient descent and Nesterov's accelerated method.

Lemma C.1 (Adversarial model: Gradient descent). *Suppose Assumption 1 holds. Let λ^+ (λ^-) be the upper (lower) edge of the Marcenko Pastur distribution (2) and P_k the residual polynomial for gradient descent. Then the adversarial model for the maximal expected squared norm of the gradient is the following.*

1. If there is no noise $\tilde{R} = 0$, then

$$\lim_{d \rightarrow \infty} \max_{\mathbf{H}} \mathbb{E}[\|\nabla f(\mathbf{x}_k)\|^2] \sim \begin{cases} \frac{R^2(\lambda^+)^2}{(k+1)^2} e^{-2}, & \text{if } \lambda^- = 0 \\ R^2(\lambda^-)^2 \left(1 - \frac{\lambda^-}{\lambda^+}\right)^{2k}, & \text{if } \lambda^- > 0. \end{cases}$$

2. If $\tilde{R} > 0$, then the following holds

$$\lim_{d \rightarrow \infty} \max_{\mathbf{H}} \mathbb{E}[\|\nabla f(\mathbf{x}_k)\|^2] \sim \begin{cases} \left[\frac{R^2(\lambda^+)^2}{4} \frac{1}{k^2} + \frac{\tilde{R}^2 \lambda^+}{2} \frac{1}{k} \right] e^{-2}, & \text{if } \lambda^- = 0 \\ \left[R^2(\lambda^-)^2 + r\tilde{R}^2 \lambda^- \right] \left(1 - \frac{\lambda^-}{\lambda^+}\right)^{2k}, & \text{if } \lambda^- > 0. \end{cases}$$

Proof. Suppose we are in the noiseless setting. By a change of variables, setting $u = \lambda/\lambda^+$, the following holds

$$\max_{\lambda \in [\lambda^-, \lambda^+]} \lambda^2 \left(1 - \frac{\lambda}{\lambda^+}\right)^{2k} = \max_{u \in \left[\frac{\lambda^-}{\lambda^+}, 1\right]} (\lambda^+)^2 u^2 (1-u)^{2k}. \quad (137)$$

Taking derivatives, we get that the maximum of the RHS occurs when $u = \frac{1}{k+1}$. For sufficiently large k and $\lambda^- > 0$, the maximum lies outside the constraint of $\left[\frac{\lambda^-}{\lambda^+}, 1\right]$. Hence the maximum occurs

on the boundary, or equivalently, where $u = \frac{\lambda^-}{\lambda^+}$. The result in the setting when $\lambda^- > 0$ immediately follows from this. When $\lambda^- = 0$, then the maximum does occur at $\frac{1}{k+1}$. Plugging this value into the RHS of (137) and noting that for sufficiently large k , $(1 - 1/(k+1))^{2k} \rightarrow e^{-2}$, we get the other result for noiseless case.

Now suppose that $\tilde{R} > 0$. By a change variables, setting $u = \lambda/\lambda^+$, we have that

$$\max_{\lambda \in [\lambda^-, \lambda^+]} \left(R^2 \lambda^2 + r \tilde{R}^2 \lambda \right) \left(1 - \frac{\lambda}{\lambda^+} \right)^{2k} = \max_{u \in \left[\frac{\lambda^-}{\lambda^+}, 1 \right]} \left\{ h(u) \stackrel{\text{def}}{=} \lambda^+ (R^2 \lambda^+ u^2 + r \tilde{R}^2 u) (1-u)^{2k} \right\}. \quad (138)$$

The derivative $h'(u)$ equals 0 at $u = 1$ (local minimum) and at solutions to the quadratic

$$2R^2 \lambda^+ (k+1) u^2 + [2r \tilde{R}^2 k + r \tilde{R}^2 - 2R^2 \lambda^+] u - r \tilde{R}^2 = 0.$$

There is only one positive root of this quadratic so

$$\lambda^* = \frac{\sqrt{(2r \tilde{R}^2 k + r \tilde{R}^2 - 2R^2 \lambda^+)^2 + 8r \tilde{R}^2 R^2 \lambda^+ (k+1)} - [2r \tilde{R}^2 k + r \tilde{R}^2 - 2R^2 \lambda^+]}{4R^2 \lambda^+ (k+1)}. \quad (139)$$

We can approximate the square root using Taylor approximation to get that

$$\begin{aligned} \frac{1}{k} \sqrt{(2r \tilde{R}^2 k + r \tilde{R}^2 - 2R^2 \lambda^+)^2 + 8r \tilde{R}^2 R^2 \lambda^+ (k+1)} &= 2r \tilde{R}^2 \left[1 + \frac{r \tilde{R}^2 - 2R^2 \lambda^+}{r \tilde{R}^2 k} + \frac{2R^2 \lambda^+}{r \tilde{R}^2 k} + \mathcal{O}(k^{-2}) \right]^{1/2} \\ \text{(Taylor approximation)} &= 2r \tilde{R}^2 \left[1 + \frac{1}{2k} + \mathcal{O}(k^{-2}) \right]. \end{aligned}$$

Putting this together into (139), we get that

$$\begin{aligned} &\frac{\sqrt{(2r \tilde{R}^2 k + r \tilde{R}^2 - 2R^2 \lambda^+)^2 + 8r \tilde{R}^2 R^2 \lambda^+ (k+1)} - [2r \tilde{R}^2 k + r \tilde{R}^2 - 2R^2 \lambda^+]}{4R^2 \lambda^+ (k+1)} \\ &= \frac{\frac{r \tilde{R}^2}{k} + \mathcal{O}(k^{-2}) - \frac{r \tilde{R}^2}{k} + \frac{2R^2 \lambda^+}{k}}{4R^2 \lambda^+ + \frac{4R^2 \lambda^+}{k}} \sim \frac{1}{2k}. \end{aligned}$$

As before for sufficiently large k and $\lambda^- > 0$, the root above lies outside the constraint of $[\frac{\lambda^-}{\lambda^+}, 1]$ and so maximum occurs on the boundary, or equivalently, $u = \frac{\lambda^-}{\lambda^+}$. The result immediately follows by plugging this u into (139). When $\lambda^- = 0$, then the maximum is the root of the above quadratic which asymptotically equals $1/k$. Plugging this value into (139) and noting for sufficiently large k that $(1 - 1/k)^{2k} = e^{-2}$, we get the result for the noiseless setting. \square

Lemma C.2 (Adversarial model: Nesterov (convex)). *Suppose Assumption 1 holds. Let λ^+ be the upper edge of the Marcenko Pastur distribution (2) and P_k the residual polynomial for gradient descent. Suppose $r = 1$. Then the adversarial model for the maximal expected squared norm of the gradient is the following.*

1. If there is no noise $\tilde{R} = 0$, then

$$\lim_{d \rightarrow \infty} \max_{\mathbf{H}} \mathbb{E} [\|\nabla f(\mathbf{x}_k)\|^2] \sim \frac{8e^{-1/2}}{\sqrt{2\pi}} (\lambda^+)^2 R^2 \frac{1}{k^{7/2}}.$$

2. If $\tilde{R} > 0$, then the following holds

$$\lim_{d \rightarrow \infty} \max_{\mathbf{H}} \mathbb{E} [\|\nabla f(\mathbf{x}_k)\|^2] \sim \|J_1^2(x)\|_{\infty} (\lambda^+)^2 \tilde{R}^2 \frac{1}{k^2}.$$

Proof. First, we claim that

$$\|\lambda^2 P_k^2(\lambda; \lambda^\pm)\|_\infty k^{7/2} \rightarrow \frac{8}{\pi} (\lambda^+)^2 \max_{x \geq 0} \{x^{1/2} e^{-x}\} = \frac{8e^{-1/2}}{\sqrt{2\pi}} (\lambda^+)^2 \quad \text{as } k \rightarrow \infty. \quad (140)$$

Using the definitions in (96) and (97) for $\tilde{P}_k(u)$ and $I_k(u)$ respectively, we can write

$$P_k(\lambda^+ u; \lambda^\pm) = \tilde{P}_k(u) = \frac{2(1-u)^{(k+1)/2}}{k\sqrt{u}} I_k(u).$$

Now by a change of variables we have the following

$$\begin{aligned} \max_{\lambda \in [0, \lambda^+]} \lambda^2 P_k^2(\lambda; \lambda^\pm) k^{7/2} &= \max_{u \in [0, 1]} (\lambda^+)^2 u^2 \tilde{P}_k^2(u) k^{7/2} \\ &= \max \left\{ \max_{u \in [0, \frac{\log^2(k)}{k}]} (\lambda^+)^2 k^{7/2} u^2 \tilde{P}_k^2(u), \max_{u \in [\frac{\log^2(k)}{k}, 1]} (\lambda^+)^2 k^{7/2} u^2 \tilde{P}_k^2(u) \right\}. \end{aligned} \quad (141)$$

Let's first consider the second term in the maximum. Here we use that $|I_k(u)|$ is bounded so that

$$\begin{aligned} \max_{u \in [\frac{\log^2(k)}{k}, 1]} (\lambda^+)^2 k^{7/2} u^2 \tilde{P}_k^2(u) &= \max_{u \in [\frac{\log^2(k)}{k}, 1]} 4(\lambda^+)^2 k^{3/2} u(1-u)^{k+1} I_k^2(u) \leq \max_{u \in [\frac{\log^2(k)}{k}, 1]} C(\lambda^+)^2 k^{3/2} u(1-u)^{k+1} \\ &\stackrel{\text{def}}{=} \max_{u \in [\frac{\log^2(k)}{k}, 1]} h(u). \end{aligned}$$

The function $h(u)$ is maximized when $u = \frac{1}{k+2}$ and hence the maximum over the constrained set occurs at the endpoint $\frac{\log^2(k)}{k}$. With this value, it is immediately clear that the maximum over $u \in [\frac{\log^2(k)}{k}, 1]$ of $(\lambda^+)^2 k^{7/2} u^2 \tilde{P}_k^2(u) \rightarrow 0$. Now we consider the first term in (141). In this regime, the polynomial $\tilde{P}_k^2(u)$ behaves like the Bessel function in (32). We further break up the interval $[0, \log^2(k)/k]$ into larger or smaller than $k^{-4/3}$. When $u \in [0, k^{-4/3}]$, Corollary A.1 says there exists a constant C such that

$$\max_{u \in [0, k^{-4/3}]} (\lambda^+)^2 u^2 k^{7/2} \left| \tilde{P}_k^2(u) - \frac{4e^{-uk} J_1^2(k\sqrt{u})}{k^2 u} \right| \leq C(\lambda^+)^2 [k^{-1/2} + k^{-1/3}] \rightarrow 0.$$

Similarly when $u \in [k^{-4/3}, \log^2(k)/k]$, Corollary A.1 yields that

$$\max_{u \in [k^{-4/3}, \log^2(k)/k]} (\lambda^+)^2 u^2 k^{7/2} \left| \tilde{P}_k^2(u) - \frac{4e^{-uk} J_1^2(k\sqrt{u})}{k^2 u} \right| \leq \max_{u \in [k^{-4/3}, \log^2(k)/k]} C(\lambda^+)^2 (uk^{5/6} + u^{3/4} k^{2/3}) \rightarrow 0.$$

Using a change of variables, the relevant asymptotic to compute is

$$\max_{u \in [0, \log^2(k)/k]} 4(\lambda^+)^2 u k^{3/2} J_1^2(k\sqrt{u}) e^{-uk} = \max_{x \in [0, \log^2(k)]} 4(\lambda^+)^2 (\sqrt{xk}) J_1^2(\sqrt{kx}) \sqrt{x} e^{-x}.$$

From the uniform boundedness of the function $y \mapsto \sqrt{y} J_1^2(\sqrt{y})$, there is a constant $C > 0$ so that

$$\max_{x \in [0, \delta]} 4(\lambda^+)^2 (\sqrt{xk}) J_1^2(\sqrt{kx}) \sqrt{x} e^{-x} \leq 4(\lambda^+)^2 C \sqrt{\delta}.$$

Moreover, the Bessel function satisfies

$$J_1(z) = \sqrt{\frac{2}{\pi z}} \cos\left(z - \frac{3\pi}{4}\right) + O(z^{-3/2}),$$

and so for any fixed $\delta > 0$

$$\max_{x \in [\delta, \log^2(k)]} 4(\lambda^+)^2 (\sqrt{xk}) J_1^2(\sqrt{kx}) \sqrt{x} e^{-x} \rightarrow \max_{x \in [\delta, \infty]} \left\{ \frac{8}{\pi} (\lambda^+)^2 \sqrt{x} e^{-x} \right\}.$$

As $\delta > 0$ is arbitrary, picking it sufficiently small completes the claim.

Next we claim that the following holds

$$\max_{\lambda \in [0, \lambda^+]} k^2 \lambda P_k^2(\lambda; \lambda^\pm) = \max_{u \in [0, 1]} k^2 \lambda^+ u \tilde{P}_k^2(u) \sim \lambda^+ \|J_1(u)\|_\infty^2 \quad \text{as } k \rightarrow \infty. \quad (142)$$

A similar argument as above using that in this regime the exponential dominates the polynomial $\tilde{P}_k^2(u)$ we have

$$\max_{u \in [\log^2(k)/k, 1]} k^2 \lambda^+ u \tilde{P}_k^2(u) \rightarrow 0 \quad \text{as } k \rightarrow \infty.$$

Now we need to consider the regime where the Bessel function (32) becomes important. We use our asymptotic in Corollary A.1 to show that the polynomial is close to the Bessel, namely,

$$\begin{aligned} \max_{u \in [0, k^{-4/3}]} \lambda^+ k^2 u \left| \tilde{P}_k^2(u) - \frac{4e^{-uk} J_1^2(k\sqrt{u})}{k^2 u} \right| &\leq C\lambda^+ [k^{-2/3} + k^{-5/2}] \rightarrow 0 \quad \text{as } k \rightarrow \infty \\ \text{and } \max_{u \in [k^{-4/3}, \log^2(k)/k]} \lambda^+ k^2 u \left| \tilde{P}_k^2(u) - \frac{4e^{-uk} J_1^2(k\sqrt{u})}{k^2 u} \right| &\leq C\lambda^+ [k^{-2/3} + k^{-1/2}] \rightarrow 0 \quad \text{as } k \rightarrow \infty. \end{aligned} \quad (143)$$

It remains to compute the maximum of the Bessel equation in (32) for u in 0 to $\log^2(k)/k$. Now there exists an absolute constant \mathcal{C} so that $|J_1(x)|^2 \leq \frac{\mathcal{C}}{|x|}$ and there is also an $\eta > 0$ so that $\max_{0 \leq x \leq \frac{1}{\eta}} |J_1^2(x)| > \eta$. Moreover, the maximizer of $J_1^2(x)$ exists. By picking R sufficiently large, we see that

$$\max_{u > \frac{R}{k^2}} 4\lambda^+ e^{-ku} J_1^2(k\sqrt{u}) \leq \frac{4\mathcal{C}\lambda^+}{k^2 u} \Big|_{u=R/k^2} = \frac{4\mathcal{C}\lambda^+}{R} < \eta.$$

This means that the maximum must occur for u between 0 and R/k^2 . Hence, by picking R sufficiently large, we have the following

$$\max_{u \in [0, R/k^2]} 4\lambda^+ e^{-ku} J_1^2(k\sqrt{u}) \rightarrow \max_{x \in [0, R]} 4\lambda^+ J_1^2(x) = 4\lambda^+ \|J_1^2(x)\|_\infty.$$

Consequently, for sufficiently large k , the maximizer of

$$\begin{aligned} \max_{u \in [0, 1]} (\lambda^+)^2 R^2 u^2 \tilde{P}_k^2(u) + \lambda^+ \tilde{R}^2 r u \tilde{P}_k^2(u) &= \max_{u \in [0, \log^2(k)/k]} 4\lambda^+ e^{-ku} J_1^2(k\sqrt{u}) \\ &\rightarrow \max_{x \in [0, R]} 4\lambda^+ J_1^2(x) = 4\lambda^+ \|J_1^2(x)\|_\infty. \end{aligned} \quad (144)$$

□

D Simulation details

For Figures 1 and 7, which show that the halting time concentrates, we perform $\frac{2^{12}}{\sqrt{d}}$ training runs for each value of d . In our initial simulations we observed that the empirical standard deviation was decreasing as $d^{-1/2}$ as the model grows. Because the larger models have a significant runtime, but very little variance in the halting time, we decided to scale the number of experiments based on this estimate of the variance.

As discussed in the text, the Student's t -distribution can produce ill-conditioned matrices with large halting times. To make the numerical experiments feasible we limit the number of iterations to 1000 steps for the GD and Nesterov experiments and discard the very few runs that have not converged by this time (less than 0.1%).

For Figure 3, which shows the convergence rates, we trained 8192 models for $d = n = 4096$ steps both with ($\tilde{R}^2 = 0.05$) and without noise. The convergence rates were estimated by fitting a line to the second half of the log-log curve.

For each run we calculate the worst-case upper bound on $\|\nabla f(\mathbf{x}_k)\|^2$ at $k = n$ using Taylor et al. [2017, Conjecture 3].

$$\|\nabla f(\mathbf{x}_k)\|^2 \leq \frac{L^2 \|\mathbf{x}_0 - \mathbf{x}^*\|^2}{(k+1)^2} \stackrel{\text{def}}{=} \text{UB}_{\text{cvx}}(\|\nabla f(\mathbf{x}_k)\|^2)$$

where \mathbf{x}^* is the argmin of f calculated using the linear solver in JAX [Bradbury et al., 2018]. To visualize the difference between the worst-case and average-case rates, we draw a log-log histogram of the ratio,

$$\frac{\text{UB}_{\text{cvx}}(\|\nabla f(\mathbf{x}_k)\|^2)}{\|\nabla f(\mathbf{x}_k)\|^2}.$$

D.1 Step sizes

In this appendix section, we discuss our choices of step sizes for logistic regression and stochastic gradient descent (SGD).

D.1.1 Logistic regression

For both gradient descent and Nesterov’s accelerated method (convex) on the least squares problem we use the step size $\frac{1}{L}$. The Lipschitz constant, L , is equal to the largest eigenvalue of \mathbf{H} which can be quickly approximated using a power iteration method.

For logistic regression the Hessian is equal to $\mathbf{A}^T \mathbf{D} \mathbf{A}$, where \mathbf{D} is the Jacobian matrix of the sigmoid activation function. Hence, the Hessian’s eigenvalues are equal to those of \mathbf{H} scaled by the diagonal entries $\mathbf{D}_{ii} = \sigma((\mathbf{A}\mathbf{x})_i)(1 - \sigma((\mathbf{A}\mathbf{x})_i))$. Since the maximum value of these entries is $\frac{1}{4}$ we use a step size of $\frac{4}{L}$ for our logistic regression experiments.

D.1.2 Stochastic gradient descent (SGD)

The least squares problem (8) can be reformulated as

$$\min_{\mathbf{x} \in \mathbb{R}^d} \frac{1}{2n} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2 = \frac{1}{2n} \sum_{i=1}^n (\mathbf{a}_i \mathbf{x} - b_i)^2, \quad (145)$$

where \mathbf{a}_i is the i th row of the matrix \mathbf{A} . We perform a mini-batch SGD, *i.e.*, at each iteration we select uniformly at random a subset of the samples $b_k \subset \{1, \dots, n\}$ and perform the update

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \frac{\alpha}{|b_k|} \sum_{i \in b_k} \nabla f_i(\mathbf{x}_k). \quad (146)$$

With a slight abuse of notation, we denote by $\nabla f_i(\mathbf{x}_k) = \frac{1}{|b_k|} \sum_{i \in b_k} \nabla f_i(\mathbf{x}_k)$ the update direction and use the shorthand $b = |b_k|$ for the mini-batch size since it is fixed across iterations. The rest of this section is devoted to choosing the step size α so that the halting time is consistent across dimensions n and d . Contrary to (full) gradient descent, the step size in SGD is dimension-dependent because a typical step size in SGD uses the variance in the gradients which grows as the dimension d increases.

Over-parametrized. If $n \leq d$ we call the model over-parametrized. In this case, the strong growth condition from Schmidt and Le Roux [2013] holds. This implies that training will converge when we use a fixed step size $\frac{2}{LB^2}$ where B is defined as a constant verifying for all \mathbf{x}

$$\max_i \{\|\nabla f_i(\mathbf{x})\|\} \leq B \|\nabla f(\mathbf{x})\|. \quad (147)$$

To estimate B we will compute the expected values of $\|\nabla f_i(\mathbf{x})\|^2$ and $\|\nabla f(\mathbf{x})\|^2$. To simplify the derivation we will assume that $\tilde{\mathbf{x}}$ and $\boldsymbol{\eta}$ are normally distributed. At iterate \mathbf{x} we then have

$$\nabla f(\mathbf{x}) = \frac{1}{n} \mathbf{A}^T (\mathbf{A}(\mathbf{x} - \tilde{\mathbf{x}}) - \boldsymbol{\eta}) \quad (148)$$

$$\nabla f(\mathbf{x}) \sim N \left(\mathbf{H}\mathbf{x}, \frac{1}{d} \mathbf{H}^2 + \frac{\tilde{R}^2}{n} \mathbf{H} \right). \quad (149)$$

Hence the expected value of $\|\nabla f(\mathbf{x})\|^2$ is $\|\mathbf{H}\mathbf{x}\|^2 + \text{tr} \left(\frac{1}{d} \mathbf{H}^2 + \frac{\tilde{R}^2}{n} \mathbf{H} \right)$. Following [Bai and Silverstein \[2010, Equation 3.1.6 and Lemma 3.1\]](#) we know that for large values of n and d , the expected trace $\frac{1}{d} \text{tr} \mathbf{H} \approx 1$ and $\frac{1}{d} \text{tr} \mathbf{H}^2 \approx 1 + r$. Further, $\mathbb{E} [\|\mathbf{H}\mathbf{x}\|^2] = (1 + r)\|\mathbf{x}\|^2$ and hence

$$\begin{aligned} \mathbb{E} [\|\nabla f(\mathbf{x})\|^2] &\approx (1 + r)\|\mathbf{x}\|^2 + (1 + r) + r\tilde{R}^2 \\ &= (1 + r)(1 + \|\mathbf{x}\|^2) + r\tilde{R}^2. \end{aligned} \quad (150)$$

We can approximate the same value for a mini-batch gradient, where

$$\mathbb{E} [\|\nabla f_i(\mathbf{x})\|^2] \approx (1 + r')(1 + \|\mathbf{x}\|^2) + r'\tilde{R}^2, \quad (151)$$

for batch size b and $r' = \frac{d}{b}$. Note that $\|\mathbf{x}\| \approx 1$ because of the normalization of both the initial point and the solution, so for our experiments we set $B^2 = \frac{2+r'(2+\tilde{R}^2)}{2+r(2+\tilde{R}^2)}$.

Under-parametrized. In the under-parametrized case SGD will not converge but reach a stationary distribution around the optimum. Given a step size $\bar{\alpha} \leq \frac{1}{LM_G}$ the expected square norm of the mini-batch gradients will converge to $\bar{\alpha}LM$ where M and M_G are constants such that $\mathbb{E} [\|\nabla f_i(\mathbf{x})\|^2] \leq M + M_G\|\nabla f(\mathbf{x})\|^2$ [[Bottou et al., 2018, Theorem 4.8, Equation 4.28](#)]. We will use rough approximations of both M and M_G . In fact, we will set $M_G = B^2 = \frac{2+3r'}{2+3r}$.

To approximate M we will estimate the norm of the mini-batch gradients at the optimum for our least squares model. Set $\mathbf{x}^* = \mathbf{A}^+\mathbf{b} = \mathbf{A}^+\boldsymbol{\eta} + \tilde{\mathbf{x}}$ where \mathbf{A}^+ is the Moore-Penrose pseudoinverse. We will write the row-sampled matrix $\tilde{\mathbf{A}}$ in mini-batch SGD $\tilde{\mathbf{A}} = \mathbf{P}\mathbf{A}$ where \mathbf{P} consists of exactly b rows of the identity matrix. Note that $\mathbf{P}^T\mathbf{P}$ is idempotent.

$$\begin{aligned} \tilde{\nabla} f(\mathbf{x}^*) &= \frac{1}{b} \tilde{\mathbf{A}}^T (\tilde{\mathbf{A}}(\mathbf{A}^+\boldsymbol{\eta} + \tilde{\mathbf{x}} - \tilde{\mathbf{x}}) - \tilde{\boldsymbol{\eta}}) \\ &= \frac{1}{b} \mathbf{A}^T \mathbf{P}^T (\mathbf{P}\mathbf{A}\mathbf{A}^+\boldsymbol{\eta} - \mathbf{P}\boldsymbol{\eta}) \\ &= \frac{1}{b} \mathbf{A}^T \mathbf{P}^T \mathbf{P} (\mathbf{A}\mathbf{A}^+ - \mathbf{I}) \boldsymbol{\eta}. \end{aligned}$$

To simplify the derivation we will again assume that $\boldsymbol{\eta}$ is normally distributed and that $\tilde{R} = 1$. Thus we have

$$\tilde{\nabla} f(\mathbf{x}^*) \sim N \left(0, \frac{1}{b^2} \mathbf{A}^T \mathbf{P}^T \mathbf{P} (\mathbf{A}\mathbf{A}^+ - \mathbf{I}) (\mathbf{A}\mathbf{A}^+ - \mathbf{I})^T \mathbf{P}^T \mathbf{P} \mathbf{A} \right). \quad (152)$$

By taking the expectation of the squared norm, we derive the following

$$\begin{aligned}
\mathbb{E} \left[\|\widetilde{\nabla} f(\mathbf{x}^*)\|^2 \right] &= \frac{1}{b^2} \text{tr} \left(\mathbf{A}^T \mathbf{P}^T \mathbf{P} (\mathbf{A} \mathbf{A}^+ - \mathbf{I}) (\mathbf{A} \mathbf{A}^+ - \mathbf{I})^T \mathbf{P}^T \mathbf{P} \mathbf{A} \right) \\
&= \frac{1}{b^2} \text{tr} \left(\mathbf{A}^T \mathbf{P}^T \mathbf{P} \mathbf{A} \mathbf{A}^+ \mathbf{A}^{+T} \mathbf{A}^T \mathbf{P}^T \mathbf{P} \mathbf{A} \right) + \\
&\quad \frac{1}{b^2} \text{tr} \left(\mathbf{A}^T \mathbf{P}^T \mathbf{P} \mathbf{P}^T \mathbf{P} \mathbf{A} \right) - \\
&\quad \frac{2}{b^2} \text{tr} \left(\mathbf{A}^T \mathbf{P}^T \mathbf{P} \mathbf{A} \mathbf{A}^+ \mathbf{P}^T \mathbf{P} \mathbf{A} \right) \\
&= \text{tr} \left(\widetilde{\mathbf{H}}^2 \mathbf{A}^+ \mathbf{A}^{+T} \right) + \frac{1}{b} \text{tr} \left(\widetilde{\mathbf{H}} \right) - \frac{2}{b} \text{tr} \left(\widetilde{\mathbf{H}} \mathbf{A}^+ \mathbf{P}^T \mathbf{P} \mathbf{A} \right) \\
&= \frac{1}{n} \text{tr} \left(\widetilde{\mathbf{H}}^2 \mathbf{H}^+ \right) + \frac{1}{b} \text{tr} \left(\widetilde{\mathbf{H}} \right) - \frac{2}{b} \text{tr} \left(\widetilde{\mathbf{H}} \mathbf{A}^+ \mathbf{A}^{+T} \mathbf{A}^T \mathbf{P}^T \mathbf{P} \mathbf{A} \right) \\
&= \frac{1}{n} \text{tr} \left(\widetilde{\mathbf{H}}^2 \mathbf{H}^+ \right) + \frac{1}{b} \text{tr} \left(\widetilde{\mathbf{H}} \right) - 2 \text{tr} \left(\widetilde{\mathbf{H}}^2 \mathbf{A}^+ \mathbf{A}^{+T} \right) \\
&= \frac{1}{b} \text{tr} \left(\widetilde{\mathbf{H}} \right) - \frac{1}{n} \text{tr} \left(\widetilde{\mathbf{H}}^2 \mathbf{H}^+ \right).
\end{aligned}$$

Now, we must find an approximation of $\frac{1}{n} \text{tr} \left(\widetilde{\mathbf{H}}^2 \mathbf{H}^+ \right)$. For $b \approx n$ we have $\widetilde{\mathbf{H}}^2 \mathbf{H}^+ \approx \mathbf{H}$ whereas for $b \approx 1$ we argue that $\widetilde{\mathbf{H}}$ and \mathbf{H}^+ can be seen as independent matrices with $\mathbf{H}^+ \approx \mathbf{I}$. We can linearly interpolate between these two extremes,

$$\mathbb{E} \left[\|\widetilde{\nabla} f(\mathbf{x}^*)\|^2 \right] \approx \frac{1}{b} \text{tr} \left(\widetilde{\mathbf{H}} \right) - \frac{b}{n} \frac{1}{n} \text{tr} \left(\mathbf{H} \right) - \left(1 - \frac{b}{n} \right) \frac{1}{n} \text{tr} \left(\widetilde{\mathbf{H}}^2 \right) \quad (153)$$

$$= r' - \frac{b}{n} r - \left(1 - \frac{b}{n} \right) r(1 + r') = (1 - r)(r' - r). \quad (154)$$

Experimentally these approximations work well. Hence in our simulations we set $M = \widetilde{R}^2(1 - r)(r' - r)$.

References

- S. Arora, S. S. Du, W. Hu, Z. Li, R. R. Salakhutdinov, and R. Wang. [On exact computation with an infinitely wide neural net](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- Z. Bai and J. Silverstein. [CLT for Linear Spectral Statistics of Large-Dimensional Sample Covariance Matrices](#). *The Annals of Probability*, 32, 01 2004.
- Z. Bai and J. Silverstein. *Spectral analysis of large dimensional random matrices*, volume 20. Springer, 2010.
- Z. D. Bai and Jack W. Silverstein. [No eigenvalues outside the support of the limiting spectral distribution of large-dimensional sample covariance matrices](#). *Ann. Probab.*, 26(1):316–345, 1998.
- Z. D. Bai and Jack W. Silverstein. [Exact separation of eigenvalues of large-dimensional sample covariance matrices](#). *Ann. Probab.*, 27(3):1536–1555, 1999.
- A. Beck and M. Teboulle. [A fast iterative shrinkage-thresholding algorithm for linear inverse problems](#). *SIAM J. Imaging Sci.*, 2(1):183–202, 2009.

- G. Behrooz, S. Krishnan, and Y. Xiao. [An Investigation into Neural Net Optimization via Hessian Eigenvalue Density](#). *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2019.
- L. Benigni and S. Péché. Eigenvalue distribution of nonlinear models of random matrices. *arXiv preprint arXiv:1904.03090*, 2019.
- K. Borgwardt. *A Probabilistic Analysis of the Simplex Method*. Springer-Verlag, Berlin, Heidelberg, 1986.
- L. Bottou, F.E. Curtis, and J. Nocedal. [Optimization methods for large-scale machine learning](#). *SIAM Review*, 60(2):223–311, 2018.
- J. Bradbury, R. Frostig, P. Hawkins, M. Johnson, C. Leary, D. Maclaurin, and S. Wanderman-Milne. [JAX: composable transformations of Python+NumPy programs](#), 2018.
- L. Chizat, E. Oyallon, and F. Bach. [On lazy training in differentiable programming](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- Y.N. Dauphin, R. Pascanu, C. Gulcehre, K. Cho, S. Ganguli, and Y. Bengio. [Identifying and attacking the saddle point problem in high-dimensional non-convex optimization](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, 2014.
- P.A. Deift and T. Trogdon. [Universality in numerical computation with random data: Case Studies, Analytical Results, and Some Speculations](#). *Abel Symposia*, 13(3):221–231, 2018.
- P.A. Deift and T. Trogdon. [The conjugate gradient algorithm on well-conditioned Wishart matrices is almost deterministic](#). *Quarterly of Applied Mathematics*, 2020.
- P.A. Deift, G. Menon, S. Olver, and T. Trogdon. [Universality in numerical computations with random data](#). *Proceedings of the National Academy of Sciences*, 2014.
- R. Durrett. *Probability: theory and examples*. Cambridge University Press, fourth edition, 2010.
- A. Edelman. [Eigenvalues and condition numbers of random matrices](#). *SIAM J. Matrix Anal. Appl.*, 9(4), 1988.
- B. Fischer. *Polynomial based iteration methods for symmetric linear systems*. Springer, 1996.
- D.A. Flanders and G. Shortley. [Numerical determination of fundamental modes](#). *Journal of Applied Physics*, 1950.
- G. Golub and R. Varga. [Chebyshev semi-iterative methods, successive overrelaxation iterative methods, and second order Richardson iterative methods](#). *Numerische Mathematik*, 1961.
- S. Gunasekar, J. Lee, D. Soudry, and N. Srebro. [Characterizing Implicit Bias in Terms of Optimization Geometry](#). *Proceedings of the 35th International Conference on Machine Learning (ICML)*, 2018.
- W. Hachem, A. Hardy, and J. Najim. [Large complex correlated Wishart matrices: fluctuations and asymptotic independence at the edges](#). *Ann. Probab.*, 44(3):2264–2348, 2016.
- T. Hastie, A. Montanari, S. Rosset, and R.J. Tibshirani. [Surprises in high-dimensional ridgeless least squares interpolation](#). *arXiv preprint arXiv:1903.08560*, 2019.
- M.R. Hestenes and E. Stiefel. [Methods of conjugate gradients for solving linear systems](#). *Journal of research of the National Bureau of Standards*, 49:409–436, 1952.

- C.A.R. Hoare. [Quicksort](#). *The Computer Journal*, 5(1):10–16, 01 1962.
- A. Jacot, F. Gabriel, and C. Hongler. [Neural tangent kernel: Convergence and generalization in neural networks](#). In *Advances in neural information processing systems (NeurIPS)*, 2018.
- A. Knowles and J. Yin. [Anisotropic local laws for random matrices](#). *Probab. Theory Related Fields*, 169(1-2):257–352, 2017.
- J. Lacotte and M. Pilanci. [Optimal Randomized First-Order Methods for Least-Squares Problems](#). *arXiv preprint arXiv:2002.09488*, 2020.
- Z. Liao and R. Couillet. [The Dynamics of Learning: A Random Matrix Approach](#). *Proceedings of the 35th International Conference on Machine Learning (ICML)*, 2018.
- C. Louart, Z. Liao, and R. Couillet. [A random matrix approach to neural networks](#). *Ann. Appl. Probab.*, 28(2):1190–1248, 2018.
- V.A. Marčenko and L.A. Pastur. [Distribution of eigenvalues for some sets of random matrices](#). *Mathematics of the USSR-Sbornik*, 1967.
- C.H. Martin and M.W. Mahoney. [Implicit self-regularization in deep neural networks: Evidence from random matrix theory and implications for learning](#). *arXiv preprint arXiv:1810.01075*, 2018.
- S. Mei and A. Montanari. [The generalization error of random features regression: Precise asymptotics and double descent curve](#). *arXiv preprint arXiv:1908.05355*, 2019.
- A. Nemirovski. [Information-based complexity of convex programming](#). *Lecture Notes*, 1995.
- Y. Nesterov. *Introductory lectures on convex optimization*. Springer, 2004.
- Y. Nesterov. [How to make the gradients small](#). *Optima* 88, pages 10–11, 2012.
- R. Novak, L. Xiao, J. Lee, Y. Bahri, G. Yang, J. Hron, D.A. Abolafia, J. Pennington, and J. Sohl-Dickstein. [Bayesian deep convolutional networks with many channels are Gaussian processes](#). In *Proceedings of the 7th International Conference on Learning Representations (ICLR 2019)*, 2019.
- V. Pappayan. [The full spectrum of deepnet Hessians at scale: Dynamics with SGD Training and Sample Size](#). *arXiv preprint arXiv:1811.07062*, 2018.
- E. Paquette and T. Trogon. [Universality for the conjugate gradient and MINRES algorithms on sample covariance matrices](#). *arXiv preprint arXiv:2007.00640*, 2020.
- F. Pedregosa and D. Scieur. [Average-case Acceleration Through Spectral Density Estimation](#). In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, 2020.
- J. Pennington and P. Worah. [Nonlinear random matrix theory for deep learning](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- B.T. Polyak. [Some methods of speeding up the convergence of iteration methods](#). *USSR Computational Mathematics and Mathematical Physics*, 04, 1964.
- A. Rahimi and B. Recht. [Random features for large-scale kernel machines](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1177–1184, 2008.
- H. Rutishauser. [Theory of gradient methods](#). In *Refined Iterative Methods for Computation of the Solution and the Eigenvalues of Self-Adjoint Boundary Value Problems*, pages 24–49. Springer, 1959.

- L. Sagun, L. Bottou, and Y. LeCun. [Eigenvalues of the Hessian in Deep Learning: Singularity and Beyond](#). *arXiv preprint arXiv:1611.07476*, 2016.
- L. Sagun, T. Trogdon, and Y. LeCun. [Universal halting times in optimization and machine learning](#). *Quarterly of Applied Mathematics*, 76:1, 09 2017.
- M. Schmidt and N. Le Roux. [Fast convergence of stochastic gradient descent under a strong growth condition](#). *arXiv preprint arXiv:1308.6370*, 2013.
- S. Smale. [On the average number of steps of the simplex method of linear programming](#). *Mathematical Programming*, 27(3):241–262, 1983.
- D. Spielman and S. Teng. [Smoothed Analysis of Algorithms: Why the Simplex Algorithm Usually Takes Polynomial Time](#). *J. ACM*, 51(3):385–463, 2004.
- W. Su, S. Boyd, and E.J. Candès. [A Differential Equation for Modeling Nesterov’s Accelerated Gradient Method: Theory and Insights](#). *Journal of Machine Learning Research*, 17(153):1–43, 2016.
- T. Tao. *Topics in random matrix theory*, volume 132. American Mathematical Soc., 2012.
- T. Tao and V. Vu. [Random Matrices: the Distribution of the Smallest Singular Values](#). *Geom. Funct. Anal.*, 20(1):260–297, 2010.
- A.B. Taylor, J.M. Hendrickx, and F. Glineur. [Smooth strongly convex interpolation and exact worst-case performance of first-order methods](#). *Mathematical Programming*, 2017.
- R. Walpole and R. Myers. *Probability and Statistics for Engineers and Scientists*. Macmill, 1989.
- A.C. Wilson, R. Roelofs, M. Stern, N. Srebro, and B. Recht. [The marginal value of adaptive gradient methods in machine learning](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.